

# Foreword

It wasn't always so clear, but the Rust program *empowerment*: no matter what kind of code you to reach farther, to program with confidence in did before.

Take, for example, “systems-level” work that deal management, data representation, and concurrent programming is seen as arcane, accessible only necessary years learning to avoid its infamous p do so with caution, lest their code be open to ex

Rust breaks down these barriers by eliminating friendly, polished set of tools to help you along t “dip down” into lower-level control can do so wit customary risk of crashes or security holes, and points of a fickle toolchain. Better yet, the langua towards reliable code that is efficient in terms of

Programmers who are already working with low ambitions. For example, introducing parallelism operation: the compiler will catch the classical m more aggressive optimizations in your code with accidentally introduce crashes or vulnerabilities.

But Rust isn't limited to low-level systems progr ergonomic enough to make CLI apps, web serve quite pleasant to write — you'll find simple exam Working with Rust allows you to build skills that you can learn Rust by writing a web app, then ap Raspberry Pi.

This book fully embraces the potential of Rust to approachable text intended to help you level up also your reach and confidence as a programme learn—and welcome to the Rust community!

— Nicholas Matsakis and Aaron Turon

# Introduction

Note: This edition of the book is the same as the 1st edition, which is available in print and ebook format from [No Starch Press](#).

Welcome to *The Rust Programming Language*, an Rust programming language helps you write fast ergonomics and low-level control are often at odds. Rust challenges that conflict. Through balancing great developer experience, Rust gives you the control (such as memory usage) without all the hassle of manual control.

## Who Rust Is For

Rust is ideal for many people for a variety of reasons. It is important for many groups.

## Teams of Developers

Rust is proving to be a productive tool for collaborating developers with varying levels of systems programming experience. It's a language prone to a variety of subtle bugs, which in most cases are caught out through extensive testing and careful code review. The Rust compiler plays a gatekeeper role by refusing to compile code with bugs, including concurrency bugs. By working around these bugs, developers spend their time focusing on the program's logic.

Rust also brings contemporary developer tools t

- Cargo, the included dependency manager, compiling, and managing dependencies packages in the ecosystem.
- Rustfmt ensures a consistent coding style across the ecosystem.
- The Rust Language Server powers IDE integrations for code completion and inline documentation.

By using these and other tools in the Rust ecosystem while writing systems-level code.

## **Students**

Rust is for students and those who are interested in learning about systems programming. Using Rust, many people have learned about top-level systems programming concepts. The community is very welcoming to newcomers and encourages questions. Through efforts such as this book, the goal is to make these concepts more accessible to more people, especially students.

## **Companies**

Hundreds of companies, large and small, use Rust. Those tasks include command line tools, web servers, embedded systems, audio and video analysis and transcoding, search engines, Internet of Things applications, and parts of the Firefox web browser.

## **Open Source Developers**

Rust is for people who want to build the Rust project, develop new Rust developer tools, and libraries. We'd love to have you.

## **People Who Value Speed and Stability**

Rust is for people who crave speed and stability. Rust ensures the speed of the programs that you can create with Rust. You write them. The Rust compiler's checks ensure safety and refactoring. This is in contrast to the brittle languages that lack these checks, which developers are often afraid to refactor. These abstractions, higher-level features that compile to efficient code. Written manually, Rust endeavors to make safe code.

The Rust language hopes to support many other goals. Rust is not merely some of the biggest stakeholders. Rust aims to eliminate the trade-offs that programmers have between safety *and* productivity, speed *and* ergonomics. Rust works for you.

## **Who This Book Is For**

This book assumes that you've written code in a language that doesn't make any assumptions about which one you're using. It's broadly accessible to those from a wide variety of backgrounds. We don't spend a lot of time talking about what programming is. If you're entirely new to programming, you would want to read a book that specifically provides an introduction to programming.

## How to Use This Book

In general, this book assumes that you're reading it from front to back. Later chapters build on concepts in earlier chapters. Some chapters delve into details on a topic; we typically revisit those details in later chapters.

You'll find two kinds of chapters in this book: concept chapters and project chapters. In concept chapters, you'll learn about an aspect of Python, and you'll build small programs together, applying what you've learned. Chapters 1 and 20 are project chapters; the rest are concept chapters.

Chapter 1 explains how to install Rust, how to write programs, how to use Cargo, Rust's package manager and build system, and an introduction to the Rust language. Here we cover the basics. The following chapters will provide additional detail. If you want to learn more about Rust, Chapter 2 is the place for that. At first, you might think that Chapter 2 covers Rust features similar to those of other programming languages, but it goes straight to Chapter 4 to learn about Rust's own features. If you're a particularly meticulous learner who prefers to learn the details, the next, you might want to skip Chapter 2 and go straight to Chapter 4. Chapter 2 when you'd like to work on a project and

Chapter 5 discusses structs and methods, and C expressions, and the `if let` control flow construct make custom types in Rust.

In Chapter 7, you'll learn about Rust's module system for organizing your code and its public Application Framework. Chapter 8 discusses some common collection data structures that the standard library provides, such as vectors, strings, and hash maps. Chapter 9 discusses Rust's memory handling philosophy and techniques.

Chapter 10 digs into generics, traits, and lifetime code that applies to multiple types. Chapter 11 i Rust’s safety guarantees is necessary to ensure y Chapter 12, we’ll build our own implementation

`grep` command line tool that searches for text in files. This chapter covers the concepts we discussed in the previous chapter.

Chapter 13 explores closures and iterators: features found in many functional programming languages. In Chapter 14, we'll talk about best practices for sharing your library. Chapter 15 discusses smart pointers that the standard library provides and their functionality.

In Chapter 16, we'll walk through different modes of concurrency. We'll talk about how Rust helps you to program in multithreaded mode and look at how Rust idioms compare to object-oriented programming that might be familiar with.

Chapter 18 is a reference on patterns and patterns of expressing ideas throughout Rust programs. It covers a range of advanced topics of interest, including unsafe code, raw pointers, types, functions, and closures.

In Chapter 20, we'll complete a project in which we'll build a multithreaded web server!

Finally, some appendixes contain useful information in a reference-like format. Appendix A covers Rust's operators and symbols, Appendix C covers derived types, Appendix D covers library, and Appendix D covers macros.

There is no wrong way to read this book: if you want to learn more, you might have to jump back to earlier chapters if you need to. Read whatever works for you.

An important part of the process of learning Rust is understanding the messages the compiler displays: these will guide you. We'll provide many examples of code that doesn't compile. The message the compiler will show you in each situation. In a random example, it may not compile! Make sure you see whether the example you're trying to run is correct. We'll lead you to the correct version of any code.

## Source Code

The source files from which this book is generated are available in the `src` directory of the book's repository.

# Getting Started

Let's start your Rust journey! There's a lot to learn somewhere. In this chapter, we'll discuss:

- Installing Rust on Linux, macOS, and Windows
- Writing a program that prints `Hello, world!`
- Using `cargo`, Rust's package manager and build system

## Installation

The first step is to install Rust. We'll download `Rustup`, a tool for managing Rust versions and associated toolchains, and then use it to install Rust.

---

Note: If you prefer not to use `rustup` for some reason, see the [installation page](#) for other options.

---

The following steps install the latest stable version of Rust. Rust's stability guarantees ensure that all the examples in this book will continue to compile with newer Rust versions. There are differences between versions, because Rust often improves itself. In other words, any newer, stable version of Rust you install will work as expected with the content of this book.

---

## Command Line Notation

In this chapter and throughout the book, we'll use a notation to represent the terminal. Lines that you should enter in a terminal are preceded by the `$` character; it indicates that the following text is a command that you need to type in the terminal. Lines that don't start with `$` typically show the output of a command. Additionally, PowerShell-specific examples will be marked with `>`.

---

## Installing `rustup` on Linux or macOS

If you're using Linux or macOS, open a terminal

```
$ curl https://sh.rustup.rs -sSf | sh
```

The command downloads a script and starts the which installs the latest stable version of Rust. You'll be prompted for a password. If the install is successful, the following

```
Rust is installed now. Great!
```

If you prefer, feel free to download the script and

The installation script automatically adds Rust to your PATH. If you want to start using Rust right away in your shell, run the following command in your shell to add

```
$ source $HOME/.cargo/env
```

Alternatively, you can add the following line to your

```
$ export PATH="$HOME/.cargo/bin:$PATH"
```

Additionally, you'll need a linker of some kind. It's common when you try to compile a Rust program and get an error that it can't execute, that means a linker isn't installed or you need to install one manually. C compilers usually come with a linker. Check your platform's documentation for how to install a C compiler. Rust packages depend on C code and will need a C compiler. We're installing one now.

## Installing **rustup** on Windows

On Windows, go to <https://www.rust-lang.org/in> for installing Rust. At some point in the installation, you'll see a message explaining that you'll also need the C++ build tools. The easiest way to acquire the build tools is to install Visual Studio. The tools are in the Other Tools and Frameworks section.

The rest of this book uses commands that work on Linux and macOS. If there are specific differences, we'll explain which

## Updating and Uninstalling

After you've installed Rust via `rustup`, updating your shell, run the following update script:

```
$ rustup update
```

To uninstall Rust and `rustup`, run the following

```
$ rustup self uninstall
```

## Troubleshooting

To check whether you have Rust installed correc

```
$ rustc --version
```

You should see the version number, commit has stable version that has been released in the foll

```
rustc x.y.z (abcabcabc yyyy-mm-dd)
```

If you see this information, you have installed Ru information and you're on Windows, check that variable. If that's all correct and Rust still isn't wo you can get help. The easiest is [the #rust IRC ch](#) can access through [Mibbit](#). At that address you c nickname we call ourselves) who can help you o [Users forum](#) and [Stack Overflow](#).

## Local Documentation

The installer also includes a copy of the docume offline. Run `rustup doc` to open the local docur

Any time a type or function is provided by the st what it does or how to use it, use the applicator documentation to find out!

## Hello, World!



Now that you've installed Rust, let's write your first program. Learning a new language to write a little program that prints "Hello, world!" to the screen, so we'll do the same here!

---

Note: This book assumes basic familiarity with Rust. If you have no specific demands about your editing or tooling, or if you prefer to use an integrated development environment, feel free to use your favorite IDE. To a certain degree of Rust support; check the IDE's documentation. The Rust team has been focusing on enabling great IDE support, and it has been made rapidly on that front!

---

## Creating a Project Directory

You'll start by making a directory to store your Rust code where your code lives, but for the exercises and examples, making a *projects* directory in your home directory is a good idea.

Open a terminal and enter the following commands to create a directory for the Hello, world! project within the home directory:

For Linux and macOS, enter this:

```
$ mkdir ~/projects
$ cd ~/projects
$ mkdir hello_world
$ cd hello_world
```

For Windows CMD, enter this:

```
> mkdir "%USERPROFILE%\projects"
> cd /d "%USERPROFILE%\projects"
> mkdir hello_world
> cd hello_world
```

For Windows PowerShell, enter this:

```
> mkdir $env:USERPROFILE\projects
> cd $env:USERPROFILE\projects
> mkdir hello_world
> cd hello_world
```

## Writing and Running a Rust Program

Next, make a new source file and call it *main.rs*. extension. If you're using more than one word in the filename, separate them. For example, use *hello\_world.rs* if you want to call it *hello\_world*.

Now open the *main.rs* file you just created and edit it.

Filename: main.rs

```
fn main() {  
    println!("Hello, world!");  
}
```

Listing 1-1: A program that prints *Hello, world*

Save the file and go back to your terminal window. Enter the following commands to compile and run the file:

```
$ rustc main.rs  
$ ./main  
Hello, world!
```

On Windows, enter the command *.\main.exe* instead of *./main*:

```
> rustc main.rs  
> .\main.exe  
Hello, world!
```

Regardless of your operating system, the string *Hello, world!* should appear in the terminal. If you don't see this output, refer back to the [Installation](#) section for ways to get help.

If *Hello, world!* did print, congratulations! You've just written your first Rust program. That makes you a Rust programmer—welcome!

## Anatomy of a Rust Program

Let's review in detail what just happened in your first piece of the puzzle:

```
fn main() {  
  
}
```

These lines define a function in Rust. The `main` code that runs in every executable Rust program named `main` that has no parameters and returns they would go inside the parentheses, `()`.

Also, note that the function body is wrapped in curly braces `{}` around all function bodies. It's good style to put the same line as the function declaration, adding

At the time of this writing, an automatic formatter is in development. If you want to stick to a standard style, you will format your code in a particular style. The Rustfmt tool with the standard Rust distribution, like you read this book, it might already be installed on your system. Check the documentation for more details.

Inside the `main` function is the following code:

```
println!("Hello, world!");
```

This line does all the work in this little program: `println!("Hello, world!");`. There are four important details to notice here. First, Rust uses semicolons `;` to end a line.

Second, `println!` calls a Rust macro. If it called a normal function, it would be entered as `println` (without the `!`). We'll discuss macros in Appendix D. For now, you just need to know that `println!` is a macro instead of a normal function.

Third, you see the `"Hello, world!"` string. We pass it to `println!`, and the string is printed to the screen.

Fourth, we end the line with a semicolon `;`, which tells the compiler to move over and the next one is ready to begin. Most lines in a Rust program end with a semicolon.

## Compiling and Running Are Separate Steps

You've just run a newly created program, so let's move on to the next step.

Before running a Rust program, you must compile it. This is done by entering the `rustc` command and passing it the name of the file to compile.

```
$ rustc main.rs
```

If you have a C or C++ background, you'll notice  
After compiling successfully, Rust outputs a binary

On Linux and macOS you can see the executable  
your shell as follows:

```
$ ls
main  main.rs
```

With PowerShell on Windows, you can use `ls` a

```
> ls
```

Directory: Path:\to\the\project

Mode	LastWriteTime	
----	-----	
-a----	6/1/2018	7:31 AM
-a----	6/1/2018	7:31 AM
-a----	6/1/2018	7:31 AM

With CMD on Windows, you would enter the foll

```
> dir /B %* the /B option says to only show
main.exe
main.pdb
main.rs
```

This shows the source code file with the `.rs` exten  
Windows, but `main` on all other platforms), and,  
debugging information with the `.pdb` extension.  
`main.exe` file, like this:

```
$ ./main # or .\main.exe on Windows
```

If `main.rs` was your Hello, world! program, this li  
your terminal.

If you're more familiar with a dynamic language,  
you might not be used to compiling and running  
an *ahead-of-time compiled* language, meaning yo  
executable to someone else, and they can run it  
you give someone a `.rb`, `.py`, or `.js` file, they need  
implementation installed (respectively). But in th  
command to compile and run your program. Ev

design.

Just compiling with `rustc` is fine for simple projects, but if you'll want to manage all the options and make sure you have all the dependencies, I'll introduce you to the Cargo tool, which will help you manage your Rust projects.

## Hello, Cargo!

Cargo is Rust's build system and package manager. It helps you manage their Rust projects because Cargo handles building your code, downloading the libraries you need, and managing those libraries. (We call libraries your code needs dependencies.)

The simplest Rust programs, like the one we've built, don't have any dependencies. So if we had built the Hello, world! program, we would use the part of Cargo that handles building your code. If you have Rust programs, you'll add dependencies, and if you have dependencies, it will be much easier to do.

Because the vast majority of Rust projects use Cargo, you're using Cargo too. Cargo comes installed on most operating systems, but if you're using a different operating system, check whether Cargo is installed by running the following command in a terminal:

```
$ cargo --version
```

If you see a version number, you have it! If you see `command not found`, look at the documentation for your operating system to determine how to install Cargo separately.

## Creating a Project with Cargo

Let's create a new project using Cargo and look at the output. We'll create a Hello, world! project. Navigate back to your *project directory* (the directory you decided to store your code). Then, on any operating system, run the following commands:

```
$ cargo new hello_cargo
$ cd hello_cargo
```

The first command creates a new directory called `hello_cargo`.

project *hello\_cargo*, and Cargo creates its files in

Go into the *hello\_cargo* directory and list the files. You will see two files and one directory for us: a *Cargo.toml* file inside. It has also initialized a new Git repository.

---

Note: Git is a common version control system. You can choose to use a different version control system or no version control system at all by using the `--vcs` flag. Run `cargo new --help` to see the options.

---

Open *Cargo.toml* in your text editor of choice. It should look like Listing 1-2.

Filename: Cargo.toml

```
[package]
name = "hello_cargo"
version = "0.1.0"
authors = ["Your Name <you@example.com>"]

[dependencies]
```

Listing 1-2: Contents of *Cargo.toml* generated by Cargo

This file is in the [TOML](#) (Tom's Obvious, Minimal Language) configuration format.

The first line, `[package]`, is a section heading that indicates that the statements are configuring a package. As we add more sections, we will add other sections.

The next three lines set the configuration information for the program: the name, the version, and who wrote the program. This information comes from your environment, so if that information changes, you will need to update this information now and then save the file.

The last line, `[dependencies]`, is the start of a section for the project's dependencies. In Rust, packages of code often depend on other crates for this project, but we will not use this dependencies section then.

Now open *src/main.rs* and take a look:

Filename: src/main.rs

```
fn main() {
    println!("Hello, world!");
}
```

Cargo has generated a Hello, world! program for Listing 1-1! So far, the differences between our previous program and the one generated by Cargo are that Cargo placed the code in the *Cargo.toml* configuration file in the top directory

Cargo expects your source files to live inside the *src* directory. The *src* directory is just for README files, license information, or anything else not related to your code. Using Cargo, there's a place for everything, and everything is in its place.

If you started a project that doesn't use Cargo, you can convert it to a project that does use Cargo by running `cargo init` in the project directory and create an appropriate *Cargo.toml* file.

## Building and Running a Cargo Project

Now let's look at what's different when we build with Cargo! From your *hello\_cargo* directory, build the following command:

```
$ cargo build
   Compiling hello_cargo v0.1.0 (file:///path/to/hello_cargo)
    Finished dev [unoptimized + debuginfo] target(s) in 0.1s
```

This command creates an executable file in *target/debug/hello\_cargo.exe* (on Windows) rather than in your current directory. To run the executable with this command:

```
$ ./target/debug/hello_cargo # or .\target\debug\hello_cargo.exe
Hello, world!
```

If all goes well, *Hello, world!* should print to the terminal. The first time you run the program also causes Cargo to create a *target* directory. The *target* directory keeps track of the exact versions of dependencies used. If a project doesn't have dependencies, so the file is a bit smaller. You can also manage this file manually; Cargo manages its contents for you.

We just built a project with `cargo build` and ran it with `./target/debug/hello_cargo`, but we can also run it with `cargo run`.

and then run the resulting executable all in one

```
$ cargo run
    Finished dev [unoptimized + debuginfo]
    Running `target/debug/hello_cargo`
Hello, world!
```

Notice that this time we didn't see output indicating `hello_cargo`. Cargo figured out that the files had changed. If you had modified your source code, Cargo would have been running it, and you would have seen this output

```
$ cargo run
    Compiling hello_cargo v0.1.0 (file:///project)
    Finished dev [unoptimized + debuginfo]
    Running `target/debug/hello_cargo`
Hello, world!
```

Cargo also provides a command called `cargo check` that checks your code to make sure it compiles but doesn't produce an executable.

```
$ cargo check
    Checking hello_cargo v0.1.0 (file:///project)
    Finished dev [unoptimized + debuginfo]
```

Why would you not want an executable? Often, you use `cargo build`, because it skips the step of producing an executable and continually checking your work while writing the code. It speeds up the process! As such, many Rustaceans run `cargo check` on their program to make sure it compiles. Then they run `cargo build` to produce the executable ready to use the executable.

Let's recap what we've learned so far about Cargo:

- We can build a project using `cargo build`
- We can build and run a project in one step
- Instead of saving the result of the build in a file, Cargo stores it in the `target/debug` directory.

An additional advantage of using Cargo is that it provides specific instructions for Linux and macOS versus Windows.

## Building for Release



When your project is finally ready for release, you compile it with optimizations. This command will instead of *target/debug*. The optimizations make turning them on lengthens the time it takes for you, but there are two different profiles: one for development, which builds quickly and often, and another for building the final binary, which won't be rebuilt repeatedly and that will run as fast as possible. For benchmarking your code's running time, be sure to benchmark with the executable in *target/release*.

## Cargo as Convention

With simple projects, Cargo doesn't provide a lot of features, but it will prove its worth as your programs become more complex. If your program is composed of multiple crates, it's much easier to manage with Cargo.

Even though the `hello_cargo` project is simple, you'll use it in the rest of your Rust career. In fact, you can use the following commands to check out the project's directory, and build:

```
$ git clone someurl.com/someproject
$ cd someproject
$ cargo build
```

For more information about Cargo, check out [its documentation](#).

## Summary

You're already off to a great start on your Rust journey. Here's how to:

- Install the latest stable version of Rust using the instructions on the Rust website
- Update to a newer Rust version
- Open locally installed documentation
- Write and run a Hello, world! program using Cargo
- Create and run a new project using the `new` command

This is a great time to build a more substantial project. In the next chapter, we'll be writing Rust code. So, in Chapter 2, we'll build a game. But rather than start by learning how common programming constructs work, we'll

Chapter 3 and then return to Chapter 2.

# Programming a Guessi

Let's jump into Rust by working through a hands introduces you to a few common Rust concepts real program. You'll learn about `let`, `match`, `mod`, external crates, and more! The following chapter detail. In this chapter, you'll practice the fundam

We'll implement a classic beginner programming how it works: the program will generate a random then prompt the player to enter a guess. After a indicate whether the guess is too low or too high print a congratulatory message and exit.

## Setting Up a New Project

To set up a new project, go to the *projects* directory and make a new project using Cargo, like so:

```
$ cargo new guessing_game
$ cd guessing_game
```

The first command, `cargo new`, takes the name as the first argument. The second command changes the directory to the new project.

Look at the generated *Cargo.toml* file:

Filename: Cargo.toml

```
[package]
name = "guessing_game"
version = "0.1.0"
authors = ["Your Name <you@example.com>"]

[dependencies]
```

If the author information that Cargo obtained from the command line is not correct, you can fix that in the file and save it again.

As you saw in Chapter 1, `cargo new` generates a *src/main.rs* file. Check out the *src/main.rs* file:

Filename: src/main.rs

```
fn main() {  
    println!("Hello, world!");  
}
```

Now let's compile this "Hello, world!" program at the `cargo run` command:

```
$ cargo run  
    Compiling guessing_game v0.1.0 (file:///...)  
    Finished dev [unoptimized + debuginfo] target(s) in 0.1s  
    Running `target/debug/guessing_game`  
Hello, world!
```

The `run` command comes in handy when you need to test your code. In this game, we'll do in this game, quickly testing each iteration.

Reopen the `src/main.rs` file. You'll be writing all the code for the game.

## Processing a Guess

The first part of the guessing game program will be to get a guess from the user and check that the input is in the expected form. We'll start by getting the user input a guess. Enter the code in Listing 2-1 into `src/main.rs`.

Filename: src/main.rs

```
use std::io;  
  
fn main() {  
    println!("Guess the number!");  
  
    println!("Please input your guess.");  
  
    let mut guess = String::new();  
  
    io::stdin().read_line(&mut guess)  
        .expect("Failed to read line");  
  
    println!("You guessed: {}", guess);  
}
```

Listing 2-1: Code that gets a guess from the user

This code contains a lot of information, so let's go through it line by line.

input and then print the result as output, we need to bring the `io` library into scope. The `io` library comes from the `std` library:

```
use std::io;
```

By default, Rust brings only a few types into the `prelude`. If a type you want to use isn't in the `prelude`, you can bring it into scope explicitly with a `use` statement. Using the `std::io` library, you can read a number of useful features, including the ability to read from and write to the standard input and output streams.

As you saw in Chapter 1, the `main` function is the entry point of the program:

```
fn main() {
```

The `fn` syntax declares a new function, the parameters in parentheses, and the curly bracket, `{`, starts the function body.

As you also learned in Chapter 1, `println!` is a macro that prints to the screen:

```
println!("Guess the number!");
```

```
println!("Please input your guess.");
```

This code is printing a prompt stating what the guess is and asking the user to input a guess.

## Storing Values with Variables

Next, we'll create a place to store the user input, a variable:

```
let mut guess = String::new();
```

Now the program is getting interesting! There's a new keyword here: `let`. That this is a `let` statement, which is used to create a new variable, you can tell from the example:

```
let foo = bar;
```

This line creates a new variable named `foo` and assigns it the value of `bar`. In Rust, variables are immutable by default. For more information, see the "Variables and Mutability" section in the book.

shows how to use `mut` before the variable name

```
let foo = 5; // immutable
let mut bar = 5; // mutable
```

---

Note: The `//` syntax starts a comment that `rustc` ignores everything in comments, which is covered in Chapter 3.

---

Now you know that `let mut guess` will introduce a mutable variable. On the other side of the equal sign (`=`) is the value of the variable, the result of calling `String::new`, a function that creates a new `String`. `String` is a string type provided by the standard library, a sequence of encoded bit of text.

The `::` syntax in the `::new` line indicates that `new` is a function on the `String` type. An associated function is implemented on a type rather than on a particular instance of a `String` *method*.

This `new` function creates a new, empty string. You might think of `new` as a common name for a function that creates a new instance of a type.

To summarize, the `let mut guess = String::new()` line creates a mutable variable that is currently bound to a new, empty string.

Recall that we included the input/output function `use std::io;` on the first line of the program. Now we can use `stdin`, on `io`:

```
io::stdin().read_line(&mut guess)
    .expect("Failed to read line");
```

If we hadn't listed the `use std::io;` line at the beginning of the program, we would have written this function call as `std::io::stdin().read_line(&mut guess)`, an instance of `std::io::Stdin`, which is a type that represents the standard input for your terminal.

The next part of the code, `.read_line(&mut guess)`, calls the `read_line` method on the standard input handle to get input from the terminal. The `read_line` method takes an argument to `read_line`: `&mut guess`.

The job of `read_line` is to take whatever the user puts in at the command line, take that input from the place that into a string, so it takes that string as an argument. It needs to be mutable so the method can change user input.

The `&` indicates that this argument is a *reference*. This way, different parts of your code access one piece of data with memory multiple times. References are a complete advantage because how safe and easy it is to use references instead of those details to finish this program. For now, variables, references are immutable by default. We use `&guess` rather than `guess` to make it mutable. (Chapter 6 covers this thoroughly.)

## Handling Potential Failure with the `Result` Type

We're not quite done with this line of code. Although it's only the first part of the string, the part is this method:

```
.expect("Failed to read line");
```

When you call a method with the `.foo()` syntax, you can use trailing whitespace to help break up long lines.

```
io::stdin().read_line(&mut guess).expect('
```

However, one long line is difficult to read, so it's better to use multiple method calls. Now let's discuss what this line does.

As mentioned earlier, `read_line` puts what the user enters into the `guess` variable, but it also returns a value—in this case, a `Result` type. The `Result` type is in its standard library: a `Result` type for submodules, such as `io::Result`.

The `Result` types are *enumerations*, often referred to as *enums*. An enumeration is a type that can have a fixed set of values, and those values are called *variants*. Chapter 6 will cover enums in more detail.

For `Result`, the variants are `Ok` or `Err`. The `Ok` variant is successful, and inside `Ok` is the successfully generated value. If the operation failed, and `Err` contains information about what failed.

The purpose of these `Result` types is to encode the result of the `Result` type, like any type, have methods. `io::Result` has an `expect` method that you can use. If the result is an `Err` value, `expect` will cause the program to panic with the message that you passed as an argument to `expect`. If the result is an `Ok`, it would likely be the result of an error coming from the system. If this instance of `io::Result` is an `Ok` value that is holding and return just that value. In this case, that value is the number of bytes in what the

If you don't call `expect` , the program will compi

```
$ cargo build  
    Compiling guessing_game v0.1.0 (file:///...)  
warning: unused `std::result::Result` which  
--> src/main.rs:10:5  
  
|  
10 |         io::stdin().read_line(&mut guess).unwrap();  
   |         ^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^  
   |  
= note: #[warn(unused_must_use)] on by
```

Rust warns that you haven't used the `Result` variable, indicating that the program hasn't handled a possible error.

The right way to suppress the warning is to actually check the error code. If you just want to crash this program when a problem occurs, you can use `assert`. You'll learn about recovering from errors in Chapter 11.

## Printing Values with `println!` Placeholder

Aside from the closing curly brackets, there's one more thing we've added so far, which is the following:

```
println!("You guessed: {}", guess);
```

This line prints the string we saved the user's input in. The `%s` is a placeholder: think of `{ }` as little crab pincers that can grab and print more than one value using curly brackets: `%s %s %s`. The first value listed after the format string, the `name`, is the first value grabbed by the first pair of curly brackets, and so on. Printing multiple values in one call to `printf` is useful for debugging.

```
let x = 5;
let y = 10;

println!("x = {} and y = {}", x, y);
```

This code would print `x = 5 and y = 10`.

## Testing the First Part

Let's test the first part of the guessing game. Run

```
$ cargo run
   Compiling guessing_game v0.1.0 (file:///...
   Finished dev [unoptimized + debuginfo]
   Running `target/debug/guessing_game`
Guess the number!
Please input your guess.
6
You guessed: 6
```

At this point, the first part of the game is done: `v` and then printing it.

## Generating a Secret Number

Next, we need to generate a secret number that number should be different every time so the game Let's use a random number between 1 and 100: doesn't yet include random number functionality. Rust team does provide a `rand crate`.

## Using a Crate to Get More Functionality

Remember that a crate is a package of Rust code *binary crate*, which is an executable. The `rand crate` code intended to be used in other programs.

Cargo's use of external crates is where it really shines. uses `rand`, we need to modify the *Cargo.toml* file dependency. Open that file now and add the following



`[dependencies]` section header that Cargo creates

Filename: Cargo.toml

```
[dependencies]
```

```
rand = "0.3.14"
```

In the *Cargo.toml* file, everything that follows a header continues until another section starts. The `[dependencies]` section tells Cargo which external crates your project depends on and the versions you require. In this case, we'll specify the version specifier `0.3.14`. Cargo understands [Semantic Versioning](#), which is a standard for writing version numbers. `0.3.14` is shorthand for `^0.3.14`, which means "any version compatible with version 0.3.14."

Now, without changing any of the code, let's build the project.

```
$ cargo build
   Updating registry `https://github.com/rust-lang/crates.io/index.json`
  Downloading rand v0.3.14
  Downloading libc v0.2.14
   Compiling libc v0.2.14
   Compiling rand v0.3.14
   Compiling guessing_game v0.1.0 (file:///home/alexander/.cargo/registry/src/github.com-1ecc6299db9ec823/guessing_game-0.1.0)
   Finished dev [unoptimized + debuginfo] target(s) in 0.41s
```

Listing 2-2: The output from running `cargo build` with an external dependency

You may see different version numbers (but the same ones as the dependencies you specified, thanks to SemVer!), and the lines may be in a different order.

Now that we have an external dependency, Cargo downloads everything from the *registry*, which is a copy of the crates that people in the Rust ecosystem post their open source projects to.

After updating the registry, Cargo checks the `[dependencies]` section for any crates you don't have yet. In this case, although we only specified a dependency, Cargo also grabbed a copy of `libc` to work. After downloading the crates, Rust compiles the project with the dependencies available.

If you immediately run `cargo build` again without changing the code, you get any output aside from the `Finished` line. Cargo

and compiled the dependencies, and you haven't changed your *Cargo.toml* file. Cargo also knows that you haven't changed the code, so it doesn't recompile that either. With no changes, you'll only see two lines of output:

If you open up the *src/main.rs* file, make a trivial change, and run the build again, you'll only see two lines of output:

```
$ cargo build
   Compiling guessing_game v0.1.0 (file:///.../guessing_game)
   Finished dev [unoptimized + debuginfo] target(s) in 0.1s
```

These lines show Cargo only updates the build of the *src/main.rs* file. Your dependencies haven't changed, so Cargo uses what it has already downloaded and compiled for you. This is the goal of a reproducible build.

## Ensuring Reproducible Builds with the *Cargo.lock* File

Cargo has a mechanism that ensures you can reproduce a build of your code or anyone else builds your code: Cargo will use the exact same versions of crates you specified until you indicate otherwise. For example, if a new version *v0.3.15* of the *rand* crate comes out and it also contains a regression that will break your code, you can update the version in *Cargo.toml* to *v0.3.15*.

The answer to this problem is the *Cargo.lock* file, which is created when you run *cargo build* and is now in your *guessing\_game* project. For the first time, Cargo figures out all the dependencies and their versions, the criteria and then writes them to the *Cargo.lock* file. In the future, Cargo will see that the *Cargo.lock* file is there rather than doing all the work of figuring out the dependencies and versions, ensuring a reproducible build automatically. In other words, you can reproduce a build until you explicitly upgrade, thanks to the *Cargo.lock* file.

## Updating a Crate to Get a New Version

When you *do* want to update a crate, Cargo provides the *cargo update* command, which will ignore the *Cargo.lock* file and figure out the latest versions of the specifications in *Cargo.toml*. If that works, Cargo will update the *Cargo.lock* file.

But by default, Cargo will only look for versions of the *rand* crate that are compatible with *v0.4.0*. If the *rand* crate has released two new versions, *v0.4.1* and *v0.4.2*, you would see the following if you ran *cargo update*:

```
$ cargo update
  Updating registry `https://github.com/
  Updating rand v0.3.14 -> v0.3.15
```

At this point, you would also notice a change in y  
version of the `rand` crate you are now using is

If you wanted to use `rand` version `0.4.0` or any  
have to update the *Cargo.toml* file to look like thi

```
[dependencies]
```

```
rand = "0.4.0"
```

The next time you run `cargo build`, Cargo will i  
and reevaluate your `rand` requirements accord  
specified.

There's a lot more to say about `Cargo` and `its ec`  
Chapter 14, but for now, that's all you need to kn  
reuse libraries, so Rustaceans are able to write s  
from a number of packages.

## Generating a Random Number

Now that you've added the `rand` crate to *Cargo*.  
step is to update *src/main.rs*, as shown in Listing

Filename: *src/main.rs*

```

extern crate rand;

use std::io;
use rand::Rng;

fn main() {
    println!("Guess the number!");

    let secret_number = rand::thread_rng()

    println!("The secret number is: {}", s

    println!("Please input your guess.");

    let mut guess = String::new();

    io::stdin().read_line(&mut guess)
        .expect("Failed to read line");

    println!("You guessed: {}", guess);
}

```

Listing 2-3: Adding code to generate a random n

First, we add a line that lets Rust know we'll be u dependency. This also does the equivalent of ca anything in the `rand` crate by placing `rand:: b`

Next, we add another `use` line: `use rand::Rng` random number generators implement, and thi those methods. Chapter 10 will cover traits in de

Also, we're adding two more lines in the middle. give us the particular random number generato local to the current thread of execution and see call the `gen_range` method on the random num defined by the `Rng` trait that we brought into sc statement. The `gen_range` method takes two n a random number between them. It's inclusive c the upper bound, so we need to specify `1` and and 100.

---

Note: You won't just know which traits to use to call from a crate. Instructions for using a cr documentation. Another neat feature of Carg `cargo doc --open` command, which will build

your dependencies locally and open it in your other functionality in the `rand` crate, for example click `rand` in the sidebar on the left.

---

The second line that we added to the code prints while we're developing the program to be able to see the final version. It's not much of a game if the program starts!

Try running the program a few times:

```
$ cargo run
   Compiling guessing_game v0.1.0 (file:///.../guessing_game)
   Finished dev [unoptimized + debuginfo] target(s) in 0.1s
   Running `target/debug/guessing_game`
Guess the number!
The secret number is: 7
Please input your guess.
4
You guessed: 4
$ cargo run
   Running `target/debug/guessing_game`
Guess the number!
The secret number is: 83
Please input your guess.
5
You guessed: 5
```

You should get different random numbers, and not 1 and 100. Great job!

## Comparing the Guess to the Secret Number

Now that we have user input and a random number, we can compare the guess to the secret number. This is shown in Listing 2-4. Note that this code won't

Filename: `src/main.rs`

```

extern crate rand;

use std::io;
use std::cmp::Ordering;
use rand::Rng;

fn main() {
    // ---snip---

    println!("You guessed: {}", guess);

    match guess.cmp(&secret_number) {
        Ordering::Less => println!("Too small"),
        Ordering::Greater => println!("Too big"),
        Ordering::Equal => println!("You won!"),
    }
}

```

Listing 2-4: Handling the possible return values of `cmp`

The first new bit here is another `use` statement that brings `std::cmp::Ordering` into scope from the standard library. `Ordering` is another enum, but the variants for `Ordering` are the three outcomes that are possible when you compare two values.

Then we add five new lines at the bottom that use the `match` expression.

The `cmp` method compares two values and can return one of the three variants of `Ordering`. It takes a reference to whatever you are comparing the `guess` to the `secret_number`. The `match` expression uses the `Ordering` enum we brought into scope with the `use` statement to decide what to do next based on the value returned from the call to `cmp` with the values in question.

A `match` expression is made up of *arms*. An arm is a pattern that should be run if the value given to the beginning of the `match` expression matches that arm's pattern. Rust takes the value given to the `match` expression and matches it against the pattern in turn. The `match` construct and patterns let you express a variety of situations your code can handle them all. These features will be covered in Chapters 17 and 18, respectively.

Let's walk through an example of what would happen here. Say that the user has guessed 50 and the secret number this time is 38. When the code compares the `guess` to the `secret_number`, it will return `Ordering::Greater`, because 50 is greater than 38. The `match` expression then gets the `Ordering::Greater` value and starts checking the patterns to see which one it matches.

the first arm's pattern, `Ordering::Less`, and see if it does not match `Ordering::Less`, so it ignores the next arm. The next arm's pattern, `Ordering::Greater`! The associated code in that arm prints `Too big!` to the screen. The `match` expression is at the last arm in this scenario.

However, the code in Listing 2-4 won't compile yet.

```
$ cargo build
   Compiling guessing_game v0.1.0 (file:///...)
error[E0308]: mismatched types
  --> src/main.rs:23:21
   |
23 |         match guess.cmp(&secret_number) {
   |                         ^^^^^^^^^^^^^^^^^ expected `&std::string::String`, found integral variable
   |
   = note: expected type `&std::string::String`
   = note: found type `i32`

error: aborting due to previous error

Could not compile `guessing_game`.
```

The core of the error states that there are *mismatched types*. However, it also has type inference information. In the line `let mut guess = String::new();`, Rust was able to infer the type of `guess` as `String` and didn't make us write the type. The error message says that the `cmp` method returns an `i32` type. A few number types can have a bit number; `u32`, an unsigned 32-bit number; `i32`, a signed 32-bit number; and others. Rust defaults to an `i32`, which is the type that caused the error. The reason for the error is that Rust cannot compare a `String` to an `i32`.

Ultimately, we want to convert the `String` to a number type so we can compare it numerically. We'll be adding the following two lines to the `main` function:

Filename: src/main.rs

```
// --snip--

let mut guess = String::new();

io::stdin().read_line(&mut guess)
    .expect("Failed to read line");

let guess: u32 = guess.trim().parse()
    .expect("Please type a number!");

println!("You guessed: {}", guess);

match guess.cmp(&secret_number) {
    Ordering::Less => println!("Too small"),
    Ordering::Greater => println!("Too big"),
    Ordering::Equal => println!("You won!")
}
```

The two new lines are:

```
let guess: u32 = guess.trim().parse()
    .expect("Please type a number!");
```

We create a variable named `guess`. But wait, do we need a variable named `guess`? It does, but Rust allows us to create a new variable named `guess` with a new type. This feature is often used to convert a value from one type to another type. For example, we can use the variable name rather than forcing us to create two variables, `guess` and `guess` for example. (Chapter 3 covers shadowing.)

We bind `guess` to the expression `guess.trim()`. The expression refers to the original `guess` that was created. The `trim` method on a `String` instance will eliminate any trailing whitespace and end. Although `u32` can contain only numerical values, we can enter to satisfy `read_line`. When the user presses a key, the character is added to the string. For example, if the user types `5\n`. The `\n` represents “newline,” the `trim` method eliminates `\n`, resulting in just `5`.

The `parse` method on strings parses a string into a number. Since the `parse` method can parse a variety of number types, we specify the type we want by using `let guess: u32`. The `u32` is an unsigned, 32-bit integer. Rust has a few built-in number types. You’ll learn about other number types in Chapter 4.



in this example program and the comparison will infer that `secret_number` should be a `u32` between two values of the same type!

The call to `parse` could easily cause an error. If, `A👍%`, there would be no way to convert that to a `parse` method returns a `Result` type, much as (discussed earlier in “Handling Potential Failure”) `Result` the same way by using the `expect` method. `Result` variant because it couldn’t create a number will crash the game and print the message we give. To convert the string to a number, it will return the `Ok` will return the number that we want from the `Ok`

Let’s run the program now!

```
$ cargo run
   Compiling guessing_game v0.1.0 (file:///.../guessing_game)
   Finished dev [unoptimized + debuginfo] target(s) in 0.1s
   Running `target/guessing_game`
Guess the number!
The secret number is: 58
Please input your guess.
76
You guessed: 76
Too big!
```

Nice! Even though spaces were added before the input, the program still works. That the user guessed 76. Run the program a few more times to see different behavior with different kinds of input: guess the secret number, guess a number that is too low, and guess a number that is too high.

We have most of the game working now, but there is one more change that we need to make by adding a loop!

## Allowing Multiple Guesses with a Loop

The `loop` keyword creates an infinite loop. We’ll use it to give the user multiple chances at guessing the number:

Filename: `src/main.rs`

```

// --snip--

println!("The secret number is: {}", s

loop {
    println!("Please input your guess.

// --snip--

match guess.cmp(&secret_number) {
    Ordering::Less => println!("To
    Ordering::Greater => println!(
    Ordering::Equal => println!("
}
}
}

```

As you can see, we've moved everything into a loop and moved it forward. Be sure to indent the lines inside the loop. Now we can run the program again. Notice that there is a new prompt for the user to input a guess. The program will ask for another guess until the user can quit!

The user could always halt the program by using Ctrl+C. There's another way to escape this insatiable machine. In the next discussion in "Comparing the Guess to the Secret Number", the program will crash. The user can use Ctrl+C to quit, as shown here:

```

$ cargo run
   Compiling guessing_game v0.1.0 (file:///
   Finished dev [unoptimized + debuginfo]
   Running `target/guessing_game`
Guess the number!
The secret number is: 59
Please input your guess.
45
You guessed: 45
Too small!
Please input your guess.
60
You guessed: 60
Too big!
Please input your guess.
59
You guessed: 59
You win!
Please input your guess.
quit
thread 'main' panicked at 'Please type a r
InvalidDigit }', src/libcore/result.rs:785
note: Run with `RUST_BACKTRACE=1` for a b
error: Process didn't exit successfully: `
101)

```

Typing `quit` actually quits the game, but so will  
 However, this is suboptimal to say the least. We  
 when the correct number is guessed.

## Quitting After a Correct Guess

Let's program the game to quit when the user w

Filename: src/main.rs

```

// --snip--

    match guess.cmp(&secret_number) {
        Ordering::Less => println!("To
        Ordering::Greater => println!(
        Ordering::Equal => {
            println!("You win!");
            break;
        }
    }
}
}
}

```

Adding the `break` line after `You win!` makes the user guess the secret number correctly. Exiting the program, because the loop is the last part of `main`.

## Handling Invalid Input

To further refine the game's behavior, rather than crashing if the user inputs a non-number, let's make the game continue guessing. We can do that by altering the type of `guess` from a `String` to a `u32`:

```
let guess: u32 = match guess.trim().parse(
    Ok(num) => num,
    Err(_) => continue,
);
```

Switching from an `expect` call to a `match` expression prevents the program from crashing on an error to handling the error. The `Result` type and `Result` is an enum that has the `match` expression here, as we did with the `Order`.

If `parse` is able to successfully turn the string into a `u32` value that contains the resulting number. That `Ok(num)` pattern, and the `match` expression will just return the value produced and put inside the `Ok` value. That number is then stored in the new `guess` variable we're creating.

If `parse` is *not* able to turn the string into a `u32`, the `Err` value contains more information about the error. The `Ok(num)` pattern in the first `match` arm, but it doesn't match the second arm. The underscore, `_`, is a catchall value that will match all `Err` values, no matter what information they contain. The program will execute the second arm's code on the next iteration of the `loop` and ask for another guess. The program ignores all errors that `parse` might encounter!

Now everything in the program should work as expected.

```
$ cargo run
   Compiling guessing_game v0.1.0 (file:///
   Running `target/guessing_game`
Guess the number!
The secret number is: 61
Please input your guess.
10
You guessed: 10
Too small!
Please input your guess.
99
You guessed: 99
Too big!
Please input your guess.
foo
Please input your guess.
61
You guessed: 61
You win!
```

Awesome! With one tiny final tweak, we will finish the program. The program is still printing the secret number. That's not good for the game. Let's delete the `println!` that outputs the secret number. Here's the final code:

Filename: src/main.rs

```

extern crate rand;

use std::io;
use std::cmp::Ordering;
use rand::Rng;

fn main() {
    println!("Guess the number!");

    let secret_number = rand::thread_rng()

    loop {
        println!("Please input your guess.");

        let mut guess = String::new();

        io::stdin().read_line(&mut guess)
            .expect("Failed to read line")

        let guess: u32 = match guess.trim().parse() {
            Ok(num) => num,
            Err(_) => continue,
        };

        println!("You guessed: {}", guess)

        match guess.cmp(&secret_number) {
            Ordering::Less => println!("Too low!"),
            Ordering::Greater => println!("Too high!"),
            Ordering::Equal => {
                println!("You win!");
                break;
            }
        }
    }
}

```

Listing 2-5: Complete guessing game code

## Summary

At this point, you’ve successfully built the guessing game.

This project was a hands-on way to introduce you to Rust’s `match`, methods, associated functions, the use of `break`. In the next few chapters, you’ll learn about these concepts that most programming languages have.

functions, and shows how to use them in Rust. Chapter 5 discusses a feature that makes Rust different from other languages: traits. Chapter 6 explains how to use traits, and Chapter 7 discusses closures, which are functions that can capture and use variables from the environment in which they were created.

## Common Programming Concepts

This chapter covers concepts that appear in almost all programming languages and how they work in Rust. Many programming languages share a common core. None of the concepts presented in this chapter are unique to Rust. We discuss them in the context of Rust and explain how they work in Rust.

Specifically, you'll learn about variables, basic types, and control flow. These foundations will be in every Rust program you write and give you a strong core to start from.

---

### Keywords

The Rust language has a set of *keywords* that are reserved for the language only, much as in other languages. Keywords have special meanings, and you'll be using them to write Rust programs; a few have no current functionality but have been reserved for functionality that might be added in the future. You can find a list of the keywords in Appendix A.

---

## Variables and Mutability

As mentioned in Chapter 2, by default variables are mutable. Rust nudges you to write your code in a way that takes advantage of the easy concurrency that Rust offers. However, you can make your variables immutable. Let's explore how and why immutability and why sometimes you might want to use mutable variables.

When a variable is immutable, once a value is bound to the variable, it can't change. To illustrate this, let's generate a new project directory by using `cargo new variables`.

Then, in your new *variables* directory, open *src/main.rs* following code that won't compile just yet:

Filename: *src/main.rs*

```
fn main() {  
    let x = 5;  
    println!("The value of x is: {}", x);  
    x = 6;  
    println!("The value of x is: {}", x);  
}
```

Save and run the program using `cargo run`. You shown in this output:

```
error[E0384]: cannot assign twice to immutable  
--> src/main.rs:4:5  
   |  
2  |     let x = 5;  
   |         - first assignment to `x`  
3  |     println!("The value of x is: {}",  
4  |     x = 6;  
   |     ^^^^^ cannot assign twice to immutable variable `x`
```

This example shows how the compiler helps you. Though compiler errors can be frustrating, they do help you do what you want it to do yet; they do *not* mean you're a bad programmer! Experienced Rustaceans still get confused by compiler errors.

The error indicates that the cause of the error is `cannot assign twice to immutable variable`. The second value to the immutable `x` variable.

It's important that we get compile-time errors with variables that we previously designated as immutable because they can't change. If one part of our code operates on the assumption that a value won't change and another part of our code changes the value, that part of the code won't do what it was designed to do. This can be difficult to track down after the fact, especially if the code changes the value only *sometimes*.

In Rust, the compiler guarantees that when you declare a variable, its value really won't change. That means that when you're debugging, you'll have to keep track of how and where a value might change to reason through.

But mutability can be very useful. Variables are immutable by default, but you can make them mutable if you need to.



in Chapter 2, you can make them mutable by adding `mut` to the variable name. In addition to allowing this value to change, you can help other readers of the code by indicating that other parts of the program can change the variable value.

For example, let's change `src/main.rs` to the following:

Filename: `src/main.rs`

```
fn main() {  
    let mut x = 5;  
    println!("The value of x is: {}", x);  
    x = 6;  
    println!("The value of x is: {}", x);  
}
```

When we run the program now, we get this:

```
$ cargo run  
   Compiling variables v0.1.0 (file:///project/src)  
   Finished dev [unoptimized + debuginfo] target(s) in 0.15s  
   Running `target/debug/variables`  
The value of x is: 5  
The value of x is: 6
```

We're allowed to change the value that `x` binds to. In some cases, you'll want to make a variable mutable. It's more convenient to write than if it had only immutable values.

There are multiple trade-offs to consider in addition to mutability. For example, in cases where you're using large data structures, creating new instances of the data structure may be faster than copying and returning the same data structure. In some programming styles, creating new instances may be easier to think through, but it comes with a worthwhile penalty for gaining that clarity.

## Differences Between Variables and Constants

Being unable to change the value of a variable is a programming concept that most other languages support. Constants are values that are bound to a specific value and cannot change, but there are a few differences between variables and constants.

First, you aren't allowed to use `mut` with constants. By default—they're always immutable.

You declare constants using the `const` keyword. The type of the value *must* be annotated. We're about to see in the next section, "Data Types," so don't worry that you must always annotate the type.

Constants can be declared in any scope, including global scope. They're useful for values that many parts of code need to know.

The last difference is that constants may be set to a literal value, the result of a function call or any other value that can be computed at compile time.

Here's an example of a constant declaration where `MAX_POINTS` and its value is set to 100,000. (Rust uses all uppercase with underscores between words to improve readability in numeric literals to improve readability):

```
const MAX_POINTS: u32 = 100_000;
```

Constants are valid for the entire time a program is running. They're declared in, making them a useful choice for values that multiple parts of the program might need to know. For example, the number of points any player of a game is allowed to score.

Naming hardcoded values used throughout your program helps convey the meaning of that value to future maintainers. If you have only one place in your code you would need to be updated in the future.

## Shadowing

As you saw in the "Comparing the Guess to the Secret Number" example, you can declare a new variable with the same name as a previous variable. The new variable shadows the previous variable. Rust uses the term *shadowed* by the second, which means that the first variable is no longer visible when the variable is used. We can shadow a variable's name and repeating the use of the `let` keyword.

Filename: src/main.rs

```
fn main() {
    let x = 5;

    let x = x + 1;

    let x = x * 2;

    println!("The value of x is: {}", x);
}
```

This program first binds `x` to a value of `5`. Then taking the original value and adding `1` so the variable statement also shadows `x`, multiplying the previous value of `12`. When we run this program, it will output:

```
$ cargo run
   Compiling variables v0.1.0 (file:///proc/1000000000/target/debug/build/variables-1000000000)
   Finished dev [unoptimized + debuginfo] target(s) in 0.12s
     Running `target/debug/variables`
The value of x is: 12
```

Shadowing is different than marking a variable as mutable. It's a compile-time error if we accidentally try to reassign to the same variable using the `let` keyword. By using `let`, we can perform a few transformations on the variable to be immutable after those transformations.

The other difference between `mut` and shadowing is that shadowing creates a new variable when we use the `let` keyword, but reuses the same name. For example, we can show how many spaces they want between some text, but we really want to store that input as a number.

```
let spaces = "   ";
let spaces = spaces.len();
```

This construct is allowed because the first `spaces` variable, which is a brand-new variable, has the same name as the first one, is a number type. So to come up with different names, such as `space`, we can reuse the simpler `spaces` name. However, if we use `mut` here, we'll get a compile-time error:

```
let mut spaces = "   ";
spaces = spaces.len();
```

The error says we're not allowed to mutate a variable.

```
error[E0308]: mismatched types
--> src/main.rs:3:14
   |
3  |     spaces = spaces.len();
   |                ^^^^^^^^^^^^^ expected &str
   |
   = note: expected type `&str`
            found type `usize`
```

Now that we've explored how variables work, let's move on to data types.

## Data Types

Every value in Rust is of a certain *data type*, which is specified so it knows how to work with that data, whether it's a scalar or a compound.

Keep in mind that Rust is a *statically typed* language, meaning the types of all variables are known at compile time. The compiler doesn't want to use a variable based on the value and how we use it; it wants to use the type. For example, when we converted a `String` to a `u32` in "Comparing the Guess to the Secret Number," we added a type annotation, like this:

```
let guess: u32 = "42".parse().expect("Not a number!");
```

If we don't add the type annotation here, Rust won't compile. It means the compiler needs more information from us:

```
error[E0282]: type annotations needed
--> src/main.rs:2:9
   |
2  |     let guess = "42".parse().expect("Not a number!");
   |         ^^^^^
   |         |
   |         cannot infer type for `_`
   |         consider giving `guess` a type
```

You'll see different type annotations for other data types.

## Scalar Types

A *scalar* type represents a single value. Rust has floating-point numbers, Booleans, and character other programming languages. Let’s jump into h

### Integer Types

An *integer* is a number without a fractional comp Chapter 2, the `u32` type. This type declaration in with should be an unsigned integer (signed integ that takes up 32 bits of space. Table 3-1 shows t variant in the Signed and Unsigned columns (for declare the type of an integer value.

Table 3-1: Integer Types in Rust

Length	Signed
8-bit	<code>i8</code>
16-bit	<code>i16</code>
32-bit	<code>i32</code>
64-bit	<code>i64</code>
128-bit	<code>i128</code>
arch	<code>isize</code>

Each variant can be either signed or unsigned an *unsigned* refer to whether it’s possible for the nu other words, whether the number needs to have will only ever be positive and can therefore be re It’s like writing numbers on paper: when the sign plus sign or a minus sign; however, when it’s saf it’s shown with no sign. Signed numbers are stor representation (if you’re unsure what this is, you explanation is outside the scope of this book).

Each signed variant can store numbers from  $-(2^{n-1})$  the number of bits that variant uses. So an `i8` c which equals -128 to 127. Unsigned variants can `u8` can store numbers from 0 to  $2^8 - 1$ , which ec

Additionally, the `isize` and `usize` types depen

program is running on: 64 bits if you're on a 64-bit system, and 32 bits if you're on a 32-bit architecture.

You can write integer literals in any of the forms shown in Table 3-2. Number literals except the byte literal allow a `u` or `i` visual separator, such as `1_000`.

Table 3-2: Integer Literals in Rust

Number literals	
Decimal	<code>9</code>
Hex	<code>0x0</code>
Octal	<code>0o0</code>
Binary	<code>0b0</code>
Byte ( <code>u8</code> only)	<code>b</code>

So how do you know which type of integer to use? Rust has some generally good choices, and integer types default to the fastest, even on 64-bit systems. The primary situation where `usize` is when indexing some sort of collection.

### Integer Overflow

Let's say that you have a `u8`, which can hold values from 0 to 255. What happens if you try to change it to `256`? This is called integer overflow. Rust has some interesting rules around this behavior. `Debug` will check for this kind of issue and will cause your program to panic. Rust uses `panic!` when a program exits with an error. `unwrap()` will also cause a panic.

In release builds, Rust does not check for overflow. This behavior is called "two's complement wrapping." In short, `256` wraps around to `0`, `257` wraps around to `1`, etc. Relying on overflow is considered an error, and `Debug` will warn you. If you want this behavior explicitly, the standard library has `wrapping` methods for it explicitly.

### Floating-Point Types

Rust also has two primitive types for *floating-point* numbers. Rust's floating-point types are `f32` and `f64`, 32 bits and 64 bits in size, respectively. The default type is `f64`. `f32` is the same speed as `f32` but is capable of more precision.

Here's an example that shows floating-point numbers

Filename: src/main.rs

```
fn main() {  
    let x = 2.0; // f64  
  
    let y: f32 = 3.0; // f32  
}
```

Floating-point numbers are represented according to their type. `f32` is a single-precision float, and `f64` has double-precision.

## Numeric Operations

Rust supports the basic mathematical operation types: addition, subtraction, multiplication, division, and remainder. The following code shows how you'd use each one in a `let` statement.

Filename: src/main.rs

```
fn main() {  
    // addition  
    let sum = 5 + 10;  
  
    // subtraction  
    let difference = 95.5 - 4.3;  
  
    // multiplication  
    let product = 4 * 30;  
  
    // division  
    let quotient = 56.7 / 32.2;  
  
    // remainder  
    let remainder = 43 % 5;  
}
```

Each expression in these statements uses a mathematical operator to combine one or more single values, which is then bound to a variable. `+`, `-`, `*`, `/`, and `%` are the operators that Rust provides.

## The Boolean Type

As in most other programming languages, a Boolean type has two possible values: `true` and `false`. The Boolean type in Rust is `bool`. Here's an example:

Filename: src/main.rs

```
fn main() {  
    let t = true;  
  
    let f: bool = false; // with explicit  
}
```

The main way to consume Boolean values is through the `if` expression. We'll cover how `if` expressions work in the next section.

Booleans are one byte in size.

## The Character Type

So far we've worked only with numbers, but Rust's `char` type is the language's most primitive alphabetic one way to use it. (Note that the `char` literal is single-quoted, as opposed to string literals, which use double quotes.)

Filename: src/main.rs

```
fn main() {  
    let c = 'z';  
    let z = 'Z';  
    let heart_eyed_cat = '😻';  
}
```

Rust's `char` type represents a Unicode Scalar Value, which is a lot more than just ASCII. Accented letters; Chinese characters; emoji; and zero-width spaces are all valid `char` values. The range from `U+0000` to `U+D7FF` and `U+E000` to `U+10FFFF` are all valid "character" ranges in Unicode, so you may not match up with what a `char` is in other languages. We'll detail in "Strings" in Chapter 8.

## Compound Types

*Compound types* can group multiple values into a single type. There are two compound types: tuples and arrays.

### The Tuple Type



A tuple is a general way of grouping together so variety of types into one compound type. Tuples they cannot grow or shrink in size.

We create a tuple by writing a comma-separated Each position in the tuple has a type, and the type don't have to be the same. We've added optional

Filename: src/main.rs

```
fn main() {  
    let tup: (i32, f64, u8) = (500, 6.4, 1)  
}
```

The variable `tup` binds to the entire tuple, because it's a compound element. To get the individual values matching to destructure a tuple value, like this:

Filename: src/main.rs

```
fn main() {  
    let tup = (500, 6.4, 1);  
  
    let (x, y, z) = tup;  
  
    println!("The value of y is: {}", y);  
}
```

This program first creates a tuple and binds it to a variable with `let` to take `tup` and turn it into a pattern. This is called *destructuring*, because it breaks the tuple into its parts. In the program, the program prints the value of `y`, which is `6.4`.

In addition to destructuring through pattern matching, you can also access tuple elements directly by using a period ( `.` ) followed by the index. For example:

Filename: src/main.rs

```
fn main() {
    let x: (i32, f64, u8) = (500, 6.4, 1);

    let five_hundred = x.0;

    let six_point_four = x.1;

    let one = x.2;
}
```

This program creates a tuple, `x`, and then make using their index. As with most programming lar

## The Array Type

Another way to have a collection of multiple val every element of an array must have the same t arrays in some other languages because arrays

In Rust, the values going into an array are writte square brackets:

Filename: src/main.rs

```
fn main() {
    let a = [1, 2, 3, 4, 5];
}
```

Arrays are useful when you want your data alloc heap (we will discuss the stack and the heap mo ensure you always have a fixed number of elem vector type, though. A vector is a similar collectio library that *is* allowed to grow or shrink in size. If array or a vector, you should probably use a vec more detail.

An example of when you might want to use an a program that needs to know the names of the n that such a program will need to add or remove because you know it will always contain 12 item:

```
let months = ["January", "February", "Marc
"July",
              "August", "September", "Octo
```

Arrays have an interesting type; it looks like this:

```
let a: [i32; 5] = [1, 2, 3, 4, 5];
```

First, there's square brackets; they look like the `[]`. There's two pieces of information, separated by a semicolon. The first is the type of each element of the array. Since all elements have the same type, it's only specified once. After the semicolon, there's a number that represents the size of the array. Since an array has a fixed size, this number is always the same. If you try to modify elements, it cannot grow or shrink.

### Accessing Array Elements

An array is a single chunk of memory allocated on the stack. You can access elements of an array using indexing, like this:

Filename: src/main.rs

```
fn main() {  
    let a = [1, 2, 3, 4, 5];  
  
    let first = a[0];  
    let second = a[1];  
}
```

In this example, the variable named `first` will get the value at index `[0]` in the array. The variable named `second` will get the value from index `[1]` in the array.

### Invalid Array Element Access

What happens if you try to access an element of an array? Say you change the example to the following, which will result in an error when it runs:

Filename: src/main.rs

```
fn main() {  
    let a = [1, 2, 3, 4, 5];  
    let index = 10;  
  
    let element = a[index];  
  
    println!("The value of element is: {}")  
}
```

Running this code using `cargo run` produces the

```
$ cargo run
   Compiling arrays v0.1.0 (file:///project/arrays)
   Finished dev [unoptimized + debuginfo] target(s) in 0.15s
   Running `target/debug/arrays`
thread '<main>' panicked at 'index out of bounds: the index is 10, which is greater than the length 10', src/main.rs:6:10
note: Run with `RUST_BACKTRACE=1` for a backtrace.
```

The compilation didn't produce any errors, but the program panicked. This is an error and didn't exit successfully. When you attempt to access an index that is out of bounds, Rust will check that the index you've specified is greater than the length, Rust will panic.

This is the first example of Rust's safety principle. In most other programming languages, this kind of check is not done, and with that, invalid memory can be accessed. Rust protects you by panicking immediately instead of allowing the memory to be accessed. We'll discuss more of Rust's error handling.

## Functions

Functions are pervasive in Rust code. You've already seen functions in the language: the `main` function, which is the entry point for all programs. You've also seen the `fn` keyword, which is used to define functions.

Rust code uses *snake case* as the conventional style for naming functions. In snake case, all letters are lowercase and underscores are used to separate words. In the program that contains an example function definition, the filename is `src/main.rs`.

Filename: `src/main.rs`

```
fn main() {
    println!("Hello, world!");

    another_function();
}

fn another_function() {
    println!("Another function.");
}
```

Function definitions in Rust start with `fn` and have a function name. The curly brackets tell the compiler where the function starts and ends.

We can call any function we've defined by entering its name in parentheses. Because `another_function` is defined from inside the `main` function. Note that we defined the `main` function in the source code; we could have defined it elsewhere. Rust doesn't care where you define your functions, or

Let's start a new binary project named `functions` and copy the `another_function` example in `src/main.rs` and run it. The output:

```
$ cargo run
   Compiling functions v0.1.0 (file:///proje.../functions)
   Finished dev [unoptimized + debuginfo] target(s) in 0.15s
  Running `target/debug/functions`
Hello, world!
Another function.
```

The lines execute in the order in which they appear in the source code. First, the "Hello, world!" message prints, and then `another_function` is printed.

## Function Parameters

Functions can also be defined to have *parameters* as part of a function's signature. When a function is called, you pass in concrete values for those parameters. Technically, these are called *arguments*, but in casual conversation, people tend to use *argument* interchangeably for either the variable name or the concrete values passed in when you call a function.

The following rewritten version of `another_function` is like in Rust:

Filename: `src/main.rs`

```
fn main() {
    another_function(5);
}

fn another_function(x: i32) {
    println!("The value of x is: {}", x);
}
```

Try running this program; you should get the fol

```
$ cargo run
Compiling functions v0.1.0 (file:///pro
Finished dev [unoptimized + debuginfo]
Running `target/debug/functions`
The value of x is: 5
```

The declaration of `another_function` has one parameter specified as `i32`. When `5` is passed to `another_function` where the pair of curly brackets were in the function signature.

In function signatures, you *must* declare the type of each parameter. This is a deliberate decision in Rust's design: requiring type annotations means the compiler almost never needs you to figure out what you mean.

When you want a function to have multiple parameters, you declare them with commas, like this:

Filename: `src/main.rs`

```
fn main() {
    another_function(5, 6);
}

fn another_function(x: i32, y: i32) {
    println!("The value of x is: {}", x);
    println!("The value of y is: {}", y);
}
```

This example creates a function with two parameters. The function then prints the values in both of its parameters. Not all parameters need to be the same type, as in this example.

Let's try running this code. Replace the program in the `src/main.rs` file with the preceding example and run it.

```
$ cargo run
   Compiling functions v0.1.0 (file:///prc
   Finished dev [unoptimized + debuginfo]
   Running `target/debug/functions`
The value of x is: 5
The value of y is: 6
```

Because we called the function with `5` as the value for `y`, the two strings are printed with the

## Function Bodies

Function bodies are made up of a series of state expression. So far, we've only covered functions you have seen an expression as part of state based language, this is an important distinction have the same distinctions, so let's look at what how their differences affect the bodies of function

## Statements and Expressions

We've actually already used statements and expressions that perform some action and do not return a value resulting value. Let's look at some examples.

Creating a variable and assigning a value to it with Listing 3-1, `let y = 6;` is a statement:

Filename: src/main.rs

```
fn main() {
    let y = 6;
}
```

Listing 3-1: A `main` function declaration containing

Function definitions are also statements; the entire function is in itself.

Statements do not return values. Therefore, you can't assign the result of a statement to another variable, as the following code tries to do

Filename: src/main.rs

```
fn main() {
    let x = (let y = 6);
}
```

When you run this program, the error you'll get is

```
$ cargo run
   Compiling functions v0.1.0 (file:///proje
error: expected expression, found statement
--> src/main.rs:2:14
   |
2  |     let x = (let y = 6);
   |                ^^^
   = note: variable declaration using `let`
```

The `let y = 6` statement does not return a value to bind to. This is different from what happens in C where the assignment returns the value of the assigned variable. In C, you can write `x = y = 6` and have both `x` and `y` have the value 6. In Rust, you cannot do this.

Expressions evaluate to something and make up the building blocks of what you'll write in Rust. Consider a simple math operation: `6 + 5`. This expression evaluates to the value `11`. In Listing 3-1, the `6` in the statement `let y = 6;` is an expression that evaluates to the value `6`. Calling a function is an expression. Call blocks that we use to create new scopes, `{ }`, are also expressions.

Filename: src/main.rs

```
fn main() {
    let x = 5;

    let y = {
        let x = 3;
        x + 1
    };

    println!("The value of y is: {}", y);
}
```

This expression:



```
{
    let x = 3;
    x + 1
}
```

is a block that, in this case, evaluates to `4`. That `let` statement. Note the `x + 1` line without a `;` most of the lines you've seen so far. Expressions you add a semicolon to the end of an expression will then not return a value. Keep this in mind as we move on to expressions next.

## Functions with Return Values

Functions can return values to the code that calls them, but we do declare their type after an arrow (`->`). The return type of a function is synonymous with the value of the final expression of a function. You can return early from a function by specifying a value, but most functions return the value of the final expression. Here's an example of a function that returns a value:

Filename: src/main.rs

```
fn five() -> i32 {
    5
}

fn main() {
    let x = five();

    println!("The value of x is: {}", x);
}
```

There are no function calls, macros, or even `let` statements in `five`, just the number `5` by itself. That's a perfectly valid expression. The function's return type is specified, too, as `-> i32`. The code should look like this:

```
$ cargo run
   Compiling functions v0.1.0 (file:///proc/12345/src)
   Finished dev [unoptimized + debuginfo] target(s) in 0.1s
   Running `target/debug/functions`
The value of x is: 5
```

The `5` in `five` is the function's return value, which is then assigned to `x` in `main`.

Let's examine this in more detail. There are two `let x = five();` shows that we're using the variable. Because the function `five` returns a `5` following:

```
let x = 5;
```

Second, the `five` function has no parameters a value, but the body of the function is a lonely `5` expression whose value we want to return.

Let's look at another example:

Filename: src/main.rs

```
fn main() {  
    let x = plus_one(5);  
  
    println!("The value of x is: {}", x);  
}  
  
fn plus_one(x: i32) -> i32 {  
    x + 1  
}
```

Running this code will print `The value of x is` the end of the line containing `x + 1`, changing i we'll get an error.

Filename: src/main.rs

```
fn main() {  
    let x = plus_one(5);  
  
    println!("The value of x is: {}", x);  
}  
  
fn plus_one(x: i32) -> i32 {  
    x + 1;  
}
```

Compiling this code produces an error, as follow

```

error[E0308]: mismatched types
--> src/main.rs:7:28
   |
 7 |     fn plus_one(x: i32) -> i32 {
   |                           ^
 8 |         x + 1;
   |         - help: consider removing the trailing semicolon
 9 |     }
   |     _^ expected i32, found ()
   |
   = note: expected type `i32`
           found type `()`

```

The main error message, “mismatched types,” re  
The definition of the function `plus_one` says tha  
statements don’t evaluate to a value, which is ex  
Therefore, nothing is returned, which contradict  
an error. In this output, Rust provides a message  
suggests removing the semicolon, which would i

## Comments

All programmers strive to make their code easy  
explanation is warranted. In these cases, progra  
their source code that the compiler will ignore b  
may find useful.

Here’s a simple comment:

```
// Hello, world.
```

In Rust, comments must start with two slashes a  
For comments that extend beyond a single line,  
line, like this:

```
// So we're doing something complicated here  

// multiple lines of comments to do it! We  

// explain what's going on.
```

Comments can also be placed at the end of lines

Filename: src/main.rs

```
fn main() {
    let lucky_number = 7; // I'm feeling i
}
```

But you'll more often see them used in this form line above the code it's annotating:

Filename: src/main.rs

```
fn main() {
    // I'm feeling lucky today.
    let lucky_number = 7;
}
```

Rust also has another kind of comment, documented in Chapter 14.

## Control Flow

Deciding whether or not to run some code depends on deciding to run some code repeatedly while a condition is true. These are the most common control blocks in most programming languages. The most common control blocks that control the flow of execution of Rust code are `if` expressions and `loop` expressions.

### `if` Expressions

An `if` expression allows you to branch your code. You provide a condition and then state, "If this condition is met, run this block of code. If the condition is not met, do not run this block of code."

Create a new project called *branches* in your *projects* directory. In the *src/main.rs* file, input the following code:

Filename: src/main.rs

```
fn main() {
    let number = 3;

    if number < 5 {
        println!("condition was true");
    } else {
        println!("condition was false");
    }
}
```

All `if` expressions start with the keyword `if`, and in this case, the condition checks whether or not the value is less than 5. The block of code we want to execute if the condition is true is immediately after the condition inside curly braces. Sometimes the conditions in `if` expressions are sometimes `match` expressions that we discussed in the “Conditional Number” section of Chapter 2.

Optionally, we can also include an `else` expression to give the program an alternative block of code to evaluate to false. If you don’t provide an `else` expression, the program will just skip the `if` block and move on.

Try running this code; you should see the following output:

```
$ cargo run
   Compiling branches v0.1.0 (file:///project/target/debug/branches)
   Finished dev [unoptimized + debuginfo] target(s) in 0.15s
   Running `target/debug/branches`
condition was true
```

Let’s try changing the value of `number` to a value greater than 5 and see what happens:

```
let number = 7;
```

Run the program again, and look at the output:

```
$ cargo run
   Compiling branches v0.1.0 (file:///project/target/debug/branches)
   Finished dev [unoptimized + debuginfo] target(s) in 0.15s
   Running `target/debug/branches`
condition was false
```

It’s also worth noting that the condition in this case isn’t a `bool`, we’ll get an error. For example:

Filename: src/main.rs

```
fn main() {  
    let number = 3;  
  
    if number {  
        println!("number was three");  
    }  
}
```

The `if` condition evaluates to a value of `3` this

```
error[E0308]: mismatched types  
--> src/main.rs:4:8  
   |  
4 |         if number {  
   |         ^^^^^^ expected bool, found int  
   |  
   = note: expected type `bool`  
           found type `{integer}`
```

The error indicates that Rust expected a `bool` but such as Ruby and JavaScript, Rust will not automatically convert integer types to a Boolean. You must be explicit and always use a Boolean condition. If we want the `if` code block to run only when the condition is true, for example, we can change the `if` expression to:

Filename: src/main.rs

```
fn main() {  
    let number = 3;  
  
    if number != 0 {  
        println!("number was something other than zero");  
    }  
}
```

Running this code will print `number was something other than zero`

## Handling Multiple Conditions with `else if`

You can have multiple conditions by combining them with `else if` expression. For example:

Filename: src/main.rs

```
fn main() {
    let number = 6;

    if number % 4 == 0 {
        println!("number is divisible by 4")
    } else if number % 3 == 0 {
        println!("number is divisible by 3")
    } else if number % 2 == 0 {
        println!("number is divisible by 2")
    } else {
        println!("number is not divisible by 4, 3, or 2")
    }
}
```

This program has four possible paths it can take following output:

```
$ cargo run
   Compiling branches v0.1.0 (file:///proj/target/debug/branches)
   Finished dev [unoptimized + debuginfo] target(s) in 0.14s
   Running `target/debug/branches`
number is divisible by 3
```

When this program executes, it checks each `if` first body for which the condition holds true. No we don't see the output `number is divisible by 4, 3, or 2` text Rust only executes the block for the first true condition and doesn't even check the rest.

Using too many `else if` expressions can clutter one, you might want to refactor your code. Chapter 3 introduces a branching construct called `match` for these cases.

## Using `if` in a `let` Statement

Because `if` is an expression, we can use it on the right-hand side of a `let` statement. Listing 3-2:

Filename: `src/main.rs`

```
fn main() {
    let condition = true;
    let number = if condition {
        5
    } else {
        6
    };

    println!("The value of number is: {}",
}

```

Listing 3-2: Assigning the result of an `if` expres

The `number` variable will be bound to a value ba expression. Run this code to see what happens:

```
$ cargo run
   Compiling branches v0.1.0 (file:///proj
   Finished dev [unoptimized + debuginfo]
   Running `target/debug/branches`
The value of number is: 5

```

Remember that blocks of code evaluate to the la by themselves are also expressions. In this case, expression depends on which block of code exe have the potential to be results from each arm c Listing 3-2, the results of both the `if` arm and t the types are mismatched, as in the following ex

Filename: src/main.rs

```
fn main() {
    let condition = true;

    let number = if condition {
        5
    } else {
        "six"
    };

    println!("The value of number is: {}",
}

```

When we try to compile this code, we'll get an er value types that are incompatible, and Rust indic problem in the program:



```

error[E0308]: if and else have incompatible types
--> src/main.rs:4:18
   |
 4 |         let number = if condition {
   |                        ^
   |                        5
   |        } else {
   |        "six"
   |        };
   |        ^ expected integral variable, found `&str`
   = note: expected type `{integer}`
           found type `&str`

```

The expression in the `if` block evaluates to an `integer`, while the expression in the `else` block evaluates to a string. This won't work because `number` has a single type. Rust needs to know at compile time what type `number` is, so it can verify at compile time that `number` is an `integer`. Rust wouldn't be able to do that if the type was determined at runtime; the compiler would be more complicated about the code if it had to keep track of multiple types.

## Repetition with Loops

It's often useful to execute a block of code more than once. Rust provides several *loops*. A loop runs through the code block and then starts immediately back at the beginning. We'll make a new project called *loops*.

Rust has three kinds of loops: `loop`, `while`, and `while let`.

### Repeating Code with `loop`

The `loop` keyword tells Rust to execute a block of code over and over until you explicitly tell it to stop.

As an example, change the `src/main.rs` file in your new project.

Filename: `src/main.rs`

```

fn main() {
    loop {
        println!("again!");
    }
}

```

When we run this program, we'll see `again!` printed over and over again. If we stop the program manually. Most terminals will allow you to `Ctrl-C` to halt a program that is stuck in a continual loop. (

```
$ cargo run
   Compiling loops v0.1.0 (file:///project/target/debug/loops)
   Finished dev [unoptimized + debuginfo] target(s) in 0.15s
   Running `target/debug/loops`
again!
again!
again!
again!
^Cagain!
```

The symbol `^C` represents where you pressed `Ctrl-C`. The word `again!` printed after the `^C`, depending on the terminal, when it received the halt signal.

Fortunately, Rust provides another, more reliable way to exit a loop. We can place the `break` keyword within the loop to tell the loop to stop. Recall that we did this in the `guessing game` section of Chapter 2 to exit the program when the user guessed the correct number.

## Returning from loops

One of the uses of a `loop` is to retry an operation if a thread completed its job. However, you might want to return from the loop to the rest of your code. If you add it to the loop, it will be returned by the broken loop.

```
fn main() {
    let mut counter = 0;

    let result = loop {
        counter += 1;

        if counter == 10 {
            break counter * 2;
        }
    };

    assert_eq!(result, 20);
}
```

## Conditional Loops with `while`

It's often useful for a program to evaluate a condition. If the condition is true, the loop runs. When the condition becomes false, the loop calls `break`, stopping the loop. This loop type is a combination of `loop`, `if`, `else`, and `break`; you can use it as you'd like.

However, this pattern is so common that Rust has called it a `while` loop. Listing 3-3 uses `while`: the loop runs down each time, and then, after the loop, it prints the message.

Filename: src/main.rs

```
fn main() {  
    let mut number = 3;  
  
    while number != 0 {  
        println!("{}", number);  
  
        number = number - 1;  
    }  
  
    println!("LIFTOFF!!!");  
}
```

Listing 3-3: Using a `while` loop to run code while a condition is true

This construct eliminates a lot of nesting that we would have with `if`, `else`, and `break`, and it's clearer. While a condition is true, the loop runs; otherwise, it exits the loop.

## Looping Through a Collection with `for`

You could use the `while` construct to loop over an array. For example, let's look at Listing 3-4:

Filename: src/main.rs

```
fn main() {  
    let a = [10, 20, 30, 40, 50];  
    let mut index = 0;  
  
    while index < 5 {  
        println!("the value is: {}", a[index]);  
  
        index = index + 1;  
    }  
}
```

#### Listing 3-4: Looping through each element of a c

Here, the code counts up through the elements then loops until it reaches the final index in the (no longer true). Running this code will print every e

```
$ cargo run
   Compiling loops v0.1.0 (file:///project
   Finished dev [unoptimized + debuginfo]
   Running `target/debug/loops`
the value is: 10
the value is: 20
the value is: 30
the value is: 40
the value is: 50
```

All five array values appear in the terminal, as expected. But when the index reaches a value of 5 at some point, the loop stops and doesn't print the sixth value from the array.

But this approach is error prone; we could cause a panic if the array length is incorrect. It's also slow, because the code has to perform the conditional check on every element on every iteration.

As a more concise alternative, you can use a `for` loop to iterate over each item in a collection. A `for` loop looks like this:

Filename: src/main.rs

```
fn main() {
    let a = [10, 20, 30, 40, 50];

    for element in a.iter() {
        println!("the value is: {}", element);
    }
}
```

#### Listing 3-5: Looping through each element of a c

When we run this code, we'll see the same output as Listing 3-4. Here, we've now increased the safety of the code and won't panic if we might result from going beyond the end of the array. We won't be missing some items.

For example, in the code in Listing 3-4, if you remembered to update the condition to `while index < a.len()` instead of `for` loop, you wouldn't need to remember to check the array length.

the number of values in the array.

The safety and conciseness of `for` loops make it a popular construct in Rust. Even in situations in which you don't know the number of times, as in the countdown example, most Rustaceans would use a `for` loop. The way to do this is which is a type provided by the standard library that represents a sequence starting from one number and ending

Here's what the countdown would look like using `for`, which we've not yet talked about, `rev`, to reverse the

Filename: src/main.rs

```
fn main() {  
    for number in (1..4).rev() {  
        println!("{}", number);  
    }  
    println!("LIFTOFF!!!");  
}
```

This code is a bit nicer, isn't it?

## Summary

You made it! That was a sizable chapter: you learned about compound data types, functions, comments, `if` statements, and how to practice with the concepts discussed in this chapter. Here are the following:

- Convert temperatures between Fahrenheit and Celsius.
- Generate the nth Fibonacci number.
- Print the lyrics to the Christmas carol "The Twelve Days of Christmas", demonstrating the advantage of the repetition in the song.

When you're ready to move on, we'll talk about `enum` types, which commonly exist in other programming languages.

## Understanding Ownership

Ownership is Rust's most unique feature, and it guarantees memory safety without needing a garbage collector. In this chapter, you'll understand how ownership works in Rust. In this

as well as several related features: borrowing, smart pointers, and memory.

## What Is Ownership?

Rust's central feature is *ownership*. Although the concept is simple, it has deep implications for the rest of the language.

All programs have to manage the way they use memory. Some languages have garbage collection that automatically manages memory as the program runs; in other languages, you must manually allocate and free the memory. Rust uses a third approach: it manages memory through a system of ownership with a set of rules that are easy to learn and time-efficient. None of the ownership features slow down the program.

Because ownership is a new concept for many programmers, it can be hard to get used to. The good news is that the more you learn about the rules of the ownership system, the more you will appreciate the code that is safe and efficient. Keep at it!

When you understand ownership, you'll have a solid foundation for the features that make Rust unique. In this chapter, we'll explore the features through some examples that focus on a very common use case.

---

## The Stack and the Heap

In many programming languages, you don't have to worry about the heap very often. But in a systems program, knowing where a value is on the stack or the heap has more consequences. In this section, we'll describe how memory behaves and why you have to make certain decisions when working with memory described in relation to the stack and the heap. We'll start with a brief explanation in preparation.

Both the stack and the heap are parts of memory that you can use at runtime, but they are structured in different ways. Values in the order it gets them and removes them. This is referred to as *last in, first out*. Think of a stack of plates: the more plates, you put them on top of the pile, and when you take one off the top. Adding or removing plates from the bottom wouldn't work as well! Adding data is called *pushing* and removing data is called *popping*.

data is called *popping off the stack*.

The stack is fast because of the way it accesses for a place to put new data or a place to get data always the top. Another property that makes stack must take up a known, fixed size.

Data with a size unknown at compile time or stored on the heap instead. The heap is less controlled than the stack, you ask for some amount of space. You find an empty spot somewhere in the heap that is big enough and returns a *pointer*, which is the address of that spot. *allocating on the heap*, sometimes abbreviated as *malloc*. Moving data onto the stack is not considered allocating. Because of its size, you can store the pointer on the stack, but you have to follow the pointer.

Think of being seated at a restaurant. When you arrive with a group of people in your group, and the staff finds an empty table and leads you there. If someone in your group comes later, they can be seated to find you.

Accessing data in the heap is slower than accessing data on the stack because you have to follow a pointer to get there. Conversely, data on the stack jumps around less in memory. Continuing the restaurant analogy, taking orders from many tables. If a waiter takes orders at one table before moving on to the next, it's efficient. If they take an order from table A, then an order from table B, then one from table C, it would be a much slower process. By the time they get to table A again, it's better if it works on data that's close to the top of the stack rather than farther away (as it can be on the heap). Managing space on the heap can also take time.

When your code calls a function, the values passed to the function (potentially, pointers to data on the heap) and the return value are pushed onto the stack. When the function is complete, the data is popped off the stack.

Keeping track of what parts of code are using data on the heap, the amount of duplicate data on the heap, and how to free it are all problems. Once you understand ownership, you won't need to manage the heap very often, but knowing that memory management exists can help explain why it works the way it does.

---

## Ownership Rules

First, let's take a look at the ownership rules. Keep in mind that we'll go through the examples that illustrate them:

- 
1. Each value in Rust has a variable that's associated with it.
  2. There can only be one owner at a time.
  3. When the owner goes out of scope, the value is deallocated.
- 

## Variable Scope

We've walked through an example of a Rust program, and now that we're past basic syntax, we won't include all the boilerplate. If you're following along, you'll have to put the following code into a function manually. As a result, our examples will focus on the actual details rather than boilerplate.

As a first example of ownership, we'll look at the concept of *scope*, the range within a program for which an item is valid. Listing 4-1 shows what it looks like this:

```
let s = "hello";
```

The variable `s` refers to a string literal, where the text is inserted into the text of our program. The variable is valid from the point it is declared until the end of the current *scope*. Listing 4-1 has a block of code where variable `s` is valid:

```
{  
    // s is not valid here  
    let s = "hello"; // s is valid from here  
    // do stuff with s  
} // this scope is now closed
```

Listing 4-1: A variable and the scope in which it is valid

In other words, there are two important points in the life of a variable:

- When `s` comes *into scope*, it is valid.
- It remains valid until it goes *out of scope*.



At this point, the relationship between scopes is different from that in other programming languages. Now we'll move forward by introducing the `String` type.

## The `String` Type

To illustrate the rules of ownership, we need a data type. The ones we covered in the “Data Types” section previously are all stored on the stack and popped off when they go out of scope, but we want to look at data that is stored on the heap and knows when to clean up that data.

We'll use `String` as the example here and connect it to ownership. These aspects also apply to `Vec` by the standard library and that you create. We'll cover `Vec` in Chapter 8.

We've already seen string literals, where a string is a sequence of characters. String literals are convenient, but they aren't suitable for all cases. You may want to use text. One reason is that they're not mutable. A string value can be known when we write our code, but you can't take user input and store it? For these situations, we use `String`. This type is allocated on the heap and can hold text that is unknown to us at compile time. You can create a `String` literal using the `from` function, like so:

```
let s = String::from("hello");
```

The double colon (`::`) is an operator that allows you to call the `from` function under the `String` type rather than `string_from`. We'll discuss this syntax more in the next chapter. In Chapter 5 and when we talk about namespacing in Chapter 7.

This kind of string *can* be mutated:

```
let mut s = String::from("hello");  
  
s.push_str(", world!"); // push_str() appends to the end  
  
println!("{}", s); // This will print `hello, world!`
```

So, what's the difference here? Why can `String` difference is how these two types deal with memory.

## Memory and Allocation

In the case of a string literal, we know the content is hardcoded directly into the final executable. This is efficient. But these properties only come from the fact that the size is known at compile time. Unfortunately, we can't put a blob of memory in the executable whose size is unknown at compile time and whose contents are determined at runtime by the program.

With the `String` type, in order to support a mutable string, we need to allocate an amount of memory on the heap, and we need to manage its contents. This means:

- The memory must be requested from the operating system.
- We need a way of returning this memory to the operating system when we're done with our `String`.

The first part is done by us: when we call `String::new`, we request the memory it needs. This is pretty much universal across all languages.

However, the second part is different. In languages with a garbage collector (GC), the GC keeps track and cleans up memory that isn't needed anymore. Without a GC, it's our responsibility to keep track of memory that's no longer being used and call code to explicitly release it. Doing this correctly has historically been a difficult task. If we forget, we'll waste memory. If we do it too early, we'll have a memory leak. If we do it twice, that's a bug too. We need to pair `new` with `free`.

Rust takes a different path: the memory is automatically managed. When a variable that owns it goes out of scope, the memory is freed. Here's a version of a function that uses a `String` instead of a string literal:

```
{
    let s = String::from("hello"); // s is a String
    // do stuff with s
}                                     // this line is never reached
                                     // because s is freed
```

There is a natural point at which we can return the memory to the operating system.

operating system: when `s` goes out of scope. We call a special function for us. This function is called `String::drop_in_place`. We can put the code to return the memory to the operating system in the closing `}`.

---

Note: In C++, this pattern of deallocating resources at the end of their lifetime is sometimes called *Resource Acquisition Is Initialization*. This function in Rust will be familiar to you if you've used C++.

---

This pattern has a profound impact on the way we write code. It's simple right now, but the behavior of code can be very different in situations when we want to have multiple variables pointing to the same memory on the heap. Let's explore some of those situations.

## Ways Variables and Data Interact: Move

Multiple variables can interact with the same data. Let's look at an example using an integer in Listing 4-2:

```
let x = 5;
let y = x;
```

Listing 4-2: Assigning the integer value of variable `x` to `y`

We can probably guess what this is doing: “bind the value 5 to `x` and bind it to `y`.” We now have two variables, `x` and `y`, both of which hold the value 5. This is indeed what is happening, because integers are known, fixed size, and these two 5 values are pointers to the same memory location.

Now let's look at the `String` version:

```
let s1 = String::from("hello");
let s2 = s1;
```

This looks very similar to the previous code, so you might expect the same: that is, the second line would bind `s2` to the same memory as `s1`. But this isn't quite what happens.

Take a look at Figure 4-1 to see what is happening. The `String` type is made up of three parts, shown on the right: a pointer to the memory that holds the contents of the string, a length, and a capacity.

on the stack. On the right is the memory on the

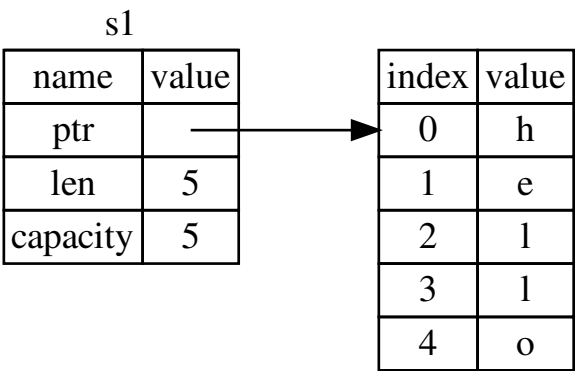


Figure 4-1: Representation in memory of a `String` object `s1`

The length is how much memory, in bytes, the object is using. The capacity is the total amount of memory that the object has received from the operating system. The difference between the two is the free space in the memory. It matters, but not in this context, so for now, it's fine.

When we assign `s1` to `s2`, the `String` data is copied into the memory. The length, and the capacity that are on the stack are the same. In other words, it looks like Figure 4-2.

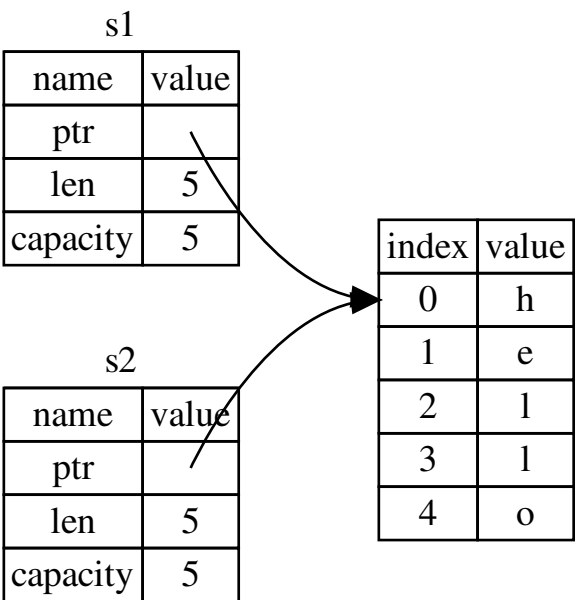


Figure 4-2: Representation in memory of the variable `s1`, its pointer, length, and capacity

The representation does *not* look like Figure 4-3,

like if Rust instead copied the heap data as well. `s2 = s1` could be very expensive in terms of ru heap were large.

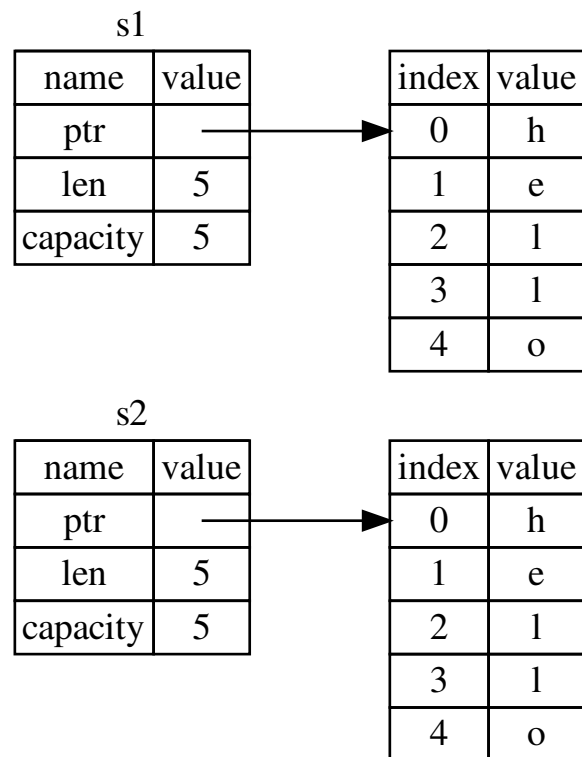


Figure 4-3: Another possibility for what `s2 = s1` data as well

Earlier, we said that when a variable goes out of `drop` function and cleans up the heap memory both data pointers pointing to the same location go out of scope, they will both try to free the same `free` error and is one of the memory safety bugs memory twice can lead to memory corruption, vulnerabilities.

To ensure memory safety, there's one more detail in Rust. Instead of trying to copy the allocated memory to be valid and, therefore, Rust doesn't need to free memory when it goes out of scope. Check out what happens when you try to work:

```
let s1 = String::from("hello");
let s2 = s1;

println!("{}", world!", s1);
```

You'll get an error like this because Rust prevents reference:

```
error[E0382]: use of moved value: `s1`
--> src/main.rs:5:28
   |
3  |     let s2 = s1;
   |             -- value moved here
4  |
5  |     println!("{}", world!", s1);
   |                                   ^^ value used here
   = note: move occurs because `s1` has type `String`, which does not implement the `Copy` trait
```

If you've heard the terms *shallow copy* and *deep copy* languages, the concept of copying the pointer, like in C, where the data probably sounds like making a shallow copy. This invalidates the first variable, instead of being called *move*. Here we would read this by saying that `s1` actually happens is shown in Figure 4-4.

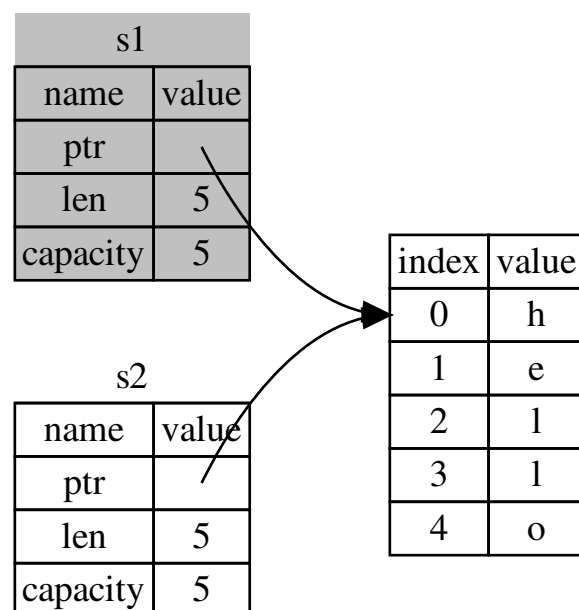


Figure 4-4: Representation in memory after `s1` is moved to `s2`

That solves our problem! With only `s2` valid, we can free the memory, and we're done.

In addition, there's a design choice that's implied by Rust: to create "deep" copies of your data. Therefore, any operation that involves copying data to be inexpensive in terms of runtime performance.

## Ways Variables and Data Interact: Clone

If we *do* want to deeply copy the heap data of the `String`, we can use a common method called `clone`. We'll see how it works, but because methods are a common feature in many languages, you've probably seen them before.

Here's an example of the `clone` method in action:

```
let s1 = String::from("hello");
let s2 = s1.clone();

println!("s1 = {}, s2 = {}", s1, s2);
```

This works just fine and explicitly produces the `String` that the heap data *does* get copied.

When you see a call to `clone`, you know that so much work is going on and that code may be expensive. It's a visual indicator that something is going on.

## Stack-Only Data: Copy

There's another wrinkle we haven't talked about yet. The `let` keyword, which was shown earlier in Listing 4-2, works on both stack and heap data.

```
let x = 5;
let y = x;

println!("x = {}, y = {}", x, y);
```

But this code seems to contradict what we just learned: `let` moves `x` into `y`, but `x` is still valid and wasn't moved into `y`.

The reason is that types such as integers that are stored entirely on the stack, so copies of the actual data means there's no reason we would want to prevent creating the variable `y`. In other words, there's no shallow copying here, so calling `clone` wouldn't be necessary for shallow copying and we can leave it out.

Rust has a special annotation called the `Copy` trait for integers that are stored on the stack (we'll talk more about this later).

type has the `Copy` trait, an older variable is still valid. To let us annotate a type with the `Copy` trait if the type has implemented the `Drop` trait. If the type needs some value goes out of scope and we add the `Copy` annotation, it will cause a compile time error. To learn about how to add the `Copy` trait, see “Derivable Traits” in Appendix C.

So what types are `Copy`? You can check the documentation, but as a general rule, any group of simple types that requires allocation or is some form of pointer are not `Copy`. The types that are `Copy`:

- All the integer types, such as `u32`.
- The Boolean type, `bool`, with values `true` and `false`.
- All the floating point types, such as `f64`.
- The character type, `char`.
- Tuples, but only if they contain types that are `Copy`. For example, `(i32, i32)` is `Copy`, but `(i32, String)` is not.

## Ownership and Functions

The semantics for passing a value to a function are different from passing a value to a variable. Passing a variable to a function does not work. Listing 4-3 has an example with variables go into and out of scope:

Filename: src/main.rs



```

fn main() {
    let s = String::from("hello"); // s c

    takes_ownership(s);           // s's
    function...                   // ...

    let x = 5;                    // x c

    makes_copy(x);                // x v
                                // but
    still                          // use

} // Here, x goes out of scope, then s. But
  // nothing
  // special happens.

fn takes_ownership(some_string: String) {
    println!("{}", some_string);
} // Here, some_string goes out of scope and
  // memory is freed.

fn makes_copy(some_integer: i32) { // some
    println!("{}", some_integer);
} // Here, some_integer goes out of scope.

```

Listing 4-3: Functions with ownership and scope

If we tried to use `s` after the call to `takes_ownership`, we would get a compile-time error. These static checks protect us from runtime errors. The compiler uses `s` and `x` to see where you can use them and prevent you from doing so.

## Return Values and Scope

Returning values can also transfer ownership. Let's add some return type annotations to those in Listing 4-3:

Filename: src/main.rs

```

fn main() {
    let s1 = gives_ownership();    //
    return                          //

    let s2 = String::from("hello");    //

    let s3 = takes_and_gives_back(s2); //
                                        //
    also                                //
} // Here, s3 goes out of scope and is dropped
    // moved, so nothing happens. s1 goes out of scope

fn gives_ownership() -> String {
    let s = String::from("hello");
    its ownership is transferred to the caller

    function

    let some_string = String::from("hello");
    scope

    some_string
    and

}

// takes_and_gives_back will take a String and return it
fn takes_and_gives_back(a_string: String) -> String {
    into

    a_string // a_string is returned and ownership is transferred to the caller
    function
}

```

Listing 4-4: Transferring ownership of return values

The ownership of a variable follows the same pattern as when another variable moves it. When a variable that owns a value goes out of scope, the value will be cleaned up by `drop` unless it is owned by another variable.

Taking ownership and then returning ownership. What if we want to let a function use a value but not take ownership of it? It's a bit annoying that anything we pass in also needs to be dropped at the end, in addition to any data resulting from the function.

want to return as well.

It's possible to return multiple values using a tuple.

Filename: src/main.rs

```
fn main() {
    let s1 = String::from("hello");

    let (s2, len) = calculate_length(s1);

    println!("The length of '{}' is {}.",
}

fn calculate_length(s: String) -> (String,
    let length = s.len(); // len() returns

    (s, length)
}
```

Listing 4-5: Returning ownership of parameters

But this is too much ceremony and a lot of work common. Luckily for us, Rust has a feature for that.

## References and Borrowing

The issue with the tuple code in Listing 4-5 is that the calling function so we can still use the `String` because the `String` was moved into `calculate_length`.

Here is how you would define and use a `calculate_length` reference to an object as a parameter instead of

Filename: src/main.rs

```
fn main() {
    let s1 = String::from("hello");

    let len = calculate_length(&s1);

    println!("The length of '{}' is {}.",
}

fn calculate_length(s: &String) -> usize {
    s.len()
}
```

First, notice that all the tuple code in the variable value is gone. Second, note that we pass `&s1` in definition, we take `&String` rather than `String`

These ampersands are *references*, and they allow taking ownership of it. Figure 4-5 shows a diagram

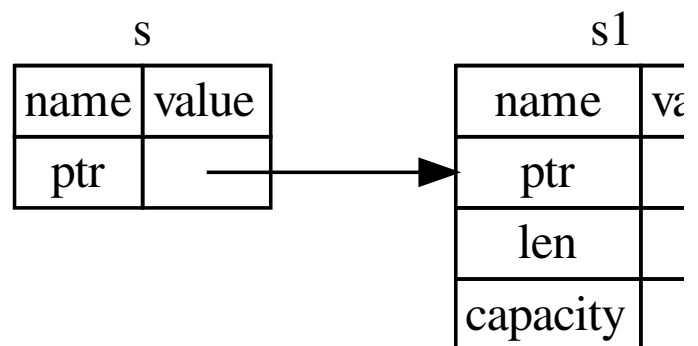


Figure 4-5: A diagram of `&String s` pointing at

---

Note: The opposite of referencing by using `&` accomplished with the dereference operator, dereference operator in Chapter 8 and discuss Chapter 15.

---

Let's take a closer look at the function call here:

```
let s1 = String::from("hello");

let len = calculate_length(&s1);
```

The `&s1` syntax lets us create a reference that *borrows* it. Because it does not own it, the value it points to goes out of scope.

Likewise, the signature of the function uses `&` to indicate that parameter `s` is a reference. Let's add some explanatory comments:

```
fn calculate_length(s: &String) -> usize {
    s.len()
} // Here, s goes out of scope. But because we only have a reference
    // it refers to, nothing happens.
```

The scope in which the variable `s` is valid is the function's scope, but we don't drop what the reference points to because we don't have ownership. When function returns, instead of the actual values, we won't need to re-allocate memory, because we never had ownership.

We call having references as function parameters *borrowing*. If you own something, you can borrow it from them. When you're done, you return it back.

So what happens if we try to modify something borrowed? Let's look at Listing 4-6. Spoiler alert: it doesn't work!

Filename: src/main.rs

```
fn main() {
    let s = String::from("hello");

    change(&s);
}

fn change(some_string: &String) {
    some_string.push_str(", world");
}
```

Listing 4-6: Attempting to modify a borrowed value

Here's the error:

```

error[E0596]: cannot borrow immutable borrow
mutable
--> error.rs:8:5
|
7 | fn change(some_string: &String) {
|                               ----- use `&n
8 |     some_string.push_str(", world");
|     ^^^^^^^^^^^^^^^^^ cannot borrow as mutab

```

Just as variables are immutable by default, so are references. To modify something we have a reference to.

## Mutable References

We can fix the error in the code from Listing 4-6

Filename: src/main.rs

```

fn main() {
    let mut s = String::from("hello");

    change(&mut s);
}

fn change(some_string: &mut String) {
    some_string.push_str(", world");
}

```

First, we had to change `s` to be `mut`. Then we had to pass `&mut s` and accept a mutable reference with `&mut String`.

But mutable references have one big restriction: you can only have one mutable reference to a particular piece of data in a particular scope.

Filename: src/main.rs

```

let mut s = String::from("hello");

let r1 = &mut s;
let r2 = &mut s;

```

Here's the error:

```

error[E0499]: cannot borrow `s` as mutable
--> borrow_twice.rs:5:19
  |
4 |     let r1 = &mut s;
  |               - first mutable borrow
5 |     let r2 = &mut s;
  |               ^ second mutable borrow
6 | }
  | - first borrow ends here

```

This restriction allows for mutation but in a very way that new Rustaceans struggle with, because most of the time you'd like.

The benefit of having this restriction is that Rust is safe at compile time. A *data race* is similar to a race condition and can occur in the following behaviors occur:

- Two or more pointers access the same data
- At least one of the pointers is being used to mutate the data
- There's no mechanism being used to synchronize access

Data races cause undefined behavior and can be very hard to debug because you're trying to track them down at runtime; Rust prevents them at compile time because it won't even compile code with a data race.

As always, we can use curly brackets to create a scope for mutable references, just not *simultaneous* ones:

```

let mut s = String::from("hello");

{
    let r1 = &mut s;

    // r1 goes out of scope here, so we can
    // use s again. No problems.

    let r2 = &mut s;
}

```

A similar rule exists for combining mutable and immutable references and results in an error:

```

let mut s = String::from("hello");

let r1 = &s; // no problem
let r2 = &s; // no problem
let r3 = &mut s; // BIG PROBLEM

```

Here's the error:

```
error[E0502]: cannot borrow `s` as mutable
  immutable
  --> borrow_thrice.rs:6:19
    |
4  |         let r1 = &s; // no problem
    |                   - immutable borrow occurs here
5  |         let r2 = &s; // no problem
6  |         let r3 = &mut s; // BIG PROBLEM
    |                   ^ mutable borrow occurs here
7  |     }
    |     - immutable borrow ends here
```

Whew! We *also* cannot have a mutable reference. Users of an immutable reference don't expect to be able to mutate the data from under them! However, multiple immutable references to the same data, where one of them who is just reading the data has the ability to affect the data, is a problem.

Even though these errors may be frustrating at first, the Rust compiler pointing out a potential bug early (at compile time) and showing you exactly where the problem is, is a good thing. It's why your data isn't what you thought it was.

## Dangling References

In languages with pointers, it's easy to erroneously free memory that references a location in memory that may have been freed. In Rust, when freeing some memory while preserving a pointer to it, the compiler guarantees that references will never be freed. If you have a reference to some data, the compiler will ensure that the data is not freed before the reference to the data does.

Let's try to create a dangling reference, which Rust will catch with an error:

Filename: src/main.rs



```
fn main() {
    let reference_to_nothing = dangle();
}

fn dangle() -> &String {
    let s = String::from("hello");

    &s
}
```

Here's the error:

```
error[E0106]: missing lifetime specifier
--> dangle.rs:5:16
|
5 | fn dangle() -> &String {
|                  ^ expected lifetime parameter
= help: this function's return type contains a borrowed value, but there is no value for it to be borrowed from
= help: consider giving it a 'static lifetime
```

This error message refers to a feature we haven't covered yet: lifetimes in detail in Chapter 10. But, if you disregard the error message, it does contain the key to why this code is problematic.

```
this function's return type contains a borrowed value
but there is no value for it to be borrowed from.
```

Let's take a closer look at exactly what's happening here.

```
fn dangle() -> &String { // dangle returns a reference to a String

    let s = String::from("hello"); // s is a String

    &s // we return a reference to the String s
} // Here, s goes out of scope, and is deallocated
// Danger!
```

Because `s` is created inside `dangle`, when the function returns, `s` is deallocated. But we tried to return a reference to `s`, which is now pointing to an invalid `String`. That's no good.

The solution here is to return the `String` directly instead of a reference to it.

```
fn no_dangle() -> String {
    let s = String::from("hello");

    s
}
```

This works without any problems. Ownership is deallocated.

## The Rules of References

Let's recap what we've discussed about references:

- At any given time, you can have *either* (but not both) one mutable reference or any number of immutable references.
- References must always be valid.

Next, we'll look at a different kind of reference: `slice`.

## The Slice Type

Another data type that does not have ownership is a `slice`, which represents a contiguous sequence of elements in a collection.

Here's a small programming problem: write a function that returns the first word it finds in that string. If the function finds the whole string must be one word, so the entire string is the first word.

Let's think about the signature of this function:

```
fn first_word(s: &String) -> ?
```

This function, `first_word`, has a `&String` as a parameter, so this is fine. But what should we return? We don't want to return a *part* of a string. However, we could return the index of the first character of that word, as shown in Listing 4-7:

Filename: src/main.rs

```
fn first_word(s: &String) -> usize {
    let bytes = s.as_bytes();

    for (i, &item) in bytes.iter().enumerate() {
        if item == b' ' {
            return i;
        }
    }

    s.len()
}
```

Listing 4-7: The `first_word` function that returns the index of the first space in a string parameter

Because we need to go through the `String` element value is a space, we'll convert our `String` to an array of bytes using the `as_bytes()` method:

```
let bytes = s.as_bytes();
```

Next, we create an iterator over the array of bytes using the `iter()` method:

```
for (i, &item) in bytes.iter().enumerate()
```

We'll discuss iterators in more detail in Chapter 10. The `iter()` method that returns each element in a collection is part of the `Iterator` trait. The result of `iter` and returns each element as part of a tuple. The first element of the tuple returned from `enumerate` is the index, and the second element is the reference to the element. This is a bit more convenient than using `len` and `get` ourselves.

Because the `enumerate` method returns a tuple, we can use a `for` loop that takes a tuple. Just like everywhere else in Rust. So in the `for` loop, we have `i` for the index in the tuple and `&item` for the reference to the element. Because we get a reference to the element from the pattern, we can use `item` directly in the loop body.

Inside the `for` loop, we search for the byte that is a space using the byte literal syntax. If we find a space, we return the index of the space by using `s.len()`:

```

        if item == b' ' {
            return i;
        }
    }
    s.len()

```

We now have a way to find out the index of the space character, but there's a problem. We're returning a `usize` on a `String` index number in the context of the `&String`. In other words, coming from the `String`, there's no guarantee that it will work in the program in Listing 4-8 that uses the `first_word` function.

Filename: src/main.rs

```

fn main() {
    let mut s = String::from("hello world");

    let word = first_word(&s); // word will have a value of 5

    s.clear(); // This empties the String,
               // but word still has the value 5 here, but it's
               // invalid!
}

```

Listing 4-8: Storing the result from calling the `first_word` function in the `String` contents

This program compiles without any errors and runs successfully. But by calling `s.clear()`, we've emptied the `String`. Because `word` isn't connected to the `String`, it still contains the value `5`. We could use that value `5` to find the first word out, but this would be a bug because the `String` is no longer valid since we saved `5` in `word`.

Having to worry about the index in `word` getting out of sync is tedious and error prone! Managing these indices is a pain. Let's create a `second_word` function. Its signature would have to be:

```

fn second_word(s: &String) -> (usize, usize)

```

Now we're tracking a starting *and* an ending index. These indices were calculated from data in a particular `String`. We now have three unrelated variables floating around: `s`, `word`, and `second_word`.

Luckily, Rust has a solution to this problem: `str::split`.

## String Slices

A *string slice* is a reference to part of a `String`, and

```
let s = String::from("hello world");  
  
let hello = &s[0..5];  
let world = &s[6..11];
```

This is similar to taking a reference to the whole bit. Rather than a reference to the entire `String`, we have a reference to a slice of `String`. The `start..end` syntax is a range that starts at `start` and ends at `end`, but not including `end`. If we wanted to include

```
let s = String::from("hello world");  
  
let hello = &s[0..=4];  
let world = &s[6..=10];
```

The `=` means that we're including the last number. The difference between `..` and `..=` is

We can create slices using a range within brackets: `[starting_index..ending_index]`, where `starting_index` is the starting position of the slice and `ending_index` is one more than the last byte of the slice. The slice data structure stores the starting position and the ending position. The difference between `ending_index` and `starting_index` corresponds to the length of the slice. For example, `let world = &s[6..11];`, `world` would be a slice of `s` and a length value of 5.

Figure 4-6 shows this in a diagram.

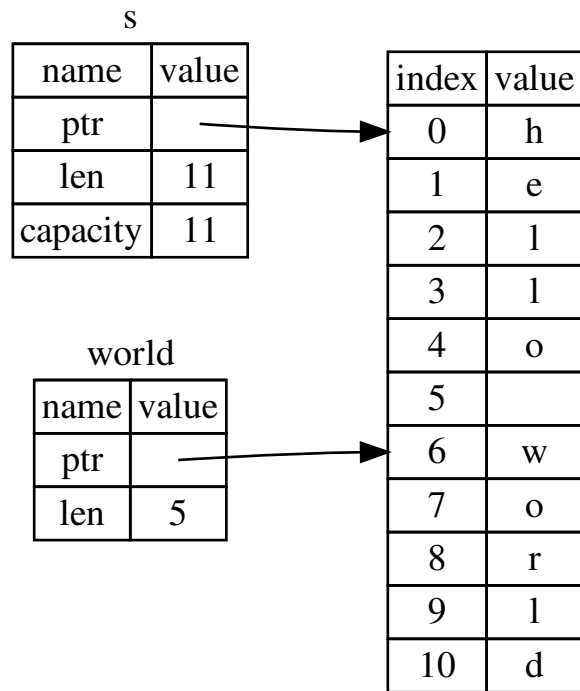


Figure 4-6: String slice referring to part of a `String`

With Rust's `..` range syntax, if you want to start the value before the two periods. In other words:

```
let s = String::from("hello");

let slice = &s[0..2];
let slice = &s[..2];
```

By the same token, if your slice includes the last trailing number. That means these are equal:

```
let s = String::from("hello");

let len = s.len();

let slice = &s[3..len];
let slice = &s[3..];
```

You can also drop both values to take a slice of the

```
let s = String::from("hello");

let len = s.len();

let slice = &s[0..len];
let slice = &s[..];
```

---

Note: String slice range indices must occur at  
If you attempt to create a string slice in the m  
program will exit with an error. For the purpo  
are assuming ASCII only in this section; a mor  
handling is in the “Strings” section of Chapter

---

With all this information in mind, let’s rewrite `fi`  
that signifies “string slice” is written as `&str`:

Filename: src/main.rs

```
fn first_word(s: &String) -> &str {
    let bytes = s.as_bytes();

    for (i, &item) in bytes.iter().enumerate() {
        if item == b' ' {
            return &s[0..i];
        }
    }

    &s[..]
}
```

We get the index for the end of the word in the s  
looking for the first occurrence of a space. When  
slice using the start of the string and the index o  
indices.

Now when we call `first_word`, we get back a si  
underlying data. The value is made up of a refer  
and the number of elements in the slice.

Returning a slice would also work for a `second_w`

```
fn second_word(s: &String) -> &str {
```

We now have a straightforward API that’s much

compiler will ensure the references into the `String` in the program in Listing 4-8, when we got the `String` then cleared the string so our index was invalid? It didn't show any immediate errors. The problem was trying to use the first word index with an empty `String` is impossible and let us know we have a problem with the slice version of `first_word` will throw a compile

Filename: src/main.rs

```
fn main() {
    let mut s = String::from("hello world");

    let word = first_word(&s);

    s.clear(); // Error!
}
```

Here's the compiler error:

```
error[E0502]: cannot borrow `s` as mutable and immutable
--> src/main.rs:6:5
   |
4  |     let word = first_word(&s);
   |                               - immutable borrow starts here
5  |
6  |     s.clear(); // Error!
   |     ^ mutable borrow occurs here
7  | }
   | - immutable borrow ends here
```

Recall from the borrowing rules that if we have a mutable reference, we cannot also take an immutable reference. Because `String`, it tries to take a mutable reference, which is an API easier to use, but it has also eliminated an error.

## String Literals Are Slices

Recall that we talked about string literals being slices. Now that we know about slices, we can properly understand them.

```
let s = "Hello, world!";
```

The type of `s` here is `&str`: it's a slice pointing to a string literal.



This is also why string literals are immutable; `&s`

## String Slices as Parameters

Knowing that you can take slices of literals and improvement on `first_word`, and that's its sign

```
fn first_word(s: &String) -> &str {
```

A more experienced Rustacean would write the allows us to use the same function on both `Str`

```
fn first_word(s: &str) -> &str {
```

If we have a string slice, we can pass that directly slice of the entire `String`. Defining a function to reference to a `String` makes our API more general functionality:

Filename: src/main.rs

```
fn main() {  
    let my_string = String::from("hello world");  
  
    // first_word works on slices of `String`  
    let word = first_word(&my_string[..]);  
  
    let my_string_literal = "hello world";  
  
    // first_word works on slices of `str`  
    let word = first_word(&my_string_literal);  
  
    // Because string literals are string slices  
    // this works too, without the slice reference  
    let word = first_word(my_string_literal);  
}
```

## Other Slices

String slices, as you might imagine, are specific to slice type, too. Consider this array:

```
let a = [1, 2, 3, 4, 5];
```

Just as we might want to refer to a part of a string or an array. We'd do so like this:

```
let a = [1, 2, 3, 4, 5];  
  
let slice = &a[1..3];
```

This slice has the type `&[i32]`. It works the same way as a reference to the first element and a length. You can use it with other collections. We'll discuss these collections in Chapter 8.

## Summary

The concepts of ownership, borrowing, and slices are checked by the compiler at compile time. The Rust language gives you a way to use data in the same way as other systems programming languages. The owner of data automatically cleans up that data when it goes out of scope, which means you don't have to write and debug extra code.

Ownership affects how lots of other parts of Rust work. We'll discuss these concepts further throughout the rest of the book. In the next chapter, we'll look at grouping pieces of data together in a `struct`.

## Using Structs to Structure Data

A *struct*, or *structure*, is a custom data type that lets you group multiple related values that make up a meaningful data structure. In an object-oriented language, a *struct* is like an object. In this chapter, we'll compare and contrast tuples with structs, and we'll discuss how to define methods and associated functions for structs. Structs and enums are the building blocks for creating new types in your programs, and they take advantage of Rust's compile-time type checking.

## Defining and Instantiating Structs

Structs are similar to tuples, which were discussed in the previous chapter.

pieces of a struct can be different types. Unlike `Vec` data so it's clear what the values mean. As a result, structs are more flexible than tuples: you don't have to rely on the order of fields to access the values of an instance.

To define a struct, we enter the keyword `struct` followed by the struct's name. The struct's name should describe the significance of the data it holds together. Then, inside curly brackets, we define the fields of the struct, which we call *fields*. For example, Listing 5-1 shows the definition of a struct for information about a user account:

```
struct User {  
    username: String,  
    email: String,  
    sign_in_count: u64,  
    active: bool,  
}
```

Listing 5-1: A `User` struct definition

To use a struct after we've defined it, we create an instance of the struct with concrete values for each of the fields. We create an instance of the struct and then add curly brackets containing the field names and their values. The field names are the names of the fields and the values are the values of the fields. We don't have to specify the fields in the instance; they are defined in the struct. In other words, the struct definition is a template, and instances fill in that template with specific values. For example, we can declare a particular user instance:

```
let user1 = User {  
    email: String::from("someone@example.com"),  
    username: String::from("someusername123"),  
    active: true,  
    sign_in_count: 1,  
};
```

Listing 5-2: Creating an instance of the `User` struct

To get a specific value from a struct, we can use the field access syntax. For example, to get the user's email address, we could use `user1.email`. If the instance is mutable, we can change the value by assigning into a particular field. Listing 5-3 shows how to change the `email` field of a mutable `User` instance:

```

let mut user1 = User {
    email: String::from("someone@example.com"),
    username: String::from("someusername123"),
    active: true,
    sign_in_count: 1,
};

user1.email = String::from("anotheremail@example.com");

```

Listing 5-3: Changing the value in the `email` field

Note that the entire instance must be mutable; I can't make certain fields as mutable.

As with any expression, we can construct a new expression in the function body to implicitly return a value. Listing 5-4 shows a `build_user` function that returns a `User` instance with a given `email` and `username`. The `active` field gets the value of `true` and the `sign_in_count` field gets the value of `1`.

```

fn build_user(email: String, username: String) -> User {
    User {
        email: email,
        username: username,
        active: true,
        sign_in_count: 1,
    }
}

```

Listing 5-4: A `build_user` function that takes an `email` and `username` and returns a `User` instance

It makes sense to name the function parameters `email` and `username` to match the struct fields, but having to repeat the `email` and `username` in the function signature is a bit tedious. If the struct had more fields, repeating them would be even more annoying. Luckily, there's a convenient shorthand for this.

## Using the Field Init Shorthand when Variable and Field Have the Same Name

Because the parameter names and the struct field names are the same in Listing 5-4, we can use the *field init shorthand* syntax. The `field init shorthand` behaves exactly the same but doesn't have the `email` and `username` parameters in the function signature.

shown in Listing 5-5.

```
fn build_user(email: String, username: String) -> User {
    User {
        email,
        username,
        active: true,
        sign_in_count: 1,
    }
}
```

Listing 5-5: A `build_user` function that uses `email` and `username` parameters have the same name

Here, we're creating a new instance of the `User` struct. We want to set the `email` field's value to the `email` parameter of the `build_user` function. Because the `email` field has the same name, we only need to write `email` rather than `email: email`.

## Creating Instances From Other Instances

It's often useful to create a new instance of a struct with the same values but changes some. You'll do this using struct update syntax.

First, Listing 5-6 shows how we create a new `User` instance using struct update syntax. We set new values for `email` and `username` and use the same values from `user1` that we created in Listing 5-5.

```
let user2 = User {
    email: String::from("another@example.com"),
    username: String::from("anotherusername"),
    active: user1.active,
    sign_in_count: user1.sign_in_count,
};
```

Listing 5-6: Creating a new `User` instance using struct update syntax

Using struct update syntax, we can achieve the same result as Listing 5-5. The syntax `..` specifies that the new struct should have the same value as the fields in the original struct.

```
let user2 = User {
    email: String::from("another@example.c
    username: String::from("anotherusern
    ..user1
};
```

Listing 5-7: Using struct update syntax to set new `User` instance but use the rest of the values from `user1` variable

The code in Listing 5-7 also creates an instance of `User` with `email` and `username` but has the same values for the other fields from `user1`.

## Tuple Structs without Named Fields to

You can also define structs that look similar to tuples. Tuple structs have the added meaning the struct name provides with their fields; rather, they just have the types of the fields. When you want to give the whole tuple a name and make it different from other tuples, and naming each field as in a tuple is redundant.

To define a tuple struct start with the `struct` keyword followed by the types in the tuple. For example, here are two tuple structs named `Color` and `Point`:

```
struct Color(i32, i32, i32);
struct Point(i32, i32, i32);

let black = Color(0, 0, 0);
let origin = Point(0, 0, 0);
```

Note that the `black` and `origin` values are different instances of different tuple structs. Each struct type has its own fields. Although the fields within the struct have the same type, a `Color` cannot take a `Point` value. Both types are made up of three `i32` values. They behave like tuples: you can destructure them into a tuple, followed by the index to access an individual field.

## Unit-Like Structs Without Any Fields

You can also define structs that don't have any fields because they behave similarly to `()`, the unit type, in situations in which you need to implement a trait that you want to store in the type itself. We

---

## Ownership of Struct Data

In the `User` struct definition in Listing 5-1, we use `&str` rather than the `String` type. This is because we want instances of this struct to own all of its data for as long as the entire struct is valid.

It's possible for structs to store references to other data. To do so requires the use of *lifetimes*, a Rust feature introduced in Rust 1.0. Lifetimes ensure that the data referenced by the struct is valid for as long as the struct is. Let's say you try to store a reference to a string, like this, which won't work:

Filename: src/main.rs

```
struct User {
    username: &str,
    email: &str,
    sign_in_count: u64,
    active: bool,
}

fn main() {
    let user1 = User {
        email: "someone@example.com",
        username: "someusername123",
        active: true,
        sign_in_count: 1,
    };
}
```

The compiler will complain that it needs lifetimes

```

error[E0106]: missing lifetime specifier
-->
  |
2 |     username: &str,
  |               ^ expected lifetime parameter

error[E0106]: missing lifetime specifier
-->
  |
3 |     email: &str,
  |           ^ expected lifetime parameter

```

In Chapter 10, we'll discuss how to fix these errors for structs, but for now, we'll fix errors like these by using references instead of references like `&str`.

---

## An Example Program Using Structs

To understand when we might want to use structs, let's write a program that calculates the area of a rectangle. We'll start with a program that uses variables, and then we'll modify it to use structs instead.

Let's make a new binary project with Cargo called `area`. The width and height of a rectangle specified in pixels and Listing 5-8 shows a short program with one way to calculate the area of a rectangle using variables. The file `src/main.rs`:

Filename: `src/main.rs`

```

fn main() {
    let width1 = 30;
    let height1 = 50;

    println!(
        "The area of the rectangle is {} square pixels",
        area(width1, height1)
    );
}

fn area(width: u32, height: u32) -> u32 {
    width * height
}

```

Listing 5-8: Calculating the area of a rectangle specified by width and height variables



Now, run this program using `cargo run`:

```
The area of the rectangle is 1500 square p
```

Even though Listing 5-8 works and figures out the `area` function with each dimension, we can do related to each other because together they des

The issue with this code is evident in the signatu

```
fn area(width: u32, height: u32) -> u32 {
```

The `area` function is supposed to calculate the function we wrote has two parameters. The parameter is expressed anywhere in our program. It would be manageable to group width and height together might do that in “The Tuple Type” section of Cha

## Refactoring with Tuples

Listing 5-9 shows another version of our program

Filename: `src/main.rs`

```
fn main() {  
    let rect1 = (30, 50);  
  
    println!(  
        "The area of the rectangle is {} s  
        area(rect1)  
    );  
}  
  
fn area(dimensions: (u32, u32)) -> u32 {  
    dimensions.0 * dimensions.1  
}
```

Listing 5-9: Specifying the width and height of th

In one way, this program is better. Tuples let us passing just one argument. But in another way, name their elements, so our calculation has become have to index into the parts of the tuple.

It doesn't matter if we mix up width and height f to draw the rectangle on the screen, it would ma

that `width` is the tuple index `0` and `height` is `1`. If we had not worked on this code, they would have to figure it out. It would be easy to forget or mix up these values and not convey the meaning of our data in our code.

## Refactoring with Structs: Adding More

We use structs to add meaning by labeling the data. We can use `using` into a data type with a name for the whole struct, as shown in Listing 5-10:

Filename: `src/main.rs`

```
struct Rectangle {
    width: u32,
    height: u32,
}

fn main() {
    let rect1 = Rectangle { width: 30, height: 50 };

    println!(
        "The area of the rectangle is {} square feet",
        area(&rect1)
    );
}

fn area(rectangle: &Rectangle) -> u32 {
    rectangle.width * rectangle.height
}
```

Listing 5-10: Defining a `Rectangle` struct

Here we've defined a struct and named it `Rectangle`. We defined the fields as `width` and `height`, both of type `u32`. Then we created a particular instance of `Rectangle` named `rect1` with `width` of 30 and `height` of 50.

Our `area` function is now defined with one parameter, `rectangle`, whose type is an immutable borrow of `Rectangle`. As mentioned in Chapter 4, we want to borrow the struct rather than own it. This way, `main` retains its ownership and can use `rect1` after the `area` function returns. Because we use the `&` in the function signature, we know that the function is taking a reference to the struct.

The `area` function accesses the `width` and `height` fields of the `Rectangle` struct. Our function signature for `area` now says exactly what it does: it takes a reference to a `Rectangle` and returns a `u32`.

of `Rectangle`, using its `width` and `height` fields. `width` and `height` are related to each other, and it gives less information than using the tuple index values of `0` and `1`. The

## Adding Useful Functionality with Derive

It'd be nice to be able to print an instance of `Rectangle` in a Rust program and see the values for all its fields. Listing 5-11 shows the `println!` macro as we have used in previous chapters. The

Filename: `src/main.rs`

```
struct Rectangle {
    width: u32,
    height: u32,
}

fn main() {
    let rect1 = Rectangle { width: 30, height: 50 };

    println!("rect1 is {}", rect1);
}
```

Listing 5-11: Attempting to print a `Rectangle` in Rust

When we run this code, we get an error with this message:

```
error[E0277]: the trait bound `Rectangle: Display` is not
satisfied
```

The `println!` macro can do many kinds of formatting. To tell `println!` to use formatting known as `Display`, we need to implement `Display` for `Rectangle`. The primitive types we've seen so far implement `Display` by default, because there's only one way you'd want to print a primitive type to a user. But with structs, the way to print them is not so clear because there are more display possibilities. Should we print the curly brackets? Should we print the field names? To avoid ambiguity, Rust doesn't try to guess what we want. Instead, we need to implement an implementation of `Display`.

If we continue reading the errors, we'll find this message:

```
`Rectangle` cannot be formatted with the `{:?}` formatter
`{:?}` instead if you are using a format string
```

Let's try it! The `println!` macro call will now look like `println!("rect1 is {:?}", rect1);`. Putting brackets tells `println!` we want to use an output trait that enables us to print our struct in a way that we can see its value while we're debugging our code.

Run the code with this change. Drat! We still get

```
error[E0277]: the trait bound `Rectangle:
    Debug` is not satisfied
```

But again, the compiler gives us a helpful note:

```
`Rectangle` cannot be formatted using `{:?}`
to print its value; consider adding `#[derive(Debug)]` to
the struct definition, or manually implementing
`Debug` for `Rectangle`
```

Rust *does* include functionality to print out debug values, but we have to explicitly opt in to make that functionality available. We can do this by adding the annotation `#[derive(Debug)]` just before the struct definition. Listing 5-12:

Filename: src/main.rs

```
#[derive(Debug)]
struct Rectangle {
    width: u32,
    height: u32,
}

fn main() {
    let rect1 = Rectangle { width: 30, height: 50 };

    println!("rect1 is {:?}", rect1);
}
```

Listing 5-12: Adding the annotation to derive the `Rectangle` struct to use debug formatting

Now when we run the program, we won't get an error. Instead, we'll get the following output:

```
rect1 is Rectangle { width: 30, height: 50 }
```

Nice! It's not the prettiest output, but it shows the struct's value, which would definitely help during debugging. It's useful to have output that's a bit easier to read.

instead of `{:?}` in the `println!` string. When we run the example, the output will look like this:

```
rect1 is Rectangle {  
  width: 30,  
  height: 50  
}
```

Rust has provided a number of traits for us to use that we can add useful behavior to our custom types. They are listed in Appendix C, “Derivable Traits.” We’ll cover how to use custom behavior as well as how to create your own traits.

Our `area` function is very specific: it only computes the area of a `Rectangle`. It would be helpful to tie this behavior more closely to our `Rectangle` type. Let’s look at how we can turn the `area` function into an `area` method.

## Method Syntax

*Methods* are similar to functions: they’re declared with a name, they can have parameters and a return value, and they are run when they’re called from somewhere else. The main difference from functions is that they’re defined within the trait object, which we cover in Chapters 6 and 17. The first parameter is always `self`, which represents the object being called on.

## Defining Methods

Let’s change the `area` function that has a `Rectangle` parameter to instead make an `area` method defined on the `Rectangle` struct. Listing 5-13:

Filename: `src/main.rs`

```

#[derive(Debug)]
struct Rectangle {
    width: u32,
    height: u32,
}

impl Rectangle {
    fn area(&self) -> u32 {
        self.width * self.height
    }
}

fn main() {
    let rect1 = Rectangle { width: 30, height: 50 };

    println!(
        "The area of the rectangle is {} square feet",
        rect1.area()
    );
}

```

Listing 5-13: Defining an `area` method on the `Rectangle` struct

To define the function within the context of `Rectangle` (implementation) block. Then we move the `area` brackets and change the first (and in this case, only) parameter in the signature and everywhere within the body. In `main` function and passed `rect1` as an argument, we call the `area` method on our `Rectangle` instance. For the instance: we add a dot followed by the method name and arguments.

In the signature for `area`, we use `&self` instead of `self`. Rust knows the type of `self` is `Rectangle` due to the `impl Rectangle` context. Note that we still need to pass `self` as we did in `&Rectangle`. Methods can take ownership of `self` immutably as we've done here, or borrow `self` as a mutable parameter.

We've chosen `&self` here for the same reason as the previous version: we don't want to take ownership, and we want to read from the struct, not write to it. If we wanted to change the width of the rectangle as part of what the method does, we would use `self` as a parameter. Having a method that takes ownership of `self` as the first parameter is rare; this technique is usually used to transform `self` into something else and you want to keep the original instance after the transformation.

The main benefit of using methods instead of function syntax and not having to repeat the type of `self` in the organization. We've put all the things we can do in the `impl` block rather than making future users of `Rectangle` in various places in the library we provide.

---

## Where's the `->` Operator?

In C and C++, two different operators are used when you're calling a method on the object directly versus a method on a pointer to the object and need to dereference. In other words, if `object` is a pointer, `object->something()` is equivalent to `(*object).something()`.

Rust doesn't have an equivalent to the `->` operator. Instead, it has *automatic referencing and dereferencing* in places in Rust that has this behavior.

Here's how it works: when you call a method on a variable, Rust automatically adds in `&`, `&mut`, or `*` so `object.method()` is equivalent to `&object.method()`. In other words, the following are the same:

```
p1.distance(&p2);
(&p1).distance(&p2);
```

The first one looks much cleaner. This automatic referencing and dereferencing is possible because methods have a clear receiver—the type and name of a method, Rust can figure out whether you're reading (`&self`), mutating (`&mut self`), or consuming (`*self`). This makes borrowing implicit for method receivers, which is ownership ergonomic in practice.

---

## Methods with More Parameters

Let's practice using methods by implementing a function on a struct. This time, we want an instance of `Rectangle` and return `true` if the second `Rectangle` is contained within the first, otherwise it should return `false`. That is, we want a function `contains` shown in Listing 5-14, once we've defined the `Rectangle` struct.

Filename: src/main.rs

```
fn main() {  
    let rect1 = Rectangle { width: 30, height: 10 };  
    let rect2 = Rectangle { width: 10, height: 5 };  
    let rect3 = Rectangle { width: 60, height: 40 };  
  
    println!("Can rect1 hold rect2? {}", rect1.can_hold(&rect2));  
    println!("Can rect1 hold rect3? {}", rect1.can_hold(&rect3));  
}
```

Listing 5-14: Using the as-yet-unwritten `can_hold` method

And the expected output would look like the following. The `rect2` are smaller than the dimensions of `rect1`, so the output is `true`. The `rect3` are larger than the dimensions of `rect1`, so the output is `false`.

```
Can rect1 hold rect2? true  
Can rect1 hold rect3? false
```

We know we want to define a method, so it will be a `fn` definition. The method name will be `can_hold`, and it will take a `Rectangle` as a parameter. We can tell what the method is doing by looking at the code that calls the method: `rect1.can_hold(&rect2)`, which is an immutable borrow to `rect2`. This makes sense because we only need to read `rect2`'s dimensions (we wouldn't need a mutable borrow), and we want to use it again after calling the `can_hold` method. The `can_hold` method will be a Boolean, and the implementation will compare the width and height of `self` to the width and height of `other`, respectively. Let's add the new `can_hold` method to Listing 5-13, shown in Listing 5-15:

Filename: src/main.rs

```
impl Rectangle {  
    fn area(&self) -> u32 {  
        self.width * self.height  
    }  
  
    fn can_hold(&self, other: &Rectangle) -> bool {  
        self.width > other.width && self.height > other.height  
    }  
}
```

Listing 5-15: Implementing the `can_hold` method for the `Rectangle` struct



When we run this code with the `main` function it produces the following output. Methods can take multiple parameters, the first of which is the `self` parameter, and those parameters work just like the parameters of a function.

## Associated Functions

Another useful feature of `impl` blocks is that we can write `impl` blocks that *don't* take `self` as a parameter. These are called *associated functions* because they're associated with the struct, not with an instance of the struct. Methods, because they don't have an instance of the struct, already used the `String::from` associated function.

Associated functions are often used for constructing a struct. For example, we could provide an associated function that takes a single dimension parameter and use that as both width and height to create a square `Rectangle` rather than having to create a `Rectangle` with two parameters.

Filename: src/main.rs

```
impl Rectangle {
    fn square(size: u32) -> Rectangle {
        Rectangle { width: size, height: size }
    }
}
```

To call this associated function, we use the `::` syntax. For example, `let sq = Rectangle::square(3);` is an example of creating a `Rectangle` struct: the `::` syntax is used for both associated functions and methods. We'll discuss modules in Chapter 7.

## Multiple `impl` Blocks

Each struct is allowed to have multiple `impl` blocks. This is equivalent to the code shown in Listing 5-16, which shows a `Rectangle` struct with two `impl` blocks:

```
impl Rectangle {
    fn area(&self) -> u32 {
        self.width * self.height
    }
}

impl Rectangle {
    fn can_hold(&self, other: &Rectangle)
        self.width > other.width && self.h
```

Listing 5-16: Rewriting Listing 5-15 using multiple

There’s no reason to separate these methods in is valid syntax. We’ll see a case in which multiple where we discuss generic types and traits.

## Summary

Structs let you create custom types that are mea structs, you can keep associated pieces of data c each piece to make your code clear. Methods let instances of your structs have, and associated fu functionality that is particular to your struct with

But structs aren’t the only way you can create cu feature to add another tool to your toolbox.

## Enums and Pattern Ma

In this chapter we’ll look at *enumerations*, also re to define a type by enumerating its possible valu enum to show how an enum can encode meanin a particularly useful enum, called `Option`, which something or nothing. Then we’ll look at how pa expression makes it easy to run different code f we’ll cover how the `if let` construct is another available to you to handle enums in your code.

Enums are a feature in many languages, but the Rust’s enums are most similar to *algebraic data t*

F#, OCaml, and Haskell.

## Defining an Enum

Let's look at a situation we might want to express that is more useful and more appropriate than structs in this case: IP addresses. Currently, two major standards are in use: version four and version six. These are the only possibilities for IP addresses we will come across: we can *enumerate* all possible IP addresses and each gets its name.

Any IP address can be either a version four or a version six address at the same time. That property of IP addresses makes an enum more appropriate, because enum values can only be one of a set of values. Version four and version six addresses are still fundamentally different, but they are treated as the same type when the code is handling an IP address.

We can express this concept in code by defining an enum, listing the possible kinds an IP address can be, and then using the *variants* of the enum:

```
enum IpAddrKind {  
    V4,  
    V6,  
}
```

`IpAddrKind` is now a custom data type that we can use to represent IP addresses.

## Enum Values

We can create instances of each of the two variants of the enum:

```
let four = IpAddrKind.V4;  
let six = IpAddrKind.V6;
```

Note that the variants of the enum are namespace-qualified using the double colon to separate the two. The reason that `IpAddrKind.V4` and `IpAddrKind.V6` are of the type `IpAddrKind` is that they are of the type `IpAddrKind`. Then, for instance, define a function that takes a

```
fn route(ip_type: IpAddrKind) { }
```

And we can call this function with either variant:

```
route(IpAddrKind::V4);  
route(IpAddrKind::V6);
```

Using enums has even more advantages. Think of the moment we don't have a way to store the actual *kind* it is. Given that you just learned about this problem as shown in Listing 6-1:

```
enum IpAddrKind {  
    V4,  
    V6,  
}  
  
struct IpAddr {  
    kind: IpAddrKind,  
    address: String,  
}  
  
let home = IpAddr {  
    kind: IpAddrKind::V4,  
    address: String::from("127.0.0.1"),  
};  
  
let loopback = IpAddr {  
    kind: IpAddrKind::V6,  
    address: String::from("::1"),  
};
```

Listing 6-1: Storing the data and `IpAddrKind` variants

Here, we've defined a struct `IpAddr` that has two fields: `IpAddrKind` (the enum we defined previously) and `String`. We have two instances of this struct. The first, `home`, has its `kind` with associated address data of `127.0.0.1`. The second, `loopback`, has the other variant of `IpAddrKind` and the address `::1` associated with it. We've used a struct to store the variant and its associated data together, so now the variant is associated with its data.

We can represent the same concept in a more concise way than an enum inside a struct, by putting data directly inside the enum. A new definition of the `IpAddr` enum says that because

associated `String` values:

```
enum IpAddr {  
    V4(String),  
    V6(String),  
}  
  
let home = IpAddr::V4(String::from("127.0.  
  
let loopback = IpAddr::V6(String::from("::
```

We attach data to each variant of the enum directly.

There's another advantage to using an enum rather than having different types and amounts of associated data. We will always have four numeric components that we wanted to store `v4` addresses as four `u8` values, but if we stored `v6` as one `String` value, we wouldn't be able to write them with ease:

```
enum IpAddr {  
    V4(u8, u8, u8, u8),  
    V6(String),  
}  
  
let home = IpAddr::V4(127, 0, 0, 1);  
  
let loopback = IpAddr::V6(String::from("::
```

We've shown several different ways to define data for version four and version six IP addresses. However, as it turns out, the way to define and encode which kind they are is so common that there is a [definition we can use!](#) Let's look at how the standard library defines the exact enum and variants that we've defined, and how it stores data inside the variants in the form of two different ways for each variant:

```

struct Ipv4Addr {
    // --snip--
}

struct Ipv6Addr {
    // --snip--
}

enum IpAddr {
    V4(Ipv4Addr),
    V6(Ipv6Addr),
}

```

This code illustrates that you can put any kind of numeric types, or structs, for example. You can use standard library types are often not much more come up with.

Note that even though the standard library contains definitions for `Ipv4Addr` and `Ipv6Addr`, we still create and use our own definition without conflict. We bring the standard library's definition into our scope. We bring `IpAddr` into scope in Chapter 7.

Let's look at another example of an enum in Listing 6-2, where the variants are types embedded in its variants:

```

enum Message {
    Quit,
    Move { x: i32, y: i32 },
    Write(String),
    ChangeColor(i32, i32, i32),
}

```

Listing 6-2: A `Message` enum whose variants each have different values

This enum has four variants with different types

- `Quit` has no data associated with it at all.
- `Move` includes an anonymous struct inside.
- `Write` includes a single `String`.
- `ChangeColor` includes three `i32` values.

Defining an enum with variants like the ones in Listing 6-2 is similar to the different kinds of struct definitions, except the variants are grouped together under a single enum.

structs could hold the same data that the previous

```
struct QuitMessage; // unit struct
struct MoveMessage {
    x: i32,
    y: i32,
}
struct WriteMessage(String); // tuple struct
struct ChangeColorMessage(i32, i32, i32);
```

But if we used the different structs, which each have their own data, it's not easy to define a function to take any of these kinds of `Message` enum defined in Listing 6-2, which is a

There is one more similarity between enums and methods on structs using `impl`, we're also able to define a method named `call` that we could define on

```
impl Message {
    fn call(&self) {
        // method body would be defined here
    }
}

let m = Message::Write(String::from("hello"));
m.call();
```

The body of the method would use `self` to get the message. In this example, we've created a variable `m` of type `Message::Write(String::from("hello"))`, and then called the `call` method when `m.call()` runs.

Let's look at another enum in the standard library, `Option`.

## The `Option` Enum and Its Advantages

In the previous section, we looked at how the `Option` type is used in the Rust type system to encode more information than just the presence or absence of a value. In this section, we explore a case study of `Option`, which is another type in the standard library. The `Option` type is used in many places in the Rust standard library. In a scenario in which a value could be something or nothing, the concept in terms of the type system means the

handled all the cases you should be handling; they are extremely common in other programming languages.

Programming language design is often thought of as including features, but the features you exclude are important. One important feature that many other languages have is *Null*. In languages with null, variables can always be null.

In his 2009 presentation “Null References: The Billion Dollar Mistake”, the inventor of null, has this to say:

---

I call it my billion-dollar mistake. At that time, when I designed the comprehensive type system for references in C++, my goal was to ensure that all use of references was safe. The temptation to put in a null reference, simply because it was so easy to implement, has led to innumerable errors, crashes, and disasters, which have probably caused a billion dollars of loss in the last forty years.

---

The problem with null values is that if you try to use a null reference, you'll get an error of some kind. Because this null check is so extremely easy to make this kind of error.

However, the concept that null is trying to express is a concept that is currently invalid or absent for some reason.

The problem isn't really with the concept but with the implementation. For example, Rust does not have nulls, but it does have an enum of a value being present or absent. This enum is in the [standard library](#) as follows:

```
enum Option<T> {  
    Some(T),  
    None,  
}
```

The `Option<T>` enum is so useful that it's even in the `std::prelude` so you don't need to bring it into scope explicitly. In addition, you can use `Some` and `None` directly without the `Option::` prefix. It's a regular enum, and `Some(T)` and `None` are still valid values.



The `<T>` syntax is a feature of Rust we haven't touched yet. The `T` is a type parameter, and we'll cover generics in more detail later. The first thing we need to know is that `<T>` means the `Some` variable can hold any piece of data of any type. Here are some examples using number types and string types:

```
let some_number = Some(5);
let some_string = Some("a string");

let absent_number: Option<i32> = None;
```

If we use `None` rather than `Some`, we need to tell the compiler what type we have, because the compiler can't infer the type by looking only at a `None` value.

When we have a `Some` value, we know that a valid value exists within the `Some`. When we have a `None` value, it's the same thing as null: we don't have a valid value. So why do we have null?

In short, because `Option<T>` and `T` (where `T` is a type) are different types. The compiler won't let us use an `Option<T>` value where a `T` is expected. For example, this code won't compile because it tries to add an `Option<i8>` to an `i8`:

```
let x: i8 = 5;
let y: Option<i8> = Some(5);

let sum = x + y;
```

If we run this code, we get an error message like the following:

```
error[E0277]: the trait bound `i8: std::ops::Add` is not satisfied
-->
   |
5  |         let sum = x + y;
   |                     ^ no implementation for i8 + Option<i8>
   |
```

Intense! In effect, this error message means that the compiler can't add an `i8` and an `Option<i8>`, because they're different types. In a type like `i8` in Rust, the compiler will ensure that it can proceed confidently without having to check for null.

when we have an `Option<i8>` (or whatever type) we have to worry about possibly not having a value, and handle that case before using the value.

In other words, you have to convert an `Option<T>` to a `T` before you can do any operations with it. Generally, this helps catch on null: assuming that something isn't null when it is.

Not having to worry about incorrectly assuming a value is present is a confident in your code. In order to have a value that you can explicitly opt in by making the type of that value `Option<T>`, you are required to explicitly handle the `None` case. Everywhere that a value has a type that isn't an `Option<T>`, you know that the value isn't null. This was a deliberate design decision to increase the pervasiveness and increase the safety of Rust code.

So, how do you get the `T` value out of a `Some` value? The `Option<T>` methods that are useful in a variety of situations are listed in the [Option documentation](#). Becoming familiar with the methods is useful in your journey with Rust.

In general, in order to use an `Option<T>` value, you have to handle each variant. You want some code that works for the `Some` value, and this code is allowed to use the inner `T` value; if you have a `None` value, and that code doesn't handle it, you have a compile-time error. The `match` expression is a control flow construct that does exactly this: it runs different code depending on which variant of `Option<T>` you use the data inside the matching value.

## The `match` Control Flow Operator

Rust has an extremely powerful control flow operator called `match`. It compares a value against a series of patterns and executes the code for the pattern matches. Patterns can be made up of literals, variables, and many other things; Chapter 18 covers all the details of what they do. The power of `match` comes from the fact that the compiler confirms that all possible values are covered.

Think of a `match` expression as being like a coin slot machine with variously sized holes along it, and each hole represents a pattern that it fits into. In the same way, values are compared against the patterns, and the code for the matching pattern is executed.

`match`, and at the first pattern the value “fits,” the block to be used during execution.

Because we just mentioned coins, let’s use them to write a function that can take an unknown `Unit` in the counting machine, determine which coin it is, and return the value. This is shown here in Listing 6-3:

```
enum Coin {
    Penny,
    Nickel,
    Dime,
    Quarter,
}

fn value_in_cents(coin: Coin) -> u32 {
    match coin {
        Coin::Penny => 1,
        Coin::Nickel => 5,
        Coin::Dime => 10,
        Coin::Quarter => 25,
    }
}
```

Listing 6-3: An enum and a `match` expression with patterns

Let’s break down the `match` in the `value_in_cents` function. The `match` keyword is followed by an expression, which in this case is `coin`. This is very similar to an expression used with `if`, but the expression needs to return a Boolean value, but `coin` in this example is the `Coin` enum that we defined in Listing 6-2.

Next are the `match` arms. An arm has two parts: a pattern and a code block. The pattern here has a value `Coin::Penny`. The `::` separates the pattern and the code to run. The code block is `=> 1`. Each arm is separated from the next with a comma.

When the `match` expression executes, it compares the value to the pattern of each arm, in order. If a pattern matches, that pattern is executed. If that pattern doesn’t match, it moves to the next arm, much as in a coin-sorting machine. In Listing 6-3, our `match` has four arms.

The code associated with each arm is an expression. The expression in the matching arm is the value that

expression.

Curly brackets typically aren't used if the match is a block where each arm just returns a value. If you want to use a match arm, you can use curly brackets. For example, "Lucky penny!" every time the method was called with a Penny, return the last value of the block, `1`:

```
fn value_in_cents(coin: Coin) -> u32 {
    match coin {
        Coin::Penny => {
            println!("Lucky penny!");
            1
        },
        Coin::Nickel => 5,
        Coin::Dime => 10,
        Coin::Quarter => 25,
    }
}
```

## Patterns that Bind to Values

Another useful feature of match arms is that they can bind to values that match the pattern. This is how we can extract data from the matched value.

As an example, let's change one of our enum variants to include a value. From 1999 through 2008, the United States minted quarters with the names of the 50 states on one side. No other coins got this extra value. We can add this information to the Quarter variant to include a `UsState` value stored inside it. Here's the code:

```
#[derive(Debug)] // So we can inspect the
enum UsState {
    Alabama,
    Alaska,
    // --snip--
}

enum Coin {
    Penny,
    Nickel,
    Dime,
    Quarter(UsState),
}
```

Listing 6-4: A `Coin` enum in which the `Quarter`

Let's imagine that a friend of ours is trying to collect our loose change by coin type, we'll also call it with each quarter so if it's one our friend doesn't have a collection.

In the match expression for this code, we add an arm that matches values of the variant `Coin::Quarter`, the `state` variable will bind to the value of that `state` in the code for that arm, like so:

```
fn value_in_cents(coin: Coin) -> u32 {
    match coin {
        Coin::Penny => 1,
        Coin::Nickel => 5,
        Coin::Dime => 10,
        Coin::Quarter(state) => {
            println!("State quarter from {}", state),
        },
    }
}
```

If we were to call `value_in_cents(Coin::Quarter(UsState::Alaska))`. When we reach the match arms, none of them match until we reach the point, the binding for `state` will be the value `UsState::Alaska` that binding in the `println!` expression, thus giving us the value of the `Coin` enum variant for `Quarter`.

## Matching with `Option<T>`

In the previous section, we wanted to get the inner value when using `Option<T>`; we can also handle `Option` with the `Coin` enum! Instead of comparing coins, we can match on `Option`, but the way that the `match` expression works is a bit different.

Let's say we want to write a function that takes a `Option<u32>` inside, adds 1 to that value. If there isn't a value (i.e. `None`), we return `None` and not attempt to perform any operations.

This function is very easy to write, thanks to `match` and `Option`.

```
fn plus_one(x: Option<i32>) -> Option<i32>
    match x {
        None => None,
        Some(i) => Some(i + 1),
    }
}

let five = Some(5);
let six = plus_one(five);
let none = plus_one(None);
```

Listing 6-5: A function that uses a `match` expression

Let's examine the first execution of `plus_one` in `plus_one(five)`, the variable `x` in the body of `Some(5)`. We then compare that against each match arm.

```
None => None,
```

The `Some(5)` value doesn't match the pattern `None`.

```
Some(i) => Some(i + 1),
```

Does `Some(5)` match `Some(i)`? Why yes it does! `Some` binds to the value contained in `Some`, so `i` takes the value 5. The second arm is then executed, so we add 1 to the value contained in `Some` with our total 6 inside.

Now let's consider the second call of `plus_one` in `plus_one(none)`. We enter the `match` and compare to the first arm.

```
None => None,
```

It matches! There's no value to add to, so the procedure returns the value on the right side of `=>`. Because the first arm matched, no further comparison is needed.

Combining `match` and enums is useful in many situations. In Rust code: `match` against an enum, bind a variable to the value, and then execute code based on it. It's a bit tricky at first, but once you get it, you'll find it in all languages. It's consistently a useful feature.

## Matches Are Exhaustive

There's one other aspect of `match` we need to discuss: the `plus_one` function that has a bug and won't compile.

```
fn plus_one(x: Option<i32>) -> Option<i32> {
    match x {
        Some(i) => Some(i + 1),
    }
}
```

We didn't handle the `None` case, so this code won't compile. Rust knows how to catch this. If we try to compile this code, we get the following error:

```
error[E0004]: non-exhaustive patterns: `None` not covered
-->
   |
6  |         match x {
   |                 ^ pattern `None` not covered
```

Rust knows that we didn't cover every possible case because we forgot! Matches in Rust are *exhaustive*: we must handle every possible case in order for the code to be valid. Especially in the context of pattern matching, this prevents us from forgetting to explicitly handle a case. It forces us to be explicit, preventing us from assuming that we have a value when we might have a `None` value, a mistake discussed earlier.

## The `_` Placeholder

Rust also has a pattern we can use when we don't want to handle a specific case. For example, a `u8` can have valid values of 0 through 255. If we want to handle values 1, 3, 5, and 7, we don't want to have to list all the other values. Fortunately, we don't have to: we can use the `_` placeholder.

```
let some_u8_value = 0u8;
match some_u8_value {
    1 => println!("one"),
    3 => println!("three"),
    5 => println!("five"),
    7 => println!("seven"),
    _ => (),
}
```

The `_` pattern will match any value. By putting it at the end of the `match` block, it matches all the possible cases that aren't specified. Since we don't have a value, so nothing will happen in the `_` case. As a result, the code will compile successfully.

nothing for all the possible values that we don't

However, the `match` expression can be a bit too care about *one* of the cases. For this situation, R

## Concise Control Flow with `if`

The `if let` syntax lets you combine `if` and `let` values that match one pattern while ignoring the 6-6 that matches on an `Option<u8>` value but o is 3:

```
let some_u8_value = Some(0u8);
match some_u8_value {
    Some(3) => println!("three"),
    _ => (),
}
```

Listing 6-6: A `match` that only cares about execu

We want to do something with the `Some(3)` ma `Some<u8>` value or the `None` value. To satisfy th `_ => ()` after processing just one variant, which

Instead, we could write this in a shorter way usir behaves the same as the `match` in Listing 6-6:

```
if let Some(3) = some_u8_value {
    println!("three");
}
```

The syntax `if let` takes a pattern and an expres the same way as a `match`, where the expression pattern is its first arm.

Using `if let` means you have less typing, less i However, you lose the exhaustive checking that `match` and `if let` depends on what you're doi whether gaining conciseness is an appropriate t checking.

In other words, you can think of `if let` as synt



when the value matches one pattern and then if

We can include an `else` with an `if let`. The block is the same as the block of code that would go with the expression that is equivalent to the `if let` and definition in Listing 6-4, where the `Quarter` variable wanted to count all non-quarter coins we see with quarters, we could do that with a `match` expression:

```
let mut count = 0;
match coin {
    Coin::Quarter(state) => println!("State quarter from {:?}", state),
    _ => count += 1,
}
```

Or we could use an `if let` and `else` expression:

```
let mut count = 0;
if let Coin::Quarter(state) = coin {
    println!("State quarter from {:?}", state);
} else {
    count += 1;
}
```

If you have a situation in which your program has a variable using a `match`, remember that `if let` is in your toolbox.

## Summary

We've now covered how to use enums to create types of enumerated values. We've shown how the static type system you use the type system to prevent errors. When you can use `match` or `if let` to extract and use values in many cases you need to handle.

Your Rust programs can now express concepts using enums. Creating custom types to use in your API will make certain your functions get only values you need.

In order to provide a well-organized API to your users and only exposes exactly what your users will need.

# Using Modules to Reuse Code

When you start writing programs in Rust, your code grows. As your code grows, you'll eventually need to organize it for reuse and better organization. By splitting your code into modules, you can make each chunk easier to understand on its own. How do you organize many functions? Rust has a module system that organizes code in an organized fashion.

In the same way that you extract lines of code into functions (and other code, like structs and enums), you can create a namespace that contains definitions of functions. Whether those definitions are visible outside the module is up to you. Here's an overview of how modules work:

- The `mod` keyword declares a new module. It is either immediately following this declaration in the file.
- By default, functions, types, constants, and other items are private. The `pub` keyword makes an item public and therefore visible outside the module.
- The `use` keyword brings modules, or the contents of a module, into scope so it's easier to refer to them.

We'll look at each of these parts to see how they work.

## `mod` and the Filesystem

We'll start our module example by making a new project. When creating a binary crate, we'll make a library crate a dependency of our project. For example, in Chapter 2 is a library crate that we used as a dependency of our project.

We'll create a skeleton of a library that provides some functionality; we'll concentrate on the organization of the code. We won't worry about what code goes in the function `communicator`. To create a library, pass the `--lib` flag to `cargo new`.

```
$ cargo new communicator --lib
$ cd communicator
```

Notice that Cargo generated *src/lib.rs* instead of the following:

Filename: *src/lib.rs*

```
#[cfg(test)]
mod tests {
    #[test]
    fn it_works() {
        assert_eq!(2 + 2, 4);
    }
}
```

Cargo creates an example test to help us get our feet wet with the `mod tests` syntax in the “Using `super` to Access Parent Modules” section of this chapter, but for now, leave this code at the end of the file.

Because we don’t have a *src/main.rs* file, there’s no `cargo run` command. Therefore, we’ll use the `cargo test` command to run our library crate’s code.

We’ll look at different options for organizing your code in a variety of situations, depending on the intent of the crate.

## Module Definitions

For our `communicator` networking library, we’ll first define a module that contains the definition of a function called `connect`. Rust starts with the `mod` keyword. Add this code above the test code:

Filename: *src/lib.rs*

```
mod network {
    fn connect() {
    }
}
```

After the `mod` keyword, we put the name of the module in curly brackets. Everything inside this block is part of the module.

`network` . In this case, we have a single function, function from code outside the `network` module and use the namespace syntax `::` like :

We can also have multiple modules, side by side example, to also have a `client` module that ha add it as shown in Listing 7-1:

Filename: `src/lib.rs`

```
mod network {  
    fn connect() {  
    }  
}  
  
mod client {  
    fn connect() {  
    }  
}
```

Listing 7-1: The `network` module and the `client` module in `src/lib.rs`

Now we have a `network::connect` function and can have completely different functionality, and with each other because they're in different modules.

In this case, because we're building a library, the building our library is `src/lib.rs`. However, in response to nothing special about `src/lib.rs`. We could also create a binary crate in the same way as we're creating a library crate. In fact, we can put modules inside of modules. As modules grow to keep related functionality organized, we can use modules to keep related functionality apart. The way you choose to organize your code and its `connect` function might make more sense if we were inside the `network` namespace instead, as

Filename: `src/lib.rs`

```

mod network {
    fn connect() {
    }

    mod client {
        fn connect() {
        }
    }
}

```

Listing 7-2: Moving the `client` module inside the `network` module

In your `src/lib.rs` file, replace the existing `mod network` with the ones in Listing 7-2, which have the `client` module nested inside `network`. The functions `network::connect` and `network::client::connect` are both named `connect`, but they don't conflict with each other because they live in different namespaces.

In this way, modules form a hierarchy. The container module is at the top level, and the submodules are at lower levels. How the hierarchy in Listing 7-1 looks like when thought of as a tree diagram:

```

communicator
├── network
└── client

```

And here's the hierarchy corresponding to the example in Listing 7-2:

```

communicator
└── network
    └── client

```

The hierarchy shows that in Listing 7-2, `client` is a submodule of `network` rather than a sibling. More complicated projects need to be organized logically in order for you to find what you need. The way you organize your project is up to you and depends on the domain you're thinking about your project's domain. Use the techniques of side modules and nested modules in whatever way makes sense for your project.

## Moving Modules to Other Files

Modules form a hierarchical structure, much like the filesystems you're used to: filesystems! We can use Rust's modules to split up Rust projects so not everything lives in `src/lib.rs`.

example, let's start with the code in Listing 7-3:

Filename: `src/lib.rs`

```
mod client {
    fn connect() {
    }
}

mod network {
    fn connect() {
    }

    mod server {
        fn connect() {
        }
    }
}
```

Listing 7-3: Three modules, `client`, `network`, and `server`, in `src/lib.rs`

The file `src/lib.rs` has this module hierarchy:

```
communicator
├── client
├── network
│   └── server
```

If these modules had many functions, and those functions were nested inside one or more other modules, the code would be difficult to scroll through this file to find a specific function. Because the functions are nested inside one or more other modules, the code inside the functions will start getting lengthy as we try to access functions from separate modules. To separate the `client`, `network`, and `server` modules into their own files.

First, let's replace the `client` module code with its own file so that `src/lib.rs` looks like code shown in Listing 7-4.

Filename: `src/lib.rs`

```
mod client;

mod network {
    fn connect() {
    }

    mod server {
        fn connect() {
        }
    }
}
```

Listing 7-4: Extracting the contents of the `client` in `src/lib.rs`

We're still *declaring* the `client` module here, but with a semicolon, we're telling Rust to look in another location for the scope of the `client` module. In other words:

```
mod client {
    // contents of client.rs
}
```

Now we need to create the external file with that function in your `src/` directory and open it. Then enter the function in the `client` module that we removed.

Filename: `src/client.rs`

```
fn connect() {
}
```

Note that we don't need a `mod` declaration in the `client` module with `mod` in `src/lib.rs`. This file is part of the `client` module. If we put a `mod client` here, we're creating our own submodule named `client`!

Rust only knows to look in `src/lib.rs` by default. If we move the function to a new file, we need to tell Rust in `src/lib.rs` to look in that file. The function needs to be defined in `src/lib.rs` and can't be defined elsewhere.

Now the project should compile successfully, although it's a bit slower. Remember to use `cargo build` instead of `cargo run` to build a crate rather than a binary crate:

```

$ cargo build
   Compiling communicator v0.1.0 (file:///
warning: function is never used: `connect`
--> src/client.rs:1:1
|
1 | / fn connect() {
2 | | }
  | |_^
  |
  = note: #[warn(dead_code)] on by default

warning: function is never used: `connect`
--> src/lib.rs:4:5
|
4 | /      fn connect() {
5 | |      }
  | |_____^

warning: function is never used: `connect`
--> src/lib.rs:8:9
|
8 | /      fn connect() {
9 | |      }
  | |_____^

```

These warnings tell us that we have functions that are never used. We'll address these warnings for now; we'll address them later in the "Visibility with `pub`" section. The good news is that the build was successful!

Next, let's extract the `network` module into its own file. In `src/lib.rs`, delete the body of the `network` module declaration, like so:

Filename: `src/lib.rs`

```

mod client;

mod network;

```

Then create a new `src/network.rs` file and enter the following code:

Filename: `src/network.rs`



```
fn connect() {
}

mod server {
    fn connect() {
    }
}
```

Notice that we still have a `mod` declaration with `server`. We still want `server` to be a submodule of `network`.

Run `cargo build` again. Success! We have one module. Because it's a submodule—that is, a module with its own file—extracting a module into a file named after that module so you can see the error. First, change `src/network.rs` to the `server` module's contents:

Filename: `src/network.rs`

```
fn connect() {
}

mod server;
```

Then create a `src/server.rs` file and enter the contents that were extracted:

Filename: `src/server.rs`

```
fn connect() {
}
```

When we try to `cargo build`, we'll get the error

```

$ cargo build
   Compiling communicator v0.1.0 (file:///
error: cannot declare a new module at this
--> src/network.rs:4:5
   |
4  | mod server;
   |         ^^^^^^^
note: maybe move this module `src/network.`
`src/network/mod.rs`
--> src/network.rs:4:5
   |
4  | mod server;
   |         ^^^^^^^
note: ... or maybe `use` the module `serve
redeclaring it
--> src/network.rs:4:5
   |
4  | mod server;
   |         ^^^^^^^

```

Listing 7-5: Error when trying to extract the `server` module

The error says we `cannot declare a new module` to the `mod server;` line in `src/network.rs`. So `src` somehow: keep reading to understand why.

The note in the middle of Listing 7-5 is actually v something we haven't yet talked about doing:

```

note: maybe move this module `network` to
`network/mod.rs`

```

Instead of continuing to follow the same file-nar can do what the note suggests:

1. Make a new *directory* named `network`, the p
2. Move the `src/network.rs` file into the new `ne`  
`src/network/mod.rs`.
3. Move the submodule file `src/server.rs` into t

Here are commands to carry out these steps:

```

$ mkdir src/network
$ mv src/network.rs src/network/mod.rs
$ mv src/server.rs src/network

```

Now when we try to run `cargo build`, compilat

though). Our module layout still looks exactly the same as the code in *src/lib.rs* in Listing 7-3:

```
communicator
├── client
├── network
└── server
```

The corresponding file layout now looks like this

```
└─ src
    ├── client.rs
    ├── lib.rs
    └── network
        ├── mod.rs
        └── server.rs
```

So when we wanted to extract the `network::server` module, we had to change the `src/network.rs` file to the `src/network/lib.rs` file. We added the `network::server` in the `network` directory in `src`. We just extract the `network::server` module into `src/network/lib.rs`. We wouldn't be able to recognize that `server` was a module of `network` if the `server.rs` file was in the `src` directory. So let's consider a different example with the following definitions are in `src/lib.rs`:

```
communicator
├── client
└── network
    └── client
```

In this example, we have three modules again: `network::client`. Following the same steps we put into files, we would create `src/client.rs` for the `client` module, we would create `src/network.rs`. But we would move the `network::client` module into a `src/client.rs` file and create a top-level `client` module! If we could put the code for the `network::client` modules in the `src/client.rs` file, we would have whether the code was for `client` or for `network`.

Therefore, in order to extract a file for the `network` module, we needed to create a directory `src/network` and a `src/network.rs` file. The code that is in the `network/mod.rs` file, and the submodule `network/client.rs` file. Now the top-level `src/client.rs` file can use the `network` module.

belongs to the `client` module.

## Rules of Module Filesystems

Let's summarize the rules of modules with regard to their filesystems:

- If a module named `foo` has no submodules, you should place the module's code in a file named `foo.rs`.
- If a module named `foo` does have submodules, you should place the module's code in a file named `foo/mod.rs`.

These rules apply recursively, so if a module named `bar` and `bar` does not have submodules, you should place its code in the `src` directory:

```
└─ foo
   └─ bar.rs (contains the declarations for bar)
      └─ mod.rs (contains the declarations for bar's submodules)
```

The modules should be declared in their parent module's scope.

Next, we'll talk about the `pub` keyword and get into how it affects module visibility.

## Controlling Visibility with `pub`

We resolved the error messages shown in Listing 10-1 by moving the `network::server` code into the `src/network/mod.rs` file and the `client::connect` code into the `src/client/mod.rs` file respectively. At that point, `cargo build` was able to compile the program, but it still produced warning messages about the `client::connect`, `network::server::connect` functions not being used.

So why are we receiving these warnings? After all, these functions are intended to be used by our `us` program, so it shouldn't matter that these connect functions are not used. The point of creating them is that they will be used by the us` program.`

To understand why this program invokes these `connect` functions, we need to look at the `communicator` library from another project, called `communicator`. To create a binary crate in the same directory as our `us` program, we need to create a file containing this code:`

Filename: `src/main.rs`

```
extern crate communicator;

fn main() {
    communicator::client::connect();
}
```

We use the `extern crate` command to bring the scope. Our package now contains *two* crates. Call it a binary crate, which is separate from the existing `src/lib.rs`. This pattern is quite common for executing a library crate, and the binary crate uses that library crate, and it's a nice separation.

From the point of view of a crate outside the `co` modules we've been creating are within a module crate, `communicator`. We call the top-level mod

Also note that even if we're using an external crate, the `extern crate` should go in our root module. In our submodules, we can refer to items from external level modules.

Right now, our binary crate just calls our library's module. However, invoking `cargo build` will not

[illegible]

Ah ha! This error tells us that the `client` module has some unused functions. It's also the first time we've run into the `unused` warning in the context of Rust. The default state of all code in Rust is `pub`. If you don't use a public function in your program, the compiler will warn you that the function is unused. If you don't use a private function, the compiler will warn you that the function has gone unused.

After you specify that a function such as `client` call to that function from your binary crate be all function is unused will go away. Marking a function will be used by code outside of your project external usage that's now possible as the function is marked public, Rust will not require that will stop warning that the function is unused.

## Making a Function Public

To tell Rust to make a function public, we add the `pub` declaration. We'll focus on fixing the warning that `main` is gone unused for now, as well as the `module `client` not found for binary crate`. Modify `src/lib.rs` to make the `client` module public.

Filename: src/lib.rs

```
pub mod client;

mod network;
```

The **pub** keyword is placed right before **mod** . Le

[illegible]

Hooray! We have a different error! Yes, different celebration. The new error shows `function `cc src/client.rs` to make client::connect public to`

Filename: src/client.rs

```
pub fn connect() {  
}
```

Now run `cargo build` again:

```
warning: function is never used: `connect`
--> src/network/mod.rs:1:1
|
1 | / fn connect() {
2 | | }
  | | ^
  |
= note: #[warn(dead_code)] on by default

warning: function is never used: `connect`
--> src/network/server.rs:1:1
|
1 | / fn connect() {
2 | | }
  | | ^
```

The code compiled, and the warning that `client`

Unused code warnings don't always indicate that code was made public: if you *didn't* want these functions to be made public, code warnings could be alerting you to code you should delete. They could also be alerting you to a bug in all places within your library where this function

But in this case, we *do* want the other two functions to be part of the public API, so let's mark them as `pub` as well to get rid of the warnings. `src/network/mod.rs` to look like the following:

Filename: `src/network/mod.rs`

```
pub fn connect() {  
}
```

```
mod server;
```

Then compile the code:

```
warning: function is never used: `connect`  
--> src/network/mod.rs:1:1  
  |  
1 | / pub fn connect() {  
2 | | }  
  | |_^  
  |  
  = note: #[warn(dead_code)] on by default
```

```
warning: function is never used: `connect`  
--> src/network/server.rs:1:1  
  |  
1 | / fn connect() {  
2 | | }  
  | |_^
```

Hmmm, we're still getting an unused function warning. The reason is that `network::connect` is set to `pub`. The reason is that the `network` module is not public, but the `network` module that the function is defined in is public. We need to change the function to be private, like so:

Filename: `src/lib.rs`

```
pub mod client;
```

```
pub mod network;
```

Now when we compile, that warning is gone:

```
warning: function is never used: `connect`
--> src/network/server.rs:1:1
   |
1  | / fn connect() {
2  | | }
   | | ^
   |
   = note: #[warn(dead_code)] on by default
```

Only one warning is left—try to fix this one on your

## Privacy Rules

Overall, these are the rules for item visibility:

- If an item is public, it can be accessed through any of the parent's child modules.
- If an item is private, it can be accessed only within any of the parent's child modules.

## Privacy Examples

Let's look at a few more privacy examples to get started. Open a new terminal window, create a new project and enter the code in Listing 7-6 into your

Filename: src/lib.rs



```

mod outermost {
    pub fn middle_function() {}

    fn middle_secret_function() {}

    mod inside {
        pub fn inner_function() {}

        fn secret_function() {}
    }
}

fn try_me() {
    outermost::middle_function();
    outermost::middle_secret_function();
    outermost::inside::inner_function();
    outermost::inside::secret_function();
}

```

Listing 7-6: Examples of private and public functions

Before you try to compile this code, make a guess: which functions will have errors. Then, try compiling the code. The code is right—and read on for the discussion of the errors.

## Looking at the Errors

The `try_me` function is in the root module of our program. The `outermost` module is private, but the second privacy rule allows the root module to access the `outermost` module because it is the parent module, as is `try_me`.

The call to `outermost::middle_function` will work. The `middle_function` is public and `try_me` is accessing `middle_function` through the `outermost` module. We determined in the previous paragraph that `outermost` is accessible.

The call to `outermost::middle_secret_function` will fail. Because `middle_secret_function` is private, the only way to access it is through the module it is defined in. The `middle` module is neither the current module of `middle_secret_function` nor is it a child module of the current module of `middle_secret_function`.

The module named `inside` is private and has no access to the root module only by its current module `outermost`. That means that `try_me` is not allowed to call `outermost::inside::inner_function` or `outermost::inside::secret_function`.

## Fixing the Errors

Here are some suggestions for changing the code. Make a guess as to whether it will fix the errors. Compile the code to see whether or not you're right and understand why. Feel free to design more experiments.

- What if the `inside` module were public?
- What if `outermost` were public and `inside` private?
- What if, in the body of `inner_function`, you wrote `::outermost::middle_secret_function()` instead of `middle_secret_function()`? This would mean that we want to refer to the module's function.

Next, let's talk about bringing items into scope with the `use` keyword.

## Referring to Names in Different Scopes

We've covered how to call functions defined with `pub` as part of the call, as in the call to the `nested_modules` function in Listing 7-7:

Filename: `src/main.rs`

```
pub mod a {
    pub mod series {
        pub mod of {
            pub fn nested_modules() {}
        }
    }
}

fn main() {
    a::series::of::nested_modules();
}
```

Listing 7-7: Calling a function by fully specifying its path

As you can see, referring to the fully qualified name of the function is a bit verbose. Rust has a keyword to make these calls more concise.

## Bringing Names into Scope with the `use` Keyword

Rust's `use` keyword shortens lengthy function calls by bringing the function name into scope.

function you want to call into scope. Here's an example of bringing the `a::series::of` module into a binary crate's `src/main.rs`

Filename: `src/main.rs`

```
pub mod a {
    pub mod series {
        pub mod of {
            pub fn nested_modules() {}
        }
    }
}

use a::series::of;

fn main() {
    of::nested_modules();
}
```

The line `use a::series::of;` means that rather than using the full path wherever we want to refer to the `of` module, we can just use `of`.

The `use` keyword brings only what we've specified into scope. That's why we still need to use the full path when we want to call the `nested_modules` function.

We could have chosen to bring the function into scope directly in the `use` as follows:

```
pub mod a {
    pub mod series {
        pub mod of {
            pub fn nested_modules() {}
        }
    }
}

use a::series::of::nested_modules;

fn main() {
    nested_modules();
}
```

Doing so allows us to exclude all the modules and functions that we don't need.

Because enums also form a sort of namespace, we can bring their variants into scope with `use` as well. For any kind of item, we can bring multiple items from one namespace into scope, separated by commas in the last position, like so:

```
enum TrafficLight {
    Red,
    Yellow,
    Green,
}

use TrafficLight::{Red, Yellow};

fn main() {
    let red = Red;
    let yellow = Yellow;
    let green = TrafficLight::Green;
}
```

We're still specifying the `TrafficLight` namespace didn't include `Green` in the `use` statement.

## Nested groups in `use` declarations

If you have a complex module tree with many different modules, and you want to import a few items from each one, it might be useful to use a single `use` declaration to keep your code clean and avoid repetition.

The `use` declaration supports nesting to help you organize your imports and glob ones. For example this snippet imports `baz` and `Bar`:

```
use foo::{
    bar::{self, Foo},
    baz::{*, quux::Bar},
};
```

## Bringing All Names into Scope with a `glob`

To bring all the items in a namespace into scope, you can use the `glob` operator, which is called the *glob operator*. This example brings all items from the `foo` namespace into scope without having to list each one specifically:

```
enum TrafficLight {
    Red,
    Yellow,
    Green,
}

use TrafficLight::*;

fn main() {
    let red = Red;
    let yellow = Yellow;
    let green = Green;
}
```

The `*` will bring into scope all the visible items in the module. You should use globs sparingly: they are convenient, but they can bring in more items than you expected and cause naming conflicts.

## Using `super` to Access a Parent Module

As you saw at the beginning of this chapter, we created a module named `communicator` and made a `tests` module for you. Let's go into the `communicator` project, open `src/lib.rs`:

Filename: `src/lib.rs`

```
pub mod client;

pub mod network;

#[cfg(test)]
mod tests {
    #[test]
    fn it_works() {
        assert_eq!(2 + 2, 4);
    }
}
```

Chapter 11 explains more about testing, but for now: we have a module named `tests` that lives inside the `communicator` module. It contains one function named `it_works`. Even though the `tests` module is just another module! So on

```

communicator
├── client
├── network
│   └── client
└── tests

```

Tests are for exercising the code within our library. We'll use the `client::connect` function from this `it_works` test to check any functionality right now. This won't work yet.

Filename: `src/lib.rs`

```

#[cfg(test)]
mod tests {
    #[test]
    fn it_works() {
        client::connect();
    }
}

```

Run the tests by invoking the `cargo test` command.

```

$ cargo test
   Compiling communicator v0.1.0 (file:///...)
error[E0433]: failed to resolve. Use of undeclared type `client`
  --> src/lib.rs:9:9
   |
9  |         client::connect();
   |         ^^^^^^^ Use of undeclared type `client`

```

The compilation failed, but why? We don't need to import the `client` module, as we did in `src/main.rs`, because we're using the `communicator` library crate here. The reason is that `client` is not in the current module, which here is `tests`. The only place where paths are relative to the crate root by default is in the `client` module in its scope!

So how do we get back up one module in the module hierarchy to use the `client::connect` function in the `tests` module? We can either use leading colons to let Rust know that we want to use the whole path, like this:

```

::client::connect();

```

Or, we can use `super` to move up one module in the hierarchy, like this:

```
super::client::connect();
```

These two options don't look that different in the module hierarchy, starting from the root every time. In those cases, using `super` to get from the current module is a shortcut. Plus, if you've specified the path from the root and then rearrange your modules by moving a module, you'll need to update the path in several places, which is annoying.

It would also be annoying to have to type `super::client::connect()` seen the tool for that solution: `use`! The `super::` gives you a shortcut to `use` so it is relative to the parent module.

For these reasons, in the `tests` module especially, `use super::client::connect();` is the best solution. So now our test looks like this:

Filename: src/lib.rs

```
#[cfg(test)]
mod tests {
    use super::client;

    #[test]
    fn it_works() {
        client::connect();
    }
}
```

When we run `cargo test` again, the test will pass and the output will be the following:

```
$ cargo test
   Compiling communicator v0.1.0 (file:///.../communicator)
   Running target/debug/communicator-926...

running 1 test
test tests::it_works ... ok

test result: ok. 1 passed; 0 failed; 0 ignored
```

## Summary

Now you know some new techniques for organizing code. To group related functionality together, keep files in a module and use `super::` to refer to the parent module.

present a tidy public API to your library users.

Next, we'll look at some collection data structure use in your nice, neat code.

## Common Collections

Rust's standard library includes a number of very useful *collections*. Most other data types represent one value, but collections contain multiple values. Unlike the built-in array, collections point to memory stored on the heap, which means they need to be known at compile time and can grow. Each kind of collection has different capabilities and characteristics, and choosing one for your current situation is a skill you'll develop over time. We'll discuss three collections that are used very often:

- A *vector* allows you to store a variable number of values.
- A *string* is a collection of characters. We've discussed it previously, but in this chapter we'll talk about it more.
- A *hash map* allows you to associate a value with a key. It's an implementation of the more general data structure called a *map*.

To learn about the other kinds of collections provided by the standard library, see the [collections documentation](#).

We'll discuss how to create and update vectors, strings, and hash maps, and what makes each special.

## Storing Lists of Values with Vectors

The first collection type we'll look at is `Vec<T>`, a *vector*. It allows you to store more than one value in a single data structure. The values are stored next to each other in memory. Vectors can only store one type of value. They are useful when you have a list of items, such as a list of items in a shopping cart.

### Creating a New Vector

To create a new, empty vector, we can call the `Vec::new` method. Here's an example: 8-1:



```
let v: Vec<i32> = Vec::new();
```

Listing 8-1: Creating a new, empty vector to hold

Note that we added a type annotation here. Because we're putting `i32` values into this vector, Rust doesn't know what kind of elements it will hold. An important point. Vectors are implemented using generics with your own types in Chapter 10. For now, the standard library can hold any type. If you want a vector of a specific type, the type is specified within angle brackets. In `v`, that the `Vec<T>` will hold elements of the type `T`.

In more realistic code, Rust can often infer the type for you if you insert values, so you rarely need to do this. You can create a `Vec<T>` that has initial values, and Rust will infer the type. The macro `vec!` will create a new vector. Listing 8-2 creates a new `Vec<i32>` that holds three values.

```
let v = vec![1, 2, 3];
```

Listing 8-2: Creating a new vector containing values

Because we've given initial `i32` values, Rust can infer the type, and the type annotation isn't necessary. Next, we'll look at how to update a vector.

## Updating a Vector

To create a vector and then add elements to it, we'll use the `push` method, shown in Listing 8-3:

```
let mut v = Vec::new();  
  
v.push(5);  
v.push(6);  
v.push(7);  
v.push(8);
```

Listing 8-3: Using the `push` method to add values to a vector

As with any variable, if we want to be able to change the vector, we need to declare it as mutable using the `mut` keyword, as discussed in Chapter 4.

inside are all of type `i32`, and Rust infers this from the `Vec<i32>` annotation.

## Dropping a Vector Drops Its Elements

Like any other `struct`, a vector is freed when it goes out of scope. Listing 8-4:

```
{  
    let v = vec![1, 2, 3, 4];  
  
    // do stuff with v  
  
} // <- v goes out of scope and is freed
```

Listing 8-4: Showing where the vector and its elements are dropped

When the vector gets dropped, all of its contents are freed. The integers it holds will be cleaned up. This may seem simple, but it can get a bit more complicated when you start to use references to elements of the vector. Let's tackle that next!

## Reading Elements of Vectors

Now that you know how to create, update, and drop a vector, reading their contents is a good next step. There are two ways to read a vector. In the examples, we've annotated the type of the returned value from these functions for extra clarity.

Listing 8-5 shows the method of accessing a value by index.

```
let v = vec![1, 2, 3, 4, 5];  
  
let third: &i32 = &v[2];
```

Listing 8-5: Using indexing syntax to access an element in a vector

Listing 8-6 shows the method of accessing a value by iterator.

```

let v = vec![1, 2, 3, 4, 5];
let v_index = 2;

match v.get(v_index) {
    Some(_) => { println!("Reachable element"); }
    None => { println!("Unreachable element"); }
}

```

Listing 8-6: Using the `get` method to access an element in a vector

Note two details here. First, we use the index value `2` to access the third element in the vector. Vectors are indexed by number, starting at zero. Second, we use the `get` method to access the element, which gives us an `Option` type. The `get` method with the index passed as an argument, `v.get(2)`, returns an `Option` type.

Rust has two ways to reference an element so you can handle errors. The `get` method behaves when you try to use an index value that is out of range. For example, let's see what a program will do if it tries to access an element at index 100.

```

let v = vec![1, 2, 3, 4, 5];

let does_not_exist = &v[100];
let does_not_exist = v.get(100);

```

Listing 8-7: Attempting to access the element at index 100 in a vector of 5 elements

When we run this code, the first `&v[100]` method will panic because it references a nonexistent element. This method is useful in a program to crash if there's an attempt to access an element out of range of a vector.

When the `get` method is passed an index that is out of range, it returns `None` without panicking. You would use this method if you want to handle an index out of range of the vector happens occasionally under normal circumstances. Then you can have logic to handle having either `Some(&element)` or `None`. In Chapter 6, we saw how to handle `Option` types. For example, the index could be computed from user input. If they accidentally enter a number that's too large, you could tell the user how many items are in the vector and give them another chance to enter a valid value. That would prevent the program from crashing due to a typo!

When the program has a valid reference, the `get` method returns an `Option` type.

and borrowing rules (covered in Chapter 4) to ensure that references to the contents of the vector remain valid. We can't have mutable and immutable references to the same memory in Listing 8-8, where we hold an immutable reference to the first element and then try to add an element to the end, which won't work.

```
let mut v = vec![1, 2, 3, 4, 5];

let first = &v[0];

v.push(6);
```

Listing 8-8: Attempting to add an element to a vector while holding an immutable reference to an item

Compiling this code will result in this error:

```
error[E0502]: cannot borrow `v` as mutable because it is also borrowed as immutable
-->
   |
4  |     let first = &v[0];
   |                                - immutable borrow occurs here
5  |
6  |     v.push(6);
   |     ^ mutable borrow occurs here
7  |
8  | }
   | - immutable borrow ends here
```

The code in Listing 8-8 might look like it should work, but it doesn't. The first element cares about what changes at the end of the vector, and the way vectors work: adding a new element onto the end requires allocating new memory and copying the old elements into it. If there's not enough room to put all the elements next to each other, the vector has to move them. In that case, the reference to the first element would be invalid. The borrowing rules prevent programs from doing this.

---

Note: For more on the implementation details of Rust's memory management, see "Rustonomicon" at <https://doc.rust-lang.org/stable/rustonomicon/>

---

## Iterating over the Values in a Vector

If we want to access each element in a vector in a safe way, we can use the `iter` method.

elements rather than use indexes to access one. We can use a `for` loop to get immutable references to each element in the vector, and then print them:

```
let v = vec![100, 32, 57];
for i in &v {
    println!("{}", i);
}
```

Listing 8-9: Printing each element in a vector by `for` loop

We can also iterate over mutable references to each element in the vector in order to make changes to all the elements. The code below iterates over each element:

```
let mut v = vec![100, 32, 57];
for i in &mut v {
    *i += 50;
}
```

Listing 8-10: Iterating over mutable references to each element in a vector

To change the value that the mutable reference points to, we use the dereference operator ( `*` ) to get to the value in memory. We'll talk more about `*` in Chapter 15.

## Using an Enum to Store Multiple Types

At the beginning of this chapter, we said that vectors store elements of the same type. This can be inconvenient; there are times when we want to store a list of items of different types. Fortunately, Rust has a way to store multiple types under the same enum type, so when we want to store multiple types in a vector, we can define and use an enum.

For example, say we want to get values from a row in a table. Some of the columns in the row contain integers, some contain floats, and some contain strings. We can define an enum whose variants represent these different types. Then all the enum variants will be considered the same type, so we can create a vector that holds that enum and iterate over it. We've demonstrated this in Listing 8-11:

```
enum SpreadsheetCell {
    Int(i32),
    Float(f64),
    Text(String),
}

let row = vec![
    SpreadsheetCell::Int(3),
    SpreadsheetCell::Text(String::from("b1")),
    SpreadsheetCell::Float(10.12),
];
```

Listing 8-11: Defining an `enum` to store values of

Rust needs to know what types will be in the vector, so it can allocate exactly how much memory on the heap will be required. One secondary advantage is that we can be explicit about the types in a vector. If Rust allowed a vector to hold any type, more of the types would cause errors with the `match` expression at the end of the vector. Using an enum plus a `match` expression at compile time that every possible case is handled.

When you're writing a program, if you don't know what types your program will get at runtime to store in a vector, you can use a trait object, which we'll cover in Chapter 17.

Now that we've discussed some of the most common collection types, let's review the API documentation for all the many useful types in the standard library. For example, in addition to `Vec::last`, `Vec::last_mut` returns the last element. Let's move on to the next chapter.

## Storing UTF-8 Encoded Text with `String`

We talked about strings in Chapter 4, but we'll look at them more in depth here. Rustaceans commonly get stuck on strings due to Rust's propensity for exposing possible errors, so the standard library has a more complex structure than many programmers give them credit for. The `String` type combines in a way that can seem difficult when you're learning new programming languages.

It's useful to discuss strings in the context of collections because `String` is implemented as a collection of bytes, plus some functionality when those bytes are interpreted as

the operations on `String` that every collection has for writing and reading. We'll also discuss the ways in which collections, namely how indexing into a `String`, differ between how people and computers interpret `s`.

## What Is a String?

We'll first define what we mean by the term *string* in the core language, which is the string slice `str` that we saw in `&str`. In Chapter 4, we talked about *string slices*, which are encoded string data stored elsewhere. String literals are a binary output of the program and are therefore

The `String` type, which is provided by Rust's standard library, is a growable, mutable, owned string type. In the core language, Rustaceans refer to “strings” in Rust, they usually use `str` slice `&str` types, not just one of those types. Although `String` and `str`, both types are used heavily in Rust's standard library. String slices are UTF-8 encoded.

Rust's standard library also includes a number of other string types: `OsString`, `OsStr`, `CString`, and `CStr`. These are for storing string data. See how those names all have owned and borrowed variants, just like the `String` and `str` types previously. These string types can store text in different ways in memory in a different way, for example. We will discuss them in this chapter; see their API documentation for more details on when each is appropriate.

## Creating a New String

Many of the same operations available with `Vec` are available with `String` as well, starting with the `new` function to create a new `String`.

```
let mut s = String::new();
```

Listing 8-11: Creating a new, empty `String`

This line creates a new empty string called `s`, which we can then use. Often, we'll have some initial data that we want to store in the string.

use the `to_string` method, which is available on `Display` trait, as string literals do. Listing 8-12 s

```
let data = "initial contents";

let s = data.to_string();

// the method also works on a literal directly
let s = "initial contents".to_string();
```

Listing 8-12: Using the `to_string` method to create a string

This code creates a string containing `initial contents`.

We can also use the function `String::from` to create a string. The code in Listing 8-13 is equivalent to the code in Listing 8-12, but using `String::from` instead of `to_string`:

```
let s = String::from("initial contents");
```

Listing 8-13: Using the `String::from` function to create a string

Because strings are used for so many things, we have many options for strings, providing us with a lot of options. So all of these things all have their place! In this case, `String::from` and `to_string` are both valid things, so which you choose is a matter of style.

Remember that strings are UTF-8 encoded, so we can store any data in them, as shown in Listing 8-14:

```
let hello = String::from("السلام عليكم");
let hello = String::from("Dobry den");
let hello = String::from("Hello");
let hello = String::from("你好");
let hello = String::from("नमस्ते");
let hello = String::from("こんにちは");
let hello = String::from("안녕하세요");
let hello = String::from("你好");
let hello = String::from("Olá");
let hello = String::from("Здравствуй");
let hello = String::from("Hola");
```

Listing 8-14: Storing greetings in different languages

All of these are valid `String` values.



## Updating a String

A `String` can grow in size and its contents can be updated using the `Vec<T>`, if you push more data into it. In addition, you can use the `+` operator or the `format!` macro to concatenate

### Appending to a String with `push_str` and `push`

We can grow a `String` by using the `push_str` method shown in Listing 8-15:

```
let mut s = String::from("foo");
s.push_str("bar");
```

Listing 8-15: Appending a string slice to a `String`

After these two lines, `s` will contain `foobar`. The `push_str` method takes a string slice as an argument because we don't necessarily want to take ownership of the slice. For example, the code in Listing 8-16 shows that it was possible to use `s2` after appending its contents to `s`.

```
let mut s1 = String::from("foo");
let s2 = "bar";
s1.push_str(s2);
println!("s2 is {}", s2);
```

Listing 8-16: Using a string slice after appending

If the `push_str` method took ownership of `s2`, `s2` would no longer be valid on the last line. However, this code works as we intended.

The `push` method takes a single character as an argument. Listing 8-17 shows code that adds the letter `l` to the end of `s`.

```
let mut s = String::from("lo");
s.push('l');
```

Listing 8-17: Adding one character to a `String`

As a result of this code, `s` will contain `lol`.

### Concatenation with the `+` Operator or the `format!` Macro

Often, you'll want to combine two existing strings as shown in Listing 8-18:

```
let s1 = String::from("Hello, ");
let s2 = String::from("world!");
let s3 = s1 + &s2; // Note s1 has been moved
```

Listing 8-18: Using the `+` operator to combine two values

The string `s3` will contain `Hello, world!` as a result. `s1` is no longer valid after the addition and the reason for this has to do with the signature of the method that gets called. The `+` operator uses the `add` method, whose signature is as follows:

```
fn add(self, s: &str) -> String {
```

This isn't the exact signature that's in the standard library, but `add` is defined using generics. Here, we're looking at concrete types substituted for the generic ones, so we can call this method with `String` values. We'll discuss generics in more depth in Chapter 15. Because `add` does not take `&String`, not `&str`, as specified in the second parameter, Listing 8-18 compiles?

First, `s2` has an `&`, meaning that we're adding a reference to a string because of the `s` parameter in the signature. We can't add two `String` values to a `String`; we can only add a `&String` to a `String`, as specified in the second parameter. Listing 8-18 compiles?

The reason we're able to use `&s2` in the call to `add` is because of `coercion`, which here turns `&s2` into `&s2[..]`. We'll discuss `coercion` in more depth in Chapter 15. Because `add` does not take `&String`, not `&str`, as specified in the second parameter, Listing 8-18 compiles?

Second, we can see in the signature that `add` takes `self` as a parameter. This means `s1` in Listing 8-18 is no longer valid after that. So although `add` will copy both strings and create a new one, this new string is a copy of `s1`, appends a copy of the contents of `s2`, and returns the result. In other words, it looks like it's making a new string. This implementation is more efficient than copying.

If we need to concatenate multiple strings, the b  
unwieldy:

```
let s1 = String::from("tic");  
let s2 = String::from("tac");  
let s3 = String::from("toe");  
  
let s = s1 + "-" + &s2 + "-" + &s3;
```

At this point, `s` will be `tic-tac-toe`. With all of  
to see what's going on. For more complicated st  
`format!` macro:

```
let s1 = String::from("tic");  
let s2 = String::from("tac");  
let s3 = String::from("toe");  
  
let s = format!("{}", s1, s2, s3);
```

This code also sets `s` to `tic-tac-toe`. The `for`  
`println!`, but instead of printing the output to  
the contents. The version of the code using `for`  
doesn't take ownership of any of its parameters.

## Indexing into Strings

In many other programming languages, accessir  
referencing them by index is a valid and commo  
access parts of a `String` using indexing syntax  
the invalid code in Listing 8-19:

```
let s1 = String::from("hello");  
let h = s1[0];
```

Listing 8-19: Attempting to use indexing syntax v

This code will result in the following error:

```

error[E0277]: the trait bound `std::string
std::ops::Index<{integer}>` is not satisfi
-->
|
3 |         let h = s1[0];
|                   ^^^^^ the type `std::strin
`{integer}`
|
= help: the trait `std::ops::Index<{inte
`std::string::String`

```

The error and the note tell the story: Rust strings are byte slices. To answer that question, we need to discuss how Rust strings are represented internally.

## Internal Representation

A `String` is a wrapper over a `Vec<u8>`. Let's look at some UTF-8 example strings from Listing 8-14. First, the `len` of the string "Hola" is 4:

```
let len = String::from("Hola").len();
```

In this case, `len` will be 4, which means the vector is 4 bytes long. Each of these letters takes 1 byte when encoded in UTF-8. The following line? (Note that this line begins with the Arabic number 3.)

```
let len = String::from("Здравствуйте").len();
```

Asked how long the string is, you might say 12. But that's not the number of bytes it takes to encode "Здравствуйте" in UTF-8. The Unicode scalar value in that string takes 2 bytes when encoded in UTF-8, so the string's bytes will not always correlate to a visual character. To demonstrate, consider this invalid Rust code:

```
let hello = "Здравствуйте";
let answer = &hello[0];
```

What should the value of `answer` be? Should it be the first character in UTF-8, the first byte of `З` is `208` and the second byte is `208`, but `208` is not a valid character on its own. If a user would want if they asked for the first letter of the string, only data that Rust has at byte index 0. Users get the first byte returned, even if the string contains only Latin letters.

that returned the byte value, it would return `10`, an unexpected value and causing bugs that might not even compile this code at all and prevents miscommunication in the development process.

## Bytes and Scalar Values and Grapheme Clusters

Another point about UTF-8 is that there are actually two ways to look at strings from Rust's perspective: as bytes, scalar values, or grapheme clusters (the closest thing to what we would call *letters*).

If we look at the Hindi word "नमस्ते" written in the vector of `u8` values that looks like this:

```
[224, 164, 168, 224, 164, 174, 224, 164, 164, 164, 224, 165, 135]
```

That's 18 bytes and is how computers ultimately store Unicode scalar values, which are what Rust's `char` type represents.

```
['न', 'म', 'स्', 'ते', 'त', 'े']
```

There are six `char` values here, but the fourth and fifth are diacritics that don't make sense on their own. If we group them into grapheme clusters, we'd get what a person would call the word:

```
["न", "म", "स्", "ते"]
```

Rust provides different ways of interpreting the data so that each program can choose the interpretation of the language the data is in.

A final reason Rust doesn't allow us to index into strings is that indexing operations are expected to always take the minimum time possible to guarantee that performance with a large string. If we walk through the contents from the beginning to the end, we can't skip over invalid characters there were.

## Slicing Strings

Indexing into a string is often a bad idea because it can be slow and error-prone.

the string-indexing operation should be: a byte cluster, or a string slice. Therefore, Rust asks you to use indices to create string slices. To be more that you want a string slice, rather than indexing can use `[]` with a range to create a string slice (

```
let hello = "Здравствуйते";  
  
let s = &hello[0..4];
```

Here, `s` will be a `&str` that contains the first 4 bytes mentioned that each of these characters was 2 bytes.

What would happen if we used `&hello[0..1]`? At runtime in the same way as if an invalid index was used.

```
thread 'main' panicked at 'byte index 1 is out of bounds for `Здравствуйते`',  
src/libcore/str.rs:105:5
```

You should use ranges to create string slices within your program.

## Methods for Iterating Over Strings

Fortunately, you can access elements in a string slice.

If you need to perform operations on individual characters, the way to do so is to use the `chars` method. Calling `chars` returns an iterator of type `char`, and you can iterate over each element:

```
for c in "नमस्ते".chars() {  
    println!("{}", c);  
}
```

This code will print the following:

न  
म  
स  
्  
त  
े

The `bytes` method returns each raw byte, which domain:

```
for b in "नमस्ते".bytes() {  
    println!("{}", b);  
}
```

This code will print the 18 bytes that make up th

```
224  
164  
// --snip--  
165  
135
```

But be sure to remember that valid Unicode sca  
than 1 byte.

Getting grapheme clusters from strings is compl  
provided by the standard library. Crates are ava  
functionality you need.

## Strings Are Not So Simple

To summarize, strings are complicated. Differen  
different choices about how to present this com  
chosen to make the correct handling of `String`  
programs, which means programmers have to p  
data upfront. This trade-off exposes more of the  
apparent in other programming languages, but  
errors involving non-ASCII characters later in you

Let's switch to something a bit less complex: has

## Storing Keys with Associated Values

The last of our common collections is the *hash map*, which is a mapping of keys of type `K` to values of type `V`. The `HashMap` class, which determines how it places these keys and values, is available in the standard library. Many programming languages support this kind of data structure, but with different names, such as `hash`, `map`, `object`, `hash table`, and so on, just to name a few.

Hash maps are useful when you want to look up values by key. You can do this with vectors, but by using a key that can be compared, you could keep track of each team's score in a hash map. The keys are the team names and the values are each team's score. Given a team name, you can look up its score.

We'll go over the basic API of hash maps in this section. The `HashMap` class is hiding in the functions defined on `HashMap<K, V>`. You can check the standard library documentation for more details.

### Creating a New Hash Map

You can create an empty hash map with `new` and `HashMap::new()`. In Listing 8-20, we're keeping track of the scores of the Blue and Yellow teams. The Blue team starts with 10 points, and the Yellow team starts with 50 points.

```
use std::collections::HashMap;

let mut scores = HashMap::new();

scores.insert(String::from("Blue"), 10);
scores.insert(String::from("Yellow"), 50);
```

Listing 8-20: Creating a new hash map and inserting values

Note that we need to first `use` the `HashMap` from the standard library. Of our three common collections, `HashMap` is the only one that's not included in the features brought into scope by `std::collections`. Hash maps also have less support from the standard library than vectors and slices. For example, you can't construct them, for example, with `vec![]`.

Just like vectors, hash maps store their data on the heap. The keys of type `String` and values of type `i32`. Like vectors, the keys of a hash map must have the same type, and all of the values must have the same type.



Another way of constructing a hash map is by using tuples, where each tuple consists of a key and a value. If we gather data into a number of collection types, it's easy to have the team names and initial scores in two separate vectors. We can use the `zip` method to create a vector of tuples where "Blue" is the key and 10 is the value. We could use the `collect` method to turn that vector into a hash map, as shown in Listing 8-21:

```
use std::collections::HashMap;

let teams = vec![String::from("Blue"), String::from("Red")];
let initial_scores = vec![10, 50];

let scores: HashMap<_, _> =
    teams.iter().zip(initial_scores.iter()).collect();
```

Listing 8-21: Creating a hash map from a list of tuples

The type annotation `HashMap<_, _>` is needed here because the `HashMap` type can be used in many different data structures and Rust doesn't know what to infer. For the parameters for the key and value, we use underscores, and Rust can infer the types that the key and value are based on the types of the data in the vectors.

## Hash Maps and Ownership

For types that implement the `Copy` trait, like `i32`, the hash map is the owner of those values. For owned values like `String`, the values in the hash map are not owned by the hash map; the `String` values are owned by the owner of those values, as demonstrated in Listing 8-22:

```
use std::collections::HashMap;

let field_name = String::from("Favorite color");
let field_value = String::from("Blue");

let mut map = HashMap::new();
map.insert(field_name, field_value);
// field_name and field_value are invalid
// and
// see what compiler error you get!
```

Listing 8-22: Showing that keys and values are owned by the owner of the hash map

We aren't able to use the variables `field_name` moved into the hash map with the call to `insert`.

If we insert references to values into the hash map, the references point to the values in the hash map. The values that the references point to live as long as the hash map is valid. We'll talk more about "References with Lifetimes" section in Chapter 10.

## Accessing Values in a Hash Map

We can get a value out of the hash map by providing a key, as shown in Listing 8-23:

```
use std::collections::HashMap;

let mut scores = HashMap::new();

scores.insert(String::from("Blue"), 10);
scores.insert(String::from("Yellow"), 50);

let team_name = String::from("Blue");
let score = scores.get(&team_name);
```

Listing 8-23: Accessing the score for the Blue team

Here, `score` will have the value that's associated with the key. The result will be `Some(&10)`. The result is wrapped in `Some` and `Option<&V>`; if there's no value for that key in the hash map, the result will be `None`. The program will need to handle the `Option` in Chapter 6.

We can iterate over each key/value pair in a hash map using `iter` with vectors, using a `for` loop:

```
use std::collections::HashMap;

let mut scores = HashMap::new();

scores.insert(String::from("Blue"), 10);
scores.insert(String::from("Yellow"), 50);

for (key, value) in &scores {
    println!("{key}: {value}");
}
```

This code will print each pair in an arbitrary order

```
Yellow: 50  
Blue: 10
```

## Updating a Hash Map

Although the number of keys and values is growing, you still have a value associated with it at a time. When you want to update a value, you have to decide how to handle the case where a key already has a value. You could replace the old value with the new value. You could keep the old value and ignore the new value. You could keep the old value if the key *doesn't* already have a value. Or you could keep the new value. Let's look at how to do each of these things.

## Overwriting a Value

If we insert a key and a value into a hash map and then insert a different value, the value associated with that key will be replaced. The code in Listing 8-24 calls `insert` twice, the hash map will contain the new pair because we're inserting the value for the Blue team.

```
use std::collections::HashMap;  
  
let mut scores = HashMap::new();  
  
scores.insert(String::from("Blue"), 10);  
scores.insert(String::from("Blue"), 25);  
  
println!("{:?}", scores);
```

Listing 8-24: Replacing a value stored with a part

This code will print `{"Blue": 25}`. The original value of 10 has been replaced.

## Only Inserting a Value If the Key Has No Value

It's common to check whether a particular key has a value for it. Hash maps have a special API for this. You can call `insert` with a key and a value, and you can want to check as a parameter. The return value is called `Entry` that represents a value that might already be in the map. You can then check whether the key for the Yellow team has a value.

doesn't, we want to insert the value 50, and the `entry` API, the code looks like Listing 8-25:

```
use std::collections::HashMap;

let mut scores = HashMap::new();
scores.insert(String::from("Blue"), 10);

scores.entry(String::from("Yellow")).or_insert(50);
scores.entry(String::from("Blue")).or_insert(50);

println!("{:?}", scores);
```

Listing 8-25: Using the `entry` method to only insert a value

The `or_insert` method on `Entry` is defined to insert a value for the corresponding `Entry` key if that key does not exist. It takes a closure parameter as the new value for this key and returns the value. This technique is much cleaner than writing `insert` twice. It also plays more nicely with the borrow checker.

Running the code in Listing 8-25 will print `{"Yellow": 50, "Blue": 10}`. The first call to `entry` will insert the key for the Yellow team because the Yellow team doesn't have a value already. The second call to `entry` will not insert a value for the Blue team because the Blue team already has the value 10.

## Updating a Value Based on the Old Value

Another common use case for hash maps is to update a value based on the old value. For instance, Listing 8-26 counts the number of times each word appears in some text. We use `entry` and increment the value to keep track of how many times we've seen a word. The first time we've seen a word, we'll first insert

```

use std::collections::HashMap;

let text = "hello world wonderful world";

let mut map = HashMap::new();

for word in text.split_whitespace() {
    let count = map.entry(word).or_insert(0);
    *count += 1;
}

println!("{:?}", map);

```

Listing 8-26: Counting occurrences of words using counts

This code will print `{"world": 2, "hello": 1, ...}`. The `or_insert` method actually returns a mutable reference (`&mut T`), and we store that mutable reference in the `count` variable. Because `count` is a mutable reference, we must first dereference `count` using the `*` operator. Since `count` goes out of scope at the end of the `for` loop, this is allowed by the borrowing rules.

## Hashing Functions

By default, `HashMap` uses a cryptographically secure hashing function to provide resistance to Denial of Service (DoS) attacks. A faster hashing algorithm is available, but the trade-off for better performance is worth it. If you profile your code and find that the hashing function is too slow for your purposes, you can specify a different *hasher*. A hasher is a type that implements the `Hasher` trait. We'll talk about traits and how to implement them in a later chapter. You don't necessarily have to implement your own hasher; many crates are shared by other Rust users that provide hashers using different hashing algorithms.

## Summary

Vectors, strings, and hash maps will provide a lot of utility in programs when you need to store, access, and manipulate data. In the exercises you should now be equipped to solve:

- Given a list of integers, use a vector and return the median (when sorted, the value in the middle); the mode (the value that occurs most often; a hash map will be helpful here); and the most common pair (two values that occur together most often).
- Convert strings to pig latin. The first consonant of each word moves to the end of the word and “ay” is added, so “first” becomes “irfstay”. Words with a vowel have “hay” added to the end instead. Keep in mind the details about UTF-8 encoding.
- Using a hash map and vectors, create a text analyzer that maps employee names to a department in a company. For example, “Engineering” or “Add Amir to Sales.” Then implement methods to find people in a department or all people in the company, sorted alphabetically.

The standard library API documentation describes many examples that hash maps have that will be helpful for these exercises.

We’re getting into more complex programs in which it’s a perfect time to discuss error handling. We’ll do that in the next chapter.

## Error Handling

Rust’s commitment to reliability extends to error handling in software, so Rust has a number of features for handling errors when something goes wrong. In many cases, Rust requires you to acknowledge the possibility of an error and take some action before proceeding. This requirement makes your program more robust and helps you to anticipate and handle them appropriately before you’ve deployed your program.

Rust groups errors into two major categories: *recoverable* and *unrecoverable*. For a recoverable error, such as a file not found or a network problem, you can report the problem to the user and retry the operation. Unrecoverable errors are symptoms of bugs, like trying to access a location that doesn’t exist.

Most languages don’t distinguish between these two types of errors in the same way, using mechanisms such as exceptions. Rust, instead, has the type `Result<T, E>` for recoverable errors, which stops execution when the program encounters an error. This chapter covers calling `panic!` first and then talking about how to recover from an error or to stop execution.

# Unrecoverable Errors with `panic!`

Sometimes, bad things happen in your code, and you need to handle them. In these cases, Rust has the `panic!` macro. When a panic occurs, the program will print a failure message, unwind the stack, and then abort. This most commonly occurs when a bug of some kind is encountered. It is clear to the programmer how to handle the error.

---

## Unwinding the Stack or Aborting in Release Mode

By default, when a panic occurs, the program will unwind the stack. Rust walks back up the stack and cleans up the state of the program. But this walking back and cleaning up can be slow. You can choose to immediately *abort*, which ends the program. The state of the program that the program was using will then need to be cleaned up by the operating system. If in your project you need to make the program as fast as possible, you can switch from unwinding to aborting by setting the `[profile.release]` `panic = 'abort'` to the appropriate `[profile.release]` in your `Cargo.toml`. For example, if you want to abort on panic in release mode:

```
[profile.release]
panic = 'abort'
```

---

Let's try calling `panic!` in a simple program:

Filename: `src/main.rs`

```
fn main() {
    panic!("crash and burn");
}
```

When you run the program, you'll see something like this:

```
$ cargo run
   Compiling panic v0.1.0 (file:///project/target/debug/panic)
   Finished dev [unoptimized + debuginfo] target(s) in 0.1s
   Running `target/debug/panic`
thread 'main' panicked at 'crash and burn'
note: Run with `RUST_BACKTRACE=1` for a backtrace
```

The call to `panic!` causes the error message on the line above. The line shows our panic message and the place in the code where it occurred.

occurred: *src/main.rs:2:4* indicates that it's the second line of the *src/main.rs* file.

In this case, the line indicated is part of our code. In other cases, the `panic!` call came from a library call, and the filename and line number reported are from someone else's code where the `panic!` macro is used. We can use the backtrace to figure out the part of the program that eventually led to the `panic!` call. We'll discuss what a backtrace is in more detail in the next chapter.

## Using a `panic!` Backtrace

Let's look at another example to see what it's like to get a backtrace from a library because of a bug in our code instead of from our code directly. Listing 9-1 has some code that attempts to access an element in a vector:

Filename: *src/main.rs*

```
fn main() {  
    let v = vec![1, 2, 3];  
  
    v[99];  
}
```

Listing 9-1: Attempting to access an element beyond the end of a vector causes a `panic!`

Here, we're attempting to access the hundredth element of a vector (because indexing starts at zero), but it's out of bounds. In this situation, Rust will panic. Using `vec![]` is supposed to create an empty vector, so using `99` is an invalid index, there's no element that Rust can access.

Other languages, like C, will attempt to give you the memory address of the element you're trying to access, even though it isn't what you want: you're accessing memory that would correspond to that element, but that memory doesn't belong to the vector. This is called a buffer overflow, and it's a security vulnerability if an attacker is able to manipulate the program to read data they shouldn't be allowed to that is sensitive.

To protect your program from this sort of vulnerability, Rust will stop execution at an index that doesn't exist, Rust will stop execution and print a backtrace to see:



```
$ cargo run
   Compiling panic v0.1.0 (file:///project
   Finished dev [unoptimized + debuginfo]
   Running `target/debug/panic`
thread 'main' panicked at 'index out of bounds
is
99', /checkout/src/liballoc/vec.rs:1555:16
note: Run with `RUST_BACKTRACE=1` for a backtrace
```

This error points at a file we didn't write, *vec.rs*. That's in the standard library. The code that gets run when you run *vec.rs*, and that is where the `panic!` is actually happening.

The next note line tells us that we can set the `RUST_BACKTRACE` environment variable to get a backtrace of exactly what happened to cause the error. It lists all the functions that have been called to get to the point of the error. They do in other languages: the key to reading the backtrace is to read until you see files you wrote. That's the backtrace. The lines above the lines mentioning your files are the lines that called your code. The lines below are code that called your code. These are the lines of standard library code, or crates that you're using. If you set the `RUST_BACKTRACE` environment variable, you'll see output similar to what you'll see:

```

$ RUST_BACKTRACE=1 cargo run
    Finished dev [unoptimized + debuginfo]
    Running `target/debug/panic`
thread 'main' panicked at 'index out of bounds:
is 99', /checkout/src/liballoc/vec.rs:155:5
stack backtrace:
 0: std::sys::imp::backtrace::tracing::imp::
    at /checkout/src/libstd/sys/unix/backtrace
/gcc_s.rs:49
 1: std::sys_common::backtrace::_print
    at /checkout/src/libstd/sys_common/backtrac
 2: std::panicking::default_hook::{{closure}}
    at /checkout/src/libstd/sys_common/backtrac
    at /checkout/src/libstd/panic.rs:321:5
 3: std::panicking::default_hook
    at /checkout/src/libstd/panic.rs:321:5
 4: std::panicking::rust_panic_with_hook
    at /checkout/src/libstd/panic.rs:398:12
 5: std::panicking::begin_panic
    at /checkout/src/libstd/panic.rs:260:5
 6: std::panicking::begin_panic_fmt
    at /checkout/src/libstd/panic.rs:241:5
 7: rust_begin_unwind
    at /checkout/src/libstd/panic.rs:260:5
 8: core::panicking::panic_fmt
    at /checkout/src/libcore/panic.rs:105:12
 9: core::panicking::panic_bounds_check
    at /checkout/src/libcore/panic.rs:161:12
10: <alloc::vec::Vec<T> as core::ops::Index<isize>::
    at /checkout/src/liballoc/vec.rs:155:5
11: panic::main
    at src/main.rs:4
12: __rust_maybe_catch_panic
    at /checkout/src/libpanic_unwind.rs:99:12
13: std::rt::lang_start
    at /checkout/src/libstd/panic.rs:260:5
    at /checkout/src/libstd/panic.rs:241:5
    at /checkout/src/libstd/rt.rs:51:5
14: main
15: __libc_start_main
16: <unknown>

```

Listing 9-2: The backtrace generated by a call to `panic!` when the environment variable `RUST_BACKTRACE` is set

That's a lot of output! The exact output you see is dependent on your operating system and Rust version. In order to get a backtrace, debug symbols must be enabled. Debug symbols are enabled by default when you run `cargo build` or `cargo run` without the `--release` flag.

In the output in Listing 9-2, line 11 of the backtrace shows the location of the panic: `src/main.rs:4`.

that's causing the problem: line 4 of *src/main.rs*. `panic`, the location pointed to by the first line message, should start investigating. In Listing 9-1, where we `panic` in order to demonstrate how to use backtraces, we request an element at index 99 from a vector that has 10 elements. In the code that panics in the future, you'll need to figure out what values to cause the panic and what the code should do.

We'll come back to `panic!` and when we should use it to handle error conditions in the "To `panic!` or Not" chapter. Next, we'll look at how to recover from errors.

## Recoverable Errors with `Result`

Most errors aren't serious enough to require the program to `panic`. Sometimes, when a function fails, it's for a reason that the program can respond to. For example, if you try to open a file that doesn't exist, you might want to create the file instead.

Recall from "Handling Potential Failure with the `Option` Enum" that the `Result` enum is defined as having two variants,

```
enum Result<T, E> {
    Ok(T),
    Err(E),
}
```

The `T` and `E` are generic type parameters: we'll discuss them in Chapter 10. What you need to know right now is that `Ok` is a value that will be returned in a success case with the type of the value, and `Err` is the type of the error that will be returned in a failure case. Because `Result` has these generic type parameters, the standard library has defined many functions that return `Result` values, where the successful value and error value we want are specified by the generic parameters.

Let's call a function that returns a `Result` value `open_file`. In Listing 9-3 we try to open a file:

Filename: *src/main.rs*

```
use std::fs::File;

fn main() {
    let f = File::open("hello.txt");
}
```

### Listing 9-3: Opening a file

How do we know `File::open` returns a `Result`? We can look at the [library API documentation](#), or we could ask the compiler. The compiler gives us an annotation that we know is *not* the return type of the function. In the code, the compiler will tell us that the types of `File::open` and `u32` are different. Then we can tell us what the type of `f` is. Let's try it! We can change `File::open` isn't of type `u32`, so let's change the

```
let f: u32 = File::open("hello.txt");
```

Attempting to compile now gives us the followin

```
error[E0308]: mismatched types
--> src/main.rs:4:18
|
4 |         let f: u32 = File::open("hello.txt
    |                                ^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^
`std::result::Result`
|
= note: expected type `u32`
       found type `std::result::Resu
std::io::Error>
```

This tells us the return type of the `File::open` function. The generic parameter `T` has been filled in here with `std::fs::File`, which is a file handle. The type `std::io::Error`.

This return type means the call to `File::open` r that we can read from or write to. The function c file might not exist, or we might not have permis `File::open` function needs to have a way to tel and at the same time give us either the file hand information is exactly what the `Result` enum co

In the case where `File::open` succeeds, the variable `ok` contains a file handle. In the case where it fails, there will be an instance of `Err` that contains more information about what happened.

We need to add to the code in Listing 9-3 to take value `File::open` returns. Listing 9-4 shows one basic tool, the `match` expression that we discuss:

Filename: src/main.rs

```
use std::fs::File;

fn main() {
    let f = File::open("hello.txt");

    let f = match f {
        Ok(file) => file,
        Err(error) => {
            panic!("There was a problem opening the file: {}", error);
        },
    };
}
```

Listing 9-4: Using a `match` expression to handle the `File::open` result returned

Note that, like the `Option` enum, the `Result` enum is imported in the prelude, so we don't need to specify `Err` variants in the `match` arms.

Here we tell Rust that when the result is `Ok`, return the `Ok` variant, and we then assign that file handle to `f`. With the `match`, we can use the file handle for reading or writing.

The other arm of the `match` handles the case where `File::open` returns an `Err`. In this example, we've chosen to call the file `hello.txt` in our current directory and we see the output from the `panic!` macro:

```
thread 'main' panicked at 'There was a problem opening the file: No such file or directory (os error 2)', src/main.rs:10:13
note: run with `RUST_BACKTRACE=1` environment variable to display a backtrace
Os { code: 2, message: "No such file or directory" }
```

As usual, this output tells us exactly what has gone wrong.

## Matching on Different Errors

The code in Listing 9-4 will `panic!` no matter what error is returned. One thing we can do instead is take different actions for different errors.

failed because the file doesn't exist, we want to create the new file. If `File::open` failed for any other reason, we didn't have permission to open the file—we still want to create the file the same way as it did in Listing 9-4. Look at Listing 9-5, which shows how to handle this.

Filename: src/main.rs

```
use std::fs::File;
use std::io::ErrorKind;

fn main() {
    let f = File::open("hello.txt");

    let f = match f {
        Ok(file) => file,
        Err(error) => match error.kind() {
            ErrorKind::NotFound => match File::create("hello.txt") {
                Ok(fc) => fc,
                Err(e) => panic!("Tried to create file, but it failed. Problem: {:?}", e),
            },
            other_error => panic!("There was another error: {:?}", other_error),
        },
    };
}
```

Listing 9-5: Handling different kinds of errors in `File::open`

The type of the value that `File::open` returns is `Result<File, io::Error>`, which is a struct provided by the standard library. The `io::Error` struct we can call to get an `io::ErrorKind` value. The `ErrorKind` enum is in the standard library and has variants representing different kinds of errors that might result from an `io` operation. The variant `ErrorKind::NotFound`, which indicates the file was not found, is the one we want. So, we `match` on `f`, but we also then have an inner `match` statement to handle the `NotFound` case.

The condition we want to check in the match guard is `error.kind() == ErrorKind::NotFound`. If it is, we create the file with `File::create`. However, because we need to add another inner `match` statement to handle the case where the file was opened, a different error message will be printed. We want the error message to stay the same so the program panics on any error.

That's a lot of `match`! `match` is very powerful, but in Chapter 13, we'll learn about closures. The `Result` type is a type that accepts a closure, and are implemented as `enum`s.

Rustacean might write this:

```
use std::fs::File;
use std::io::ErrorKind;

fn main() {
    let f = File::open("hello.txt").map_err(|error| {
        if error.kind() == ErrorKind::NotFound {
            File::create("hello.txt").unwrap_or_else(|_| {
                panic!("Tried to create file but it doesn't exist");
            })
        } else {
            panic!("There was a problem opening the file");
        }
    });
}
```

Come back to this example after you've read Chapter 9. The `map_err` and `unwrap_or_else` methods do in this example. There's many more of these methods that can be used when dealing with errors. We'll be looking at some others in the next chapter.

## Shortcuts for Panic on Error: `unwrap` and `expect`

Using `match` works well enough, but it can be a bit verbose to communicate intent well. The `Result<T, E>` type has a lot of methods on it to do various tasks. One of those methods, `unwrap`, that is implemented just like the `match` statement. If the `Result` value is the `Ok` variant, `unwrap` will return the value. If the `Result` is the `Err` variant, `unwrap` will call the `panic!` macro. Here's an example of `unwrap` in action:

Filename: src/main.rs

```
use std::fs::File;

fn main() {
    let f = File::open("hello.txt").unwrap();
}
```

If we run this code without a `hello.txt` file, we'll see a panic. This call that the `unwrap` method makes:

```
thread 'main' panicked at 'called `Result::
Error {
repr: 0s { code: 2, message: "No such file
src/libcore/result.rs:906:4
```

Another method, `expect`, which is similar to `unwrap` but provides an error message. Using `expect` instead of `unwrap` can convey your intent and make tracking down the source of the error. The syntax of `expect` looks like this:

Filename: src/main.rs

```
use std::fs::File;

fn main() {
    let f = File::open("hello.txt").expect(
}
```

We use `expect` in the same way as `unwrap`: to unwrap a `Result` and `panic!` macro. The error message used by `expect` is the parameter that we pass to `expect`, rather than what `unwrap` uses. Here's what it looks like:

```
thread 'main' panicked at 'Failed to open
code:
2, message: "No such file or directory" }
```

Because this error message starts with the text `Failed to open hello.txt`, it will be easier to figure out which `unwrap` is causing the panic. If we use `unwrap` in the same place, the panic print the same message.

## Propagating Errors

When you're writing a function whose implementation might fail, instead of handling the error within this function, you can return the error to the calling code so that it can decide what to do. This gives more control to the calling code, where the logic dictates how the error should be handled in the context of your code.

For example, Listing 9-6 shows a function that returns an error if a file doesn't exist or can't be read; this function will return



called this function:

Filename: src/main.rs

```
use std::io;
use std::io::Read;
use std::fs::File;

fn read_username_from_file() -> Result<String> {
    let f = File::open("hello.txt");

    let mut f = match f {
        Ok(file) => file,
        Err(e) => return Err(e),
    };

    let mut s = String::new();

    match f.read_to_string(&mut s) {
        Ok(_) => Ok(s),
        Err(e) => Err(e),
    }
}
```

Listing 9-6: A function that returns errors to the caller

This function can be written in a much shorter way. We'll explore a lot of it manually in order to explore error handling in a different way. Let's look at the return type of the function. The return type is `Result<String>`. This means the function is returning a value of type `Result`. The generic parameter `T` has been filled in with the concrete type `String`. The generic type `E` has been filled in with the concrete type `io::Error`. If the operation succeeds without any problems, the code that calls the function will receive a value that holds a `String`—the username that the function was looking for. If the function encounters any problems, the code that calls the function will receive an `Err` value that holds an instance of `io::Error` that describes what the problems were. We chose `io::Error` because that happens to be the type of the error returned by the operations we're calling in this function's body: `File::open` and the `read_to_string` method.

The body of the function starts by calling the `File::open` method. The `Result` value returned with a `match` statement. Instead of calling `panic!` in the `Err` case, we rewrap the error value from `File::open` back to the `Result` value. If `File::open` succeeds, we store the file

continue.

Then we create a new `String` in variable `s` and the file handle in `f` to read the contents of the file. The `read_to_string` method also returns a `Result` because it might fail. If it succeeds, then our function has succeeded, and that's now in `s` wrapped in an `Ok`. If `read_to_string` fails, we return the error value of `File::open` in the same way that we returned the error value of `File::open`. However, we don't return `Ok(s)` because this is the last expression in the function.

The code that calls this code will then handle getting a `String` or an `Err` value that contains an `io::Error`. The calling code will do with those values. If the calling code wants to panic and crash the program, use a default `panic!` from somewhere other than a file, for example. This is what the calling code is actually trying to do, so it's better to pass information upward for it to handle appropriately.

This pattern of propagating errors is so common that Rust has a question mark operator `?` to make this easier.

## A Shortcut for Propagating Errors: the `?` Operator

Listing 9-7 shows an implementation of `read_username_from_file` functionality as it had in Listing 9-6, but this implementation uses the `?` operator:

Filename: src/main.rs

```
use std::io;
use std::io::Read;
use std::fs::File;

fn read_username_from_file() -> Result<String> {
    let mut f = File::open("hello.txt")?;
    let mut s = String::new();
    f.read_to_string(&mut s)?;
    Ok(s)
}
```

Listing 9-7: A function that returns errors to the caller

The `?` placed after a `Result` value is defined to

`match` expressions we defined to handle the `Result` of the `Result` is an `Ok`, the value inside the `Ok` expression, and the program will continue. If the `Err` is returned from the whole function as if we had a `value` gets propagated to the calling code.

There is a difference between what the `match` error values taken by `?` go through the `from` function in the standard library, which is used to convert error types. When `?` calls the `from` function, the error type is converted to the error type defined in the return type of the current function. The `from` function returns one error type to represent all the ways a function might fail for many different reasons. As long as the `from` function is implemented to define how to convert itself to the error type, the conversion happens automatically.

In the context of Listing 9-7, the `?` at the end of the `read_to_string` call converts the value inside an `Ok` to the variable `f`. If an error is returned from the whole function and give any `Err` value to the caller, the `?` at the end of the `read_to_string` call.

The `?` operator eliminates a lot of boilerplate and makes the implementation simpler. We could even shorten the `read_to_string` calls immediately after the `?`, as shown in Listing 9-8.

Filename: src/main.rs

```
use std::io;
use std::io::Read;
use std::fs::File;

fn read_username_from_file() -> Result<String> {
    let mut s = String::new();

    File::open("hello.txt")?.read_to_string(&s)?;

    Ok(s)
}
```

Listing 9-8: Chaining method calls after `?`

We've moved the creation of the new `String` in `read_username_from_file` to the beginning of the function, so that part hasn't changed. Instead of creating a variable `s` and then calling `read_to_string` directly onto the result of `File::open`, we now call `read_to_string` directly onto the result of `File::open` and then call `Ok(s)` at the end of the `read_to_string` call, and

the username in `s` when both `File::open` and returning errors. The functionality is again the same; this is just a different, more ergonomic way to write it.

Speaking of different ways to write this function, here's a shorter one:

Filename: `src/main.rs`

```
use std::io;
use std::io::Read;
use std::fs;

fn read_username_from_file() -> Result<String> {
    fs::read_to_string("hello.txt")
}
```

Listing 9-9: Using `fs::read_to_string`

Reading a file into a string is a fairly common operation. Rust provides a convenience function called `fs::read_to_string`, which reads the contents of the file, and puts them into a `String`, then returns it. Of course, this doesn't give us the error handling, so we did it the hard way at first.

## The `?` Operator Can Only Be Used in Functions

The `?` operator can only be used in functions because it is defined to work in the same way as `match`. The part of the `match` that requires a `return Err(e)`, so the return type of the function must be compatible with this `return`.

Let's look at what happens if we use `?` in the `main` function's return type of `()`:

```
use std::fs::File;

fn main() {
    let f = File::open("hello.txt")?;
}
```

When we compile this code, we get the following error:

```

error[E0277]: the trait bound `(): std::ops::Try` is not satisfied
--> src/main.rs:4:13
   |
4  |         let f = File::open("hello.txt"?);
   |                                ^^^^^^^^^
   |                                the `?` operator can only
   |                                be used with types that
   |                                implement the std::ops::Try trait
   |
   = help: the trait std::ops::Try is not implemented for ()
   = note: required by std::ops::Try::from_result

```

This error points out that we're only allowed to use `?` with `Result`. In functions that don't return `Result`, if you return `Result`, you'll need to use a `match` or `unwrap` or the `Result` instead of using `?` to potentially propagate the error.

Now that we've discussed the details of calling `panic!`, let's return to the topic of how to decide which is appropriate for a given situation.

## To `panic!` or Not to `panic!`

So how do you decide when you should call `panic!` or return a `Result`? When code panics, there's no way to recover from the error situation, whether there's a possible way to make the decision on behalf of the code calling the function or not. When you choose to return a `Result`, you're giving the caller options rather than making the decision for it. The caller can then attempt to recover in a way that's appropriate for the situation. Returning an `Err` value in this case is unrecoverable, so it's a recoverable error into an unrecoverable one. The default choice when you're defining a function that can fail is to return a `Result`.

In rare situations, it's more appropriate to write a function that panics. Let's explore why it's appropriate to panic in some cases and tests. Then we'll discuss situations in which it's not possible, but you as a human can. The chapter provides guidelines on how to decide whether to panic in a given situation.

## Examples, Prototype Code, and Tests

When you're writing an example to illustrate some error handling code in the example as well can make it clear that it's understood that a call to a method like `unwrap` is a placeholder for the way you'd want your application to handle errors based on what the rest of your code is doing.

Similarly, the `unwrap` and `expect` methods are useful when you're ready to decide how to handle errors. They are for when you're ready to make your program more robust.

If a method call fails in a test, you'd want the whole test to fail, which isn't the functionality under test. Because `panic` is used for testing, calling `unwrap` or `expect` is exactly what should happen.

## Cases in Which You Have More Information

It would also be appropriate to call `unwrap` when you're sure the `Result` will have an `Ok` value, but it's not clear that you understand. You'll still have a `Result` value that represents the operation you're calling still has the possibility of failing, even if it's logically impossible in your particular situation. If you're inspecting the code that you'll never have an `Error`, you can call `unwrap`. Here's an example:

```
use std::net::IpAddr;

let home: IpAddr = "127.0.0.1".parse().unwrap();
```

We're creating an `IpAddr` instance by parsing a string. Since `127.0.0.1` is a valid IP address, so it's acceptable to use `unwrap`. Having a hardcoded, valid string doesn't change the fact that we still get a `Result` value, and the compiler will still allow the `Error` variant is a possibility because the compiler doesn't know this string is always a valid IP address. If the IP address was taken from a user rather than being hardcoded into the program and the user input was not validated, we'd definitely want to handle the `Result` in a more robust way.

## Guidelines for Error Handling

It's advisable to have your code panic when it's in a bad state. In this context, a *bad state* is when the program is in a state where it's not clear what the next step should be.

contract, or invariant has been broken, such as values, or missing values are passed to your code

- The bad state is not something that's expected
- Your code after this point needs to rely on
- There's not a good way to encode this information

If someone calls your code and passes in values that don't meet the contract, one possible choice might be to call `panic!` and alert the programmer. This is useful for debugging their code so they can fix it during development. It's not always appropriate if you're calling external code that is not under your control. It's inappropriate if the invalid state that you have no way of fixing.

However, when failure is expected, it is more appropriate to return an error. You can make a `panic!` call. Examples include a parser returning an error, or an HTTP request returning a status that indicates you should handle the error. Returning a `Result` indicates that failure is an expected outcome. Your code must decide how to handle.

When your code performs operations on values that are valid first and panic if the values aren't valid. This is a good practice. Attempting to operate on invalid data can expose the main reason the standard library will call `panic!` is memory access: trying to access memory that doesn't exist. This is a common security problem. Functions that panic if the inputs don't meet the contract behavior is only guaranteed if the inputs meet the contract. When the contract is violated makes sense because it indicates a caller-side bug and it's not a kind of error that you have to explicitly handle. In fact, there's no reason for the calling *programmers* need to fix the code. Code that panics on a violation will cause a panic, should be explained in the function.

However, having lots of error checks in all of your code is annoying. Fortunately, you can use Rust's type system (the compiler does) to do many of the checks for you. For example, if you have a parameter, you can proceed with your code assuming you already ensured you have a valid value. For example, `Option`, your program expects to have *something* or *nothing*. It doesn't have to handle two cases for the `Some` and `None` cases, but one case for definitely having a value. Code trying to access `None` won't even compile, so your function doesn't have to handle it. Another example is using an unsigned integer type where the parameter is never negative.

Let's take the idea of using Rust's type system to step further and look at creating a custom type for the game in Chapter 2 in which our code asked the user to guess a number between 1 and 100. We never validated that the user's guess was a valid number before checking it against our secret number; we just assumed it was positive. In this case, the consequences were not too severe: a guess of 0 or "Too low" would still be correct. But it would lead the user toward valid guesses and have different behavior than a guess that's out of range versus when a user types, for example, "abc".

One way to do this would be to parse the guess into a `u32`, allow potentially negative numbers, and then add a check for the range, like so:

```
loop {
    // --snip--

    let guess: i32 = match guess.trim().parse() {
        Ok(num) => num,
        Err(_) => continue,
    };

    if guess < 1 || guess > 100 {
        println!("The secret number will be between 1 and 100.");
        continue;
    }

    match guess.cmp(&secret_number) {
        // --snip--
    }
}
```

The `if` expression checks whether our value is outside the range, and calls `continue` to start the next iteration with another guess. After the `if` expression, we can compare `guess` to the secret number knowing that it's a valid number.

However, this is not an ideal solution: if it was actually used in a program that operated on values between 1 and 100, and it had a performance requirement, having a check like this in every full iteration would have a significant impact on performance).

Instead, we can make a new type and put the validation logic into a function that creates an instance of the type rather than repeating the validation logic in every function that takes a guess. Listing 9-9 shows one way to create an instance of `Guess` if the `new` function is called with a valid string.



```

pub struct Guess {
    value: i32,
}

impl Guess {
    pub fn new(value: i32) -> Guess {
        if value < 1 || value > 100 {
            panic!("Guess value must be between 1 and 100: {value}");
        }

        Guess {
            value
        }
    }

    pub fn value(&self) -> i32 {
        self.value
    }
}

```

Listing 9-10: A `Guess` type that will only continue if the value is between 1 and 100.

First, we define a struct named `Guess` that has a `value` field of type `i32`. This is where the number will be stored.

Then we implement an associated function named `new` that creates new instances of `Guess` values. The `new` function is named `value` of type `i32` and to return a `Guess` value. The `new` function tests `value` to make sure it's between 1 and 100. If the test fails, we make a `panic!` call, which will alert the calling code that they have a bug they need to fix. If the `value` outside this range would violate the constraints of the program. The conditions in which `Guess::new` might panic are documented in the facing API documentation; we'll cover documenting the possibility of a `panic!` in the API documentation in the next chapter. If `value` does pass the test, we create a new `Guess` value with the `value` parameter and return the `Guess`.

Next, we implement a method named `value` that returns the `value` field of the `Guess` instance. This kind of method is called a *getter*, because its purpose is to get some data from the instance. This method is necessary because the `value` field of `Guess` is private, so code outside the `Guess` struct is not allowed to set `value` directly: code outside the `Guess` struct can only use the `new` function to create an instance of `Guess`, thereby

to have a `value` that hasn't been checked by the function.

A function that has a parameter or returns only then declare in its signature that it takes or returns. It wouldn't need to do any additional checks in its

## Summary

Rust's error handling features are designed to help. The `panic!` macro signals that your program is in a state where the process should stop instead of trying to proceed. The `Result` enum uses Rust's type system to indicate that your code could recover from an error. You can use `Result` code that it needs to handle potential success or failure. Using `Result` in the appropriate situations will make your code avoid inevitable problems.

Now that you've seen useful ways that the standard library uses `Option` and `Result` enums, we'll talk about how to use them in your code.

## Generic Types, Traits, and Modules

Every programming language has tools for expressing common concepts. In Rust, one such tool is *generics*. Generics let you write code for concrete types or other properties. When we're talking about the behavior of generics or how they relate to other concepts, we'll be in their place when compiling and running the code.

Similar to the way a function takes parameters and returns values, code on multiple concrete values, functions can take generic types instead of a concrete type, like `i32` or `String`. We'll see generics in Chapter 6 with `Option<T>`, Chapter 8 with `Vec<T>`, and Chapter 9 with `Result<T, E>`. In this chapter, we'll define our own types, functions, and methods with generic parameters.

First, we'll review how to extract a function to reuse it. Then we'll use the same technique to make a generic function that works for many types of their parameters. We'll also explain how to use traits and enum definitions.

Then you'll learn how to use *traits* to define behavior. You'll also learn how to combine traits with generic types to constrain a type to have a particular behavior, as opposed to just any type.

Finally, we'll discuss *lifetimes*, a variety of generic types that deal with how references relate to each other. Lifetimes help you write safe code in many situations while still enabling the compiler to optimize your code.

## Removing Duplication by Extracting Functions

Before diving into generics syntax, let's first look at how to remove duplication. We won't involve generic types by extracting a function. Instead, we'll extract a generic function! In the same way that you can extract a loop into a function, you'll start to recognize duplication and extract it into a function.

Consider a short program that finds the largest number in a list. Listing 10-1.

Filename: src/main.rs

```
fn main() {  
    let number_list = vec![34, 50, 25, 100];  
  
    let mut largest = number_list[0];  
  
    for number in number_list {  
        if number > largest {  
            largest = number;  
        }  
    }  
  
    println!("The largest number is {}", largest);  
}
```

Listing 10-1: Code to find the largest number in a list

This code stores a list of integers in the variable `number_list`. It iterates over each number in the list in a variable named `number`. It compares each `number` to the current largest number in the list, and if the current number is larger than the current largest number, it replaces the number in that variable with the current number. After the loop finishes, `largest` should hold the largest number, which is 100.

To find the largest number in two different lists, we can reuse the code in Listing 10-1 and use the same logic at two different places, as shown in Listing 10-2.

Filename: src/main.rs

```
fn main() {  
    let number_list = vec![34, 50, 25, 100];  
  
    let mut largest = number_list[0];  
  
    for number in number_list {  
        if number > largest {  
            largest = number;  
        }  
    }  
  
    println!("The largest number is {}", largest);  
  
    let number_list = vec![102, 34, 6000, 4567];  
  
    let mut largest = number_list[0];  
  
    for number in number_list {  
        if number > largest {  
            largest = number;  
        }  
    }  
  
    println!("The largest number is {}", largest);  
}
```

Listing 10-2: Code to find the largest number in two different lists

Although this code works, duplicating code is tedious and error-prone. We have to update the code in multiple places when we want to change it.

To eliminate this duplication, we can create an abstraction that operates on any list of integers given to it in a parameter. This makes the code clearer and lets us express the concept of finding the largest number more abstractly.

In Listing 10-3, we extracted the code that finds the largest number into a function named `largest`. Unlike the code in Listing 10-1, which only finds the largest number in one particular list, this program can find the largest number in any list.

Filename: src/main.rs

```

fn largest(list: &[i32]) -> i32 {
    let mut largest = list[0];

    for &item in list.iter() {
        if item > largest {
            largest = item;
        }
    }

    largest
}

fn main() {
    let number_list = vec![34, 50, 25, 100];

    let result = largest(&number_list);
    println!("The largest number is {}", result);

    let number_list = vec![102, 34, 6000, 1024];

    let result = largest(&number_list);
    println!("The largest number is {}", result);
}

```

Listing 10-3: Abstracted code to find the largest item in a slice

The `largest` function has a parameter called `list` that is a slice of `i32` values that we might pass into the function. When we call the function, the code runs on the specific values that we pass in.

In sum, here are the steps we took to change the code in Listing 10-3:

1. Identify duplicate code.
2. Extract the duplicate code into the body of a function and return values of that code in the function.
3. Update the two instances of duplicated code to call the function.

Next, we'll use these same steps with generics to make the code more flexible. In the same way that the function body can operate on specific values, generics allow code to operate on a range of values.

For example, say we had two functions: one that finds the largest item in a slice of `i32` values and one that finds the largest item in a slice of `f32` values. How can we eliminate that duplication? Let's find out!

# Generic Data Types

We can use generics to create definitions for iter which we can then use with many different conc to define functions, structs, enums, and method how generics affect code performance.

## In Function Definitions

When defining a function that uses generics, we the function where we would usually specify the return value. Doing so makes our code more fle to callers of our function while preventing code

Continuing with our `largest` function, Listing 10 find the largest value in a slice.

Filename: src/main.rs

```

fn largest_i32(list: &[i32]) -> i32 {
    let mut largest = list[0];

    for &item in list.iter() {
        if item > largest {
            largest = item;
        }
    }

    largest
}

fn largest_char(list: &[char]) -> char {
    let mut largest = list[0];

    for &item in list.iter() {
        if item > largest {
            largest = item;
        }
    }

    largest
}

fn main() {
    let number_list = vec![34, 50, 25, 100];

    let result = largest_i32(&number_list);
    println!("The largest number is {}", result);

    let char_list = vec!['y', 'm', 'a', 'c'];

    let result = largest_char(&char_list);
    println!("The largest char is {}", result);
}

```

Listing 10-4: Two functions that differ only in the signatures

The `largest_i32` function is the one we extract `largest i32` in a slice. The `largest_char` function. The function bodies have the same code, so let's introduce a generic type parameter in a single

To parameterize the types in the new function with a type parameter, just as we do for the value parameter. We'll use the identifier `T` as a type parameter name. But we'll use longer parameter names in Rust are short, often just a single letter. A convention is CamelCase. Short for "type," `T` is the convention used by programmers.

When we use a parameter in the body of the function, we use the parameter name in the signature so the compiler can find it. Similarly, when we use a type parameter name in the signature, we declare the type parameter name before we use it. In this case, we place type name declarations inside the function signature, before the parameter list, like this:

```
fn largest<T>(list: &[T]) -> T {
```

We read this definition as: the function `largest` has one parameter named `list`, which will return a value of the same type as the elements in `list`.

Listing 10-5 shows the combined `largest` function definition and its signature. The listing also shows how to use the function on a slice of `i32` values or `char` values. Note that this code doesn't compile yet; we'll see why later in this chapter.

Filename: src/main.rs

```
fn largest<T>(list: &[T]) -> T {
    let mut largest = list[0];

    for &item in list.iter() {
        if item > largest {
            largest = item;
        }
    }

    largest
}

fn main() {
    let number_list = vec![34, 50, 25, 100];

    let result = largest(&number_list);
    println!("The largest number is {}", result);

    let char_list = vec!['y', 'm', 'a', 'c'];

    let result = largest(&char_list);
    println!("The largest char is {}", result);
}
```

Listing 10-5: A definition of the `largest` function that doesn't compile yet

If we compile this code right now, we'll get this error:



```

error[E0369]: binary operation `>` cannot
--> src/main.rs:5:12
   |
5  |         if item > largest {
   |                ^^^^^^^^^^^^^^^^^
   |
   = note: an implementation of `std::cmp::`
`T`

```

The note mentions `std::cmp::PartialOrd`, which we will explore in the next section. For now, this error states that `T` does not implement `PartialOrd` for all possible types that `T` could be. Because we want to use `T` in the body, we can only use types whose values can be compared. Since comparisons, the standard library has the `std::cmp::PartialOrd` trait implemented on types (see Appendix C for more on traits). To make `T` implement `PartialOrd`, that a generic type has a particular trait in the “Where” clause. We will explore other ways of using generic type parameters in the next section.

## In Struct Definitions

We can also define structs to use a generic type parameter using the `<T>` syntax. Listing 10-6 shows how to define a struct that holds coordinate values of any type.

Filename: src/main.rs

```

struct Point<T> {
    x: T,
    y: T,
}

fn main() {
    let integer = Point { x: 5, y: 10 };
    let float = Point { x: 1.0, y: 4.0 };
}

```

Listing 10-6: A `Point<T>` struct that holds `x` and `y` coordinates of any type

The syntax for using generics in struct definition is similar to function definitions. First, we declare the name of the type parameter just after the name of the struct. Then we can use the type parameter in the struct definition where we would otherwise specify concrete types.

Note that because we’ve used only one generic type parameter, the definition says that the `Point<T>` struct is generic for all types that implement `PartialOrd`.

`x` and `y` are *both* that same type, whatever that type is. If `x` is an instance of a `Point<T>` that has values of different types, it won't compile.

Filename: src/main.rs

```
struct Point<T> {
    x: T,
    y: T,
}

fn main() {
    let wont_work = Point { x: 5, y: 4.0 }
}
```

Listing 10-7: The fields `x` and `y` must be the same generic data type `T`.

In this example, when we assign the integer value 5 to `x`, that the generic type `T` will be an integer for this line. But when we specify 4.0 for `y`, which we've defined to have the same type as `x`, we get a mismatch error like this:

```
error[E0308]: mismatched types
--> src/main.rs:7:38
   |
 7 |         let wont_work = Point { x: 5, y: 4.0 };
   |         ^
   |         found
   |         floating-point variable
   |         = note: expected type `{integer}`
   |                found type `{float}`
```

To define a `Point` struct where `x` and `y` are both of the same type, we can use multiple generic type parameters. We can change the definition of `Point` to be generic over two types: `T` and `U`.

Filename: src/main.rs

```

struct Point<T, U> {
    x: T,
    y: U,
}

fn main() {
    let both_integer = Point { x: 5, y: 10 }
    let both_float = Point { x: 1.0, y: 4.0 }
    let integer_and_float = Point { x: 5, y: 10.0 }
}

```

Listing 10-8: A `Point<T, U>` generic over two types, allowing for different types

Now all the instances of `Point` shown are allowed to use any type for the parameters in a definition as you want, but using generics can be hard to read. When you need lots of generic types, your code needs restructuring into smaller pieces.

## In Enum Definitions

As we did with structs, we can define enums to have multiple variants. Let's take another look at the `Option<T>` enum, which we used in Chapter 6:

```

enum Option<T> {
    Some(T),
    None,
}

```

This definition should now make more sense to you. It's an enum that is generic over type `T` and has two variants: a `Some` variant of type `T`, and a `None` variant that doesn't hold a value. Because this enum, we can express the abstract concept of having an optional value. Because `Option<T>` is generic, we can use this abstraction for any optional value.

Enums can use multiple generic types as well. The `Vec<T, A>` we used in Chapter 9 is one example:

```
enum Result<T, E> {
    Ok(T),
    Err(E),
}
```

The `Result` enum is generic over two types, `T` which holds a value of type `T`, and `Err`, which I definition makes it convenient to use the `Result` operation that might succeed (return a value of of some type `E`). In fact, this is what we used to was filled in with the type `std::fs::File` when `E` was filled in with the type `std::io::Error` w file.

When you recognize situations in your code with that differ only in the types of the values they hc generic types instead.

## In Method Definitions

We can implement methods on structs and enum generic types in their definitions, too. Listing 10- defined in Listing 10-6 with a method named `x`

Filename: src/main.rs

```
struct Point<T> {
    x: T,
    y: T,
}

impl<T> Point<T> {
    fn x(&self) -> &T {
        &self.x
    }
}

fn main() {
    let p = Point { x: 5, y: 10 };

    println!("p.x = {}", p.x());
}
```

Listing 10-9: Implementing a method named `x` return a reference to the `x` field of type `T`

Here, we've defined a method named `x` on `Poi` data in the field `x`.

Note that we have to declare `T` just after `impl` implementing methods on the type `Point<T>`. If `impl`, Rust can identify that the type in the angle brackets is a generic type rather than a concrete type.

We could, for example, implement methods only on `Point<T>` instances with any generic type `T` rather than on `Point<T>` instances with any concrete type `f32`, meaning we don't declare an `impl` block for `f32`.

```
impl Point<f32> {  
    fn distance_from_origin(&self) -> f32 {  
        (self.x.powi(2) + self.y.powi(2)).sqrt()  
    }  
}
```

Listing 10-10: An `impl` block that only applies to the type `Point<f32>` for the generic type parameter `T`

This code means the type `Point<f32>` will have the `distance_from_origin` method and other instances of `Point` will not have this method defined. The method returns the distance from the point at coordinates (0.0, 0.0) and uses math only for floating point types.

Generic type parameters in a struct definition are also used in that struct's method signatures. For example, the `mixup` method on the `Point<T, U>` struct from Listing 10-9 takes `Point` as a parameter, which might have different types. When calling `mixup` on a `Point`, the method creates a new `Point` from the `self` `Point` (of type `T`) and the `y` value from the parameter `Point` (of type `U`).

Filename: src/main.rs

```

struct Point<T, U> {
    x: T,
    y: U,
}

impl<T, U> Point<T, U> {
    fn mixup<V, W>(self, other: Point<V, W>)
        Point {
            x: self.x,
            y: other.y,
        }
}

fn main() {
    let p1 = Point { x: 5, y: 10.4 };
    let p2 = Point { x: "Hello", y: 'c'};

    let p3 = p1.mixup(p2);

    println!("p3.x = {}, p3.y = {}", p3.x, p3.y);
}

```

Listing 10-11: A method that uses different generic parameters

In `main`, we've defined a `Point` that has an `i32` for `x` (with value `5`) and a `f64` for `y` (with value `10.4`). The `p2` variable is a `Point` (with value `"Hello"`) and a `char` for `y` (with value `'c'`). The `mixup` method call with argument `p2` gives us `p3`, which will have an `i32` for `x` and a `char` for `y`. The `p3` variable will have a `char` for `y`, because the `println!` macro call will print `p3.x = 5, p3.y = c`.

The purpose of this example is to demonstrate that generic parameters are declared with `impl` and some are defined in the struct definition. Here, the generic parameters `T` and `U` are declared with `impl` and they go with the struct definition. The generic parameters `V` and `W` are defined in the `fn mixup`, because they're only relevant to the `mixup` method.

## Performance of Code Using Generics

You might be wondering whether there is a runtime performance cost to using generic types. The good news is that Rust implements generics in a way that your code doesn't run any slower using generic types than using concrete types.

Rust accomplishes this by performing monomorphization, which means that the compiler replaces generic types with their concrete types at compile time.

generics at compile time. *Monomorphization* is the process of converting generic code into specific code by filling in the concrete types.

In this process, the compiler does the opposite of what it does for generic code. In the generic function in Listing 10-5: the compiler looks at the generic code and generates code for the concrete types. The generated code is called and generates code for the concrete types.

Let's look at how this works with an example that uses the `Option<T>` enum:

```
let integer = Some(5);
let float = Some(5.0);
```

When Rust compiles this code, it performs monomorphization. The compiler reads the values that have been used and identifies two kinds of `Option<T>`: one is `i32` and the other is `f64`. It then expands the generic definition of `Option<T>` into specific definitions for `i32` and `f64`, thereby replacing the generic definition with the specific definitions.

The monomorphized version of the code looks like this. The `Option<T>` is replaced with the specific definitions for `i32` and `f64`.

Filename: src/main.rs

```
enum Option_i32 {
    Some(i32),
    None,
}

enum Option_f64 {
    Some(f64),
    None,
}

fn main() {
    let integer = Option_i32::Some(5);
    let float = Option_f64::Some(5.0);
}
```

Because Rust compiles generic code into code that is specific to the concrete types, we pay no runtime cost for using generics. When we use generics, we would have had to duplicate each definition by hand. Monomorphization makes Rust's generics extremely efficient.

# Traits: Defining Shared Behavior

A *trait* tells the Rust compiler about functionality with other types. We can use traits to define shared behavior. We can use trait bounds to specify that a generic can have certain behavior.

---

Note: Traits are similar to a feature often called a *trait* in other languages, although with some differences.

---

## Defining a Trait

A type's behavior consists of the methods we can call on it. If we share the same behavior, we can define a trait. Traits are a way to group method signatures and behaviors necessary to accomplish some purpose.

For example, let's say we have multiple structs that hold text: a `NewsArticle` struct that holds a news article, a `Tweet` that can have at most 280 characters, and a `Reply` whether it was a new tweet, a retweet, or a reply to another tweet.

We want to make a media aggregator library that can handle these types. The methods for each type might be stored in a `NewsArticle` or `Tweet` instance. We need to request that summary method on an instance. Listing 10-12 shows the trait that expresses this behavior.

Filename: `src/lib.rs`

```
pub trait Summary {  
    fn summarize(&self) -> String;  
}
```

Listing 10-12: A `Summary` trait that consists of the `summarize` method

Here, we declare a trait using the `trait` keyword. The trait is named `Summary` in this case. Inside the curly brackets, we describe the behaviors of the types that implement the trait:

```
fn summarize(&self) -> String.
```



After the method signature, instead of providing brackets, we use a semicolon. Each type implements custom behavior for the body of the method. The type that has the `Summary` trait will have the method signature exactly.

A trait can have multiple methods in its body: the trait body is one line per method and each line ends in a semicolon.

## Implementing a Trait on a Type

Now that we've defined the desired behavior using the `Summary` trait, we implement it on the types in our media aggregator. For the implementation of the `Summary` trait on the `NewsArticle` struct, we define `headline`, the author, and the location to create the `Tweet` struct, we define `summarize` as the implementation of the trait, assuming that tweet content is already available.

Filename: `src/lib.rs`

```
pub struct NewsArticle {
    pub headline: String,
    pub location: String,
    pub author: String,
    pub content: String,
}

impl Summary for NewsArticle {
    fn summarize(&self) -> String {
        format!("{}", by {} ({}), self.headline, self.location)
    }
}

pub struct Tweet {
    pub username: String,
    pub content: String,
    pub reply: bool,
    pub retweet: bool,
}

impl Summary for Tweet {
    fn summarize(&self) -> String {
        format!("{}", self.username, self.content)
    }
}
```

### Listing 10-13: Implementing the `Summary` trait on

Implementing a trait on a type is similar to implementing a function. The main difference is that after `impl`, we put the trait name instead of the function name. We use the `for` keyword, and then specify the name of the type that implements the trait. Within the `impl` block, we put the methods that the trait has defined. Instead of adding a semicolon at the end of the curly brackets and fill in the method body with the code for the methods of the trait to have for the particular type.

After implementing the trait, we can call the methods of the trait. In this case, we can call the `summary` method and `Tweet` in the same way we call regular methods.

```
let tweet = Tweet {  
    username: String::from("horse_ebooks"),  
    content: String::from("of course, as you know, there are no  
people"),  
    reply: false,  
    retweet: false,  
};  
  
println!("1 new tweet: {}", tweet.summary());
```

This code prints

```
1 new tweet: horse_ebooks: of course, as you know, there are no people
```

Note that because we defined the `Summary` trait on `Tweet` and `Vec<T>` in the same `lib.rs` in Listing 10-13, they're in the same `lib.rs`. If `lib.rs` is for a crate we've called `aggregator` and we want to use the `Summary` trait in the `aggregator` crate's functionality to implement the `Summary` trait for another type in the `aggregator` crate's scope. They would need to import the trait from the `aggregator` crate by specifying `use aggregator::Summary;` and then implement `Summary` for their type. The `Summary` trait is defined in the `aggregator` crate, so we can implement it for another type in the `aggregator` crate, which is what we did in Listing 10-12.

One restriction to note with trait implementation is that you can only implement a trait on a type if either the trait or the type is local to the crate. For example, we can't implement standard library traits like `Display` on our `aggregator` crate functionality, because the `Display` trait is defined in the `std` crate. We can also implement `Summary` on `Vec<T>` in the `aggregator` crate, because the trait `Summary` is local to our `aggregator` crate.

But we can't implement external traits on external types. For example, we can't implement the `Display` trait on `Vec<T>` within the `aggregator` crate.

`Display` and `Vec<T>` are defined in the standard `aggregator` crate. This restriction is part of a principle and more specifically the *orphan rule*, so named for its present. This rule ensures that other people's crates can't implement the trait `Display` for `Vec<T>` and vice versa. Without the rule, two crates could implement `Display` for `Vec<T>` and Rust wouldn't know which implementation to use.

## Default Implementations

Sometimes it's useful to have default behavior for a trait instead of requiring implementations for all types. To implement the trait on a particular type, we can use `impl` to provide default behavior.

Listing 10-14 shows how to specify a default string `Summary` trait instead of only defining the `summarize` method. See Listing 10-12.

Filename: `src/lib.rs`

```
pub trait Summary {  
    fn summarize(&self) -> String {  
        String::from("(Read more...)")  
    }  
}
```

Listing 10-14: Definition of a `Summary` trait with a `summarize` method

To use a default implementation to summarize a `NewsArticle` instead of defining a custom implementation, we specify an `impl` for `Summary` for `NewsArticle` like this:

Even though we're no longer defining the `summarize` method directly, we've provided a default implementation that implements the `Summary` trait. As a result, we can create an instance of `NewsArticle`, like this:

```

let article = NewsArticle {
    headline: String::from("Penguins win t
    location: String::from("Pittsburgh, PA
    author: String::from("Iceburgh"),
    content: String::from("The Pittsburgh
    hockey team in the NHL."),
};

println!("New article available! {}", arti

```

This code prints `New article available!` (Read

Creating a default implementation for `summariz`  
 anything about the implementation of `Summary`  
 reason is that the syntax for overriding a default  
 syntax for implementing a trait method that doe

Default implementations can call other methods  
 methods don't have a default implementation. It  
 useful functionality and only require implement  
 example, we could define the `Summary` trait to h  
 whose implementation is required, and then def  
 default implementation that calls the `summarize`

```

pub trait Summary {
    fn summarize_author(&self) -> String;

    fn summarize(&self) -> String {
        format!("(Read more from {}...)",
    }
}

```

To use this version of `Summary`, we only need to  
 implement the trait on a type:

```

impl Summary for Tweet {
    fn summarize_author(&self) -> String {
        format!("@{}", self.username)
    }
}

```

After we define `summarize_author`, we can call  
 struct, and the default implementation of `summa`  
`summarize_author` that we've provided. Because  
`summarize_author`, the `Summary` trait has given  
 method without requiring us to write any more

```
let tweet = Tweet {
    username: String::from("horse_ebooks")
    content: String::from("of course, as y
people"),
    reply: false,
    retweet: false,
};

println!("1 new tweet: {}", tweet.summariz
```

This code prints `1 new tweet: (Read more from`

Note that it isn't possible to call the default implementation of that same method.

## Traits as arguments

Now that you know how to define traits and implement them, let's explore how to use traits to accept arguments as parameters.

For example, in Listing 10-13, we implemented the `NewsArticle` and `Tweet` types. We can define a function that takes an argument of type `Summary` and calls its `summarize` method on its parameter `item`, which is of some type that implements the `Summary` trait. To do this, we can use the `impl Trait` syntax.

```
pub fn notify(item: impl Summary) {
    println!("Breaking news! {}", item.summarize());
}
```

In the body of `notify`, we can call any methods defined by the `Summary` trait, like `summarize`.

## Trait Bounds

The `impl Trait` syntax works for short examples, but it's not ideal for larger programs. This is called a 'trait bound', and it looks like this:

```
pub fn notify<T: Summary>(item: T) {
    println!("Breaking news! {}", item.summarize());
}
```

This is equivalent to the example above, but it adds a trait bound with the declaration of the generic type parameter `T` in angle brackets. Because of the trait bound on `T`, we can call `summarize` on `item`.

instance of `NewsArticle` or `Tweet`. Code that calls `notify` like a `String` or an `i32`, won't compile, because `Summary`.

When should you use this form over `impl Trait`? In the shorter examples, trait bounds are nice for more context. If you wanted to take two things that implement `Summary`:

```
pub fn notify(item1: impl Summary, item2: impl Summary) { ... }  
pub fn notify<T: Summary>(item1: T, item2: T) { ... }
```

The version with the bound is a bit easier. In general, it makes your code the most understandable.

### Multiple trait bounds with `+`

We can specify multiple trait bounds on a generic function. For example, to use display formatting on the type `T`. If we have a `summarize` method, we can use `T: Summary + Display`. That means `T` that implements `Summary` and `Display`. This can be written as:

### `where` clauses for clearer code

However, there are downsides to using too many trait bounds. Functions with multiple generic parameters can get cluttered with trait bound information between a function's name and its signature, making it hard to read. For this reason, specifying trait bounds inside a `where` clause after the function signature is often preferred. Instead of writing this:

```
fn some_function<T: Display + Clone, U: Clone>(t: T, u: U) { ... }
```

we can use a `where` clause, like this:

```
fn some_function<T, U>(t: T, u: U) -> i32  
    where T: Display + Clone,  
           U: Clone + Debug  
{ ... }
```

This function's signature is less cluttered in that the parameters and return type are close together, similar to a function without generics.

## Returning Traits

We can use the `impl Trait` syntax in return position that implements a trait:

```
fn returns_summarizable() -> impl Summary
    Tweet {
        username: String::from("horse_ebook"),
        content: String::from("of course,
people"),
        reply: false,
        retweet: false,
    }
}
```

This signature says, "I'm going to return something that implements the `Summary` trait, but I'm not going to tell you the exact type. The caller doesn't know that.

Why is this useful? In chapter 13, we're going to rely heavily on traits: closures, and iterators. These functions are long, and the compiler knows, or types that are very, very long. You can write "this returns an `Iterator`" without needing to write out the full type.

This only works if you have a single type that you can return. If you have multiple types, this would *not* work:

```
fn returns_summarizable(switch: bool) -> impl Summary {
    if switch {
        NewsArticle {
            headline: String::from("Penguin
Championship!"),
            location: String::from("Pittsburgh"),
            author: String::from("Iceburgh"),
            content: String::from("The Pittsburgh
the best
hockey team in the NHL."),
        }
    } else {
        Tweet {
            username: String::from("horse_ebook"),
            content: String::from("of course,
know, people"),
            reply: false,
            retweet: false,
        }
    }
}
```

Here, we try to return either a `NewsArticle` or a `Video`. This introduces some restrictions around how `impl Trait` works. To learn more, see the Rust documentation until Chapter 17, "trait objects".

## Fixing the `largest` Function with Trait

Now that you know how to specify the behavior of a type parameter's bounds, let's return to Listing 1. The `largest` function that uses a generic type parameter `T`. In the original code, we received this error:

```
error[E0369]: binary operation `>` cannot be applied to type `T`
--> src/main.rs:5:12
   |
5  |         if item > largest {
   |               ^^^^^^^^^^^^^
   |               = note: an implementation of `std::cmp::PartialOrd` is not implemented for `T`
```

In the body of `largest` we wanted to compare `item` with `largest` using the greater than (`>`) operator. Because that operator is defined in the standard library trait `std::cmp::PartialOrd`, we need to specify the trait bounds for `T` so the `largest` function can compare. We don't need to bring `PartialOrd` into scope with a `use` prelude. Change the signature of `largest` to look like this:

```
fn largest<T: PartialOrd>(list: &[T]) -> T {
```

This time when we compile the code, we get a different error:



```

error[E0508]: cannot move out of type `[T]
--> src/main.rs:2:23
  |
2 |         let mut largest = list[0];
  |                                ^^^^^^^^
  |                                |
  |                                cannot move out
  |                                help: consider using `list.get(0)`
  |                                or `list[0].clone()`

error[E0507]: cannot move out of borrowed content
--> src/main.rs:4:9
  |
4 |         for &item in list.iter() {
  |             ^----
  |             ||
  |             |hint: to prevent move, use `list.iter().cloned()`
  |             cannot move out of borrowed content

```

The key line in this error is **cannot move out of** our non-generic versions of the **largest** function **largest i32** or **largest char**. As discussed in the “Stack-4, types like **i32** and **char** that have a known size implement the **Copy** trait. But when we made the **list** parameter to have **Copy** trait. Consequently, we wouldn’t be able to move **list** into the **largest** variable, resulting in this error

To call this code with only those types that implement **Copy** to the trait bounds of **T**! Listing 10-15 shows the **largest** function that will compile as long as the types we pass into the function implement the **PartialOrd** trait. **PartialOrd** **char** do.

Filename: src/main.rs

```

fn largest<T: PartialOrd + Copy>(list: &[T])
    let mut largest = list[0];

    for &item in list.iter() {
        if item > largest {
            largest = item;
        }
    }

    largest
}

fn main() {
    let number_list = vec![34, 50, 25, 100];

    let result = largest(&number_list);
    println!("The largest number is {}", result);

    let char_list = vec!['y', 'm', 'a', 'c'];

    let result = largest(&char_list);
    println!("The largest char is {}", result);
}

```

Listing 10-15: A working definition of the `largest` type that implements the `PartialOrd` and `Copy`

If we don't want to restrict the `largest` function to types that implement the `Copy` trait, we could specify that `T` has the trait `Clone`. We could clone each value in the slice when we want to compare it. Using the `clone` function means we have to allocate memory for each clone. This can be slow if we're working with large amounts of data.

Another way we could implement `largest` is for it to return a `T` value in the slice. If we change the return type to `T` and change the body of the function to return a reference to the largest element, we can avoid heap allocations. We could also avoid heap allocations by using the `Copy` trait bounds and we could avoid heap allocations by using alternate solutions on your own!

## Using Trait Bounds to Conditionally Implement Methods

By using a trait bound with an `impl` block that implements methods conditionally for types that implement the trait, we can implement methods conditionally. For example, the type `Pair<T>` in Listing 10-16 always implements the `cmp_display` method only if `T` implements the `cmp_display` method.

the `PartialOrd` trait that enables comparison and printing.

```
use std::fmt::Display;

struct Pair<T> {
    x: T,
    y: T,
}

impl<T> Pair<T> {
    fn new(x: T, y: T) -> Self {
        Self {
            x,
            y,
        }
    }
}

impl<T: Display + PartialOrd> Pair<T> {
    fn cmp_display(&self) {
        if self.x >= self.y {
            println!("The largest member is x")
        } else {
            println!("The largest member is y")
        }
    }
}
```

Listing 10-16: Conditionally implement methods based on trait bounds

We can also conditionally implement a trait for a type based on trait bounds. Implementations of a trait on any type that implement a trait are called *blanket implementations* and are extensively used in the standard library. For example, the standard library implements the `ToString` trait for any type that implements the `Display` trait. The `impl` block in Listing 10-17 shows this code:

```
impl<T: Display> ToString for T {
    // --snip--
}
```

Because the standard library has this blanket implementation, the `to_string` method defined by the `ToString` trait is available for any type that implements the `Display` trait. For example, we can turn integer values like this because integers implement `Display`.

```
let s = 3.to_string();
```

Blanket implementations appear in the document “Implementors” section.

Traits and trait bounds let us write code that use duplication but also specify to the compiler that particular behavior. The compiler can then use that that all the concrete types used with our code provide that behavior. In dynamically typed languages, we would get an error on a type that the type didn’t implement. But Rust so we’re forced to fix the problems before our code we don’t have to write code that checks for behavior already checked at compile time. Doing so improves give up the flexibility of generics.

Another kind of generic that we’ve already been ensuring that a type has the behavior we want, let’s be valid as long as we need them to be. Let’s look at

## Validating References with Lifetimes

One detail we didn’t discuss in the “References and Borrowing” chapter is that every reference in Rust has a *lifetime*, which is valid. Most of the time, lifetimes are implicit and types are inferred. We must annotate types when the similar way, we must annotate lifetimes when they are related in a few different ways. Rust requires us to use generic lifetime parameters to ensure the actual lifetime definitely be valid.

The concept of lifetimes is somewhat different from other languages, arguably making lifetimes Rust’s most unique feature. We won’t cover lifetimes in their entirety in this chapter but you might encounter lifetime syntax so you can become familiar with the “Advanced Lifetimes” section in Chapter 19 for more details.

## Preventing Dangling References with Lifetimes

The main aim of lifetimes is to prevent dangling references, which are references that point to memory that has been freed.

reference data other than the data it's intended  
Listing 10-17, which has an outer scope and an i

```
{
    let r;

    {
        let x = 5;
        r = &x;
    }

    println!("r: {}", r);
}
```

Listing 10-17: An attempt to use a reference wh

---

Note: The examples in Listings 10-17, 10-18, a  
giving them an initial value, so the variable na  
first glance, this might appear to be in conflict  
However, if we try to use a variable before giv  
time error, which shows that Rust indeed doe

---

The outer scope declares a variable named `r` w  
scope declares a variable named `x` with the init  
we attempt to set the value of `r` as a reference  
we attempt to print the value in `r`. This code w  
referring to has gone out of scope before we try

```
error[E0597]: `x` does not live long enough
--> src/main.rs:7:5
   |
6  |         r = &x;
   |               - borrow occurs here
7  |     }
   |     ^ `x` dropped here while still borrowed
...
10 | }
   | - borrowed value needs to live until the end of the function
```

The variable `x` doesn't "live long enough." The `r`  
when the inner scope ends on line 7. But `r` is st  
its scope is larger, we say that it "lives longer." If  
would be referencing memory that was deallocate  
anything we tried to do with `r` wouldn't work co  
that this code is invalid? It uses a borrow checke

## The Borrow Checker

The Rust compiler has a *borrow checker* that confirms all borrows are valid. Listing 10-18 shows the same code as Listing 10-17, but with lifetime annotations showing the lifetimes of the variables `r` and `x`.

```
{
    let r; // -----+---
           //          |
    {
        let x = 5; // -+--- 'b |
        r = &x;    //  |      |
    }             // -+      |
                 //          |
    println!("r: {}", r); //          |
}                     // -----+
```

Listing 10-18: Annotations of the lifetimes of `r` and `x` respectively

Here, we've annotated the lifetime of `r` with `'a` and the lifetime of `x` with `'b`. As you can see, the inner `'b` block is much smaller than the outer `'a` block. At compile time, Rust compares the size of the two lifetimes and rejects the code because `'b` is shorter than `'a`: the variable `x` goes out of scope long as the reference `r` is still in scope.

Listing 10-19 fixes the code so it doesn't have a lifetime error without any errors.

```
{
    let x = 5; // -----+---
              //          |
    let r = &x; // --+--- 'a |
              //  |      |
    println!("r: {}", r); //  |      |
                        // --+      |
}                       // -----+
```

Listing 10-19: A valid reference because the data and the reference have the same lifetime

Here, `x` has the lifetime `'a`, which in this case is the same as the lifetime of the reference `r` because Rust knows that the reference `r` is valid as long as `x` is valid.

Now that you know where the lifetimes of references live, you can use `String::from` to ensure references will always be valid. You can also use `String::from` for parameters and return values in the context of functions.

## Generic Lifetimes in Functions

Let's write a function that returns the longer of two string slices and return a string slice. After we write the function, the code in Listing 10-20 should print

Filename: src/main.rs

```
fn main() {  
    let string1 = String::from("abcd");  
    let string2 = "xyz";  
  
    let result = longest(string1.as_str(), string2);  
    println!("The longest string is {}", result);  
}
```

Listing 10-20: A `main` function that calls the `longest` function with two string slices

Note that we want the function to take string slices, not `String`. We don't want the `longest` function to take ownership of the strings; we allow the function to accept slices of a `String` (like `string1.as_str()`) as well as string literals (which is what `string2` is).

Refer to the “String Slices as Parameters” section in Chapter 10 about why the parameters we use in Listing 10-20 are the best choice.

If we try to implement the `longest` function as shown in Listing 10-21, it won't compile.

Filename: src/main.rs

```
fn longest(x: &str, y: &str) -> &str {  
    if x.len() > y.len() {  
        x  
    } else {  
        y  
    }  
}
```

Listing 10-21: An implementation of the `longest` function that takes two string slices but does not yet compile

Instead, we get the following error that talks about

```
error[E0106]: missing lifetime specifier
--> src/main.rs:1:33
   |
1  | fn longest(x: &str, y: &str) -> &str {
   |                                     ^ expected
   |
   = help: this function's return type contains a lifetime, but no lifetime parameter
   = help: you can use the `'_` lifetime to indicate that the return value borrows for the duration of the function call
```

The help text reveals that the return type needs a lifetime parameter because Rust can't tell whether the reference being returned borrows from `x` or `y`. Actually, we don't know either, because the `if` block returns a reference to `x` and the `else` block returns a reference to `y`.

When we're defining this function, we don't know which parameter will be passed into this function, so we don't know whether the reference will execute. We also don't know the concrete lifetime of the parameter passed in, so we can't look at the scopes as we can't determine whether the reference we return will borrow from the parameter. We can't determine this either, because it doesn't know how the parameter relates to the lifetime of the return value. To fix this, we need to add lifetime parameters that define the relationship between the parameter and the return value so the checker can perform its analysis.

## Lifetime Annotation Syntax

Lifetime annotations don't change how long any function can accept any type when the signature of a function can accept references with any lifetime parameter. Lifetime annotations describe the relationship between references to each other without affecting the lifetime of the references.

Lifetime annotations have a slightly unusual syntax. They must start with an apostrophe ( `'` ) and are usually followed by a lowercase letter. Most people use the name `'a`. We place lifetime annotations after the `&` of a reference, using a space to separate the annotation from the reference's type.

Here are some examples: a reference to an `i32` with lifetime `'a`, a reference to an `i32` that has a lifetime parameter `'a`, and a reference to an `i32` that also has the lifetime `'a`.



```

&i32          // a reference
&'a i32       // a reference with an explicit lifetime
&'a mut i32   // a mutable reference with an explicit lifetime

```

One lifetime annotation by itself doesn't have much meaning. Lifetime annotations are meant to tell Rust how generic lifetimes and references relate to each other. For example, let's say a function has a parameter `first` that is a reference to an `i32`. If the function also has another parameter named `second` that is a reference to an `i32`, it must have the lifetime `'a`. The lifetime annotations in the function signature for `second` must both live as long as that generic lifetime `'a`.

## Lifetime Annotations in Function Signatures

Now let's examine lifetime annotations in the context of function signatures. With generic type parameters, we need to declare the lifetime of the return value. Angle brackets between the function name and the return type allow us to express in this signature that all the references in the return value must have the same lifetime. We'll use the `longest` function to show this, as shown in Listing 10-22.

Filename: `src/main.rs`

```

fn longest<'a>(x: &'a str, y: &'a str) -> &'a str {
    if x.len() > y.len() {
        x
    } else {
        y
    }
}

```

Listing 10-22: The `longest` function definition shows how the return type signature must have the same lifetime `'a` as the parameters.

This code should compile and produce the result shown in Listing 10-20. The `main` function in Listing 10-20.

The function signature now tells Rust that for so long as the references `x` and `y` live, the return value will live at least as long as lifetime `'a`. These constraints enforce that the return value will live as long as the references. Remember, when we specify the lifetime in the function signature, we're not changing the lifetimes of the references themselves.

Rather, we're specifying that the borrow checker adhere to these constraints. Note that the `longest` function doesn't specify exactly how long `x` and `y` will live, only that something that will satisfy this signature.

When annotating lifetimes in functions, the annotations are not in the function body. Rust can analyze the code to help. However, when a function has references to variables that live outside the function, it becomes almost impossible for Rust to figure out the lifetimes of the return values on its own. The lifetimes might be inferred, but they might not be. This is why we need to annotate the lifetimes of the return values.

When we pass concrete references to `longest`, the lifetime `'a` substituted for `'a` is the part of the scope of `x` and `y` that both references live in. In other words, the generic lifetime `'a` will get the smaller of the lifetimes of `x` and `y`. Because we use the same lifetime parameter `'a`, the return value will have the length of the smaller of the lifetimes of `x` and `y`.

Let's look at how the lifetime annotations restrict references that have different concrete lifetimes in the following example.

Filename: src/main.rs

```
fn main() {
    let string1 = String::from("long string is long");

    {
        let string2 = String::from("xyz");
        let result = longest(string1.as_str(), string2.as_str());
        println!("The longest string is {}", result);
    }
}
```

Listing 10-23: Using the `longest` function with references that have different concrete lifetimes

In this example, `string1` is valid until the end of the `main` function, and `string2` is valid until the end of the inner scope, and `result` is valid until the end of the inner scope. Run this code, and you'll see the output of this code; it will compile and print `The longest string is long string is long.`

Next, let's try an example that shows that the lifetime of the return value must be the smaller lifetime of the two arguments.

`result` variable outside the inner scope but `lea` `result` variable inside the scope with `string2` uses `result` outside the inner scope, after the i Listing 10-24 will not compile.

Filename: src/main.rs

```
fn main() {  
    let string1 = String::from("long strir  
    let result;  
    {  
        let string2 = String::from("xyz");  
        result = longest(string1.as_str(),  
    }  
    println!("The longest string is {}", r  
}
```

Listing 10-24: Attempting to use `result` after `s`

When we try to compile this code, we'll get this e

```
error[E0597]: `string2` does not live long  
--> src/main.rs:15:5  
14 |         result = longest(string1.as_s  
15 |     }  
   |     ^ `string2` dropped here while st  
16 |     println!("The longest string is {  
17 | }  
   | - borrowed value needs to live until
```

The error shows that for `result` to be valid for would need to be valid until the end of the outer annotated the lifetimes of the function parameter lifetime parameter `'a`.

As humans, we can look at this code and see that and therefore `result` will contain a reference to gone out of scope yet, a reference to `string1` v statement. However, the compiler can't see that We've told Rust that the lifetime of the reference the same as the smaller of the lifetimes of the re borrow checker disallows the code in Listing 10-reference.

Try designing more experiments that vary the va

passed in to the `longest` function and how the hypotheses about whether or not your experiment will pass before you compile; then check to see if you're right.

## Thinking in Terms of Lifetimes

The way in which you need to specify lifetime parameters for a function is doing. For example, if we changed the `longest` function to always return the first parameter rather than the longest, we wouldn't need to specify a lifetime on the `y` parameter at compile time:

Filename: `src/main.rs`

```
fn longest<'a>(x: &'a str, y: &str) -> &'a str {
    x
}
```

In this example, we've specified a lifetime parameter for the return type, but not for the parameter `y`, because `y` has no relationship with the lifetime of `x` or the return value.

When returning a reference from a function, the lifetime of the reference needs to match the lifetime parameter for the function. If the reference returned does *not* refer to one of the variables created within this function, which would be a dangling reference, it will go out of scope at the end of the function. Consider the implementation of the `longest` function that we saw earlier:

Filename: `src/main.rs`

```
fn longest<'a>(x: &str, y: &str) -> &'a str {
    let result = String::from("really long");
    result.as_str()
}
```

Here, even though we've specified a lifetime parameter for the return type, the implementation will fail to compile because the lifetime of the parameters does not match the lifetime of the return value. Here is the error message:

```

error[E0597]: `result` does not live long
--> src/main.rs:3:5
   |
3 |     result.as_str()
   |     ^^^^^^^ does not live long enough
4 | }
   | - borrowed value only lives until here
   |
note: borrowed value must be valid for the
function body at 1:1...
--> src/main.rs:1:1
   |
1 | / fn longest<'a>(x: &str, y: &str) ->
2 | |     let result = String::from("really long");
3 | |     result.as_str()
4 | | }
   | |_^

```

The problem is that `result` goes out of scope at the end of the `longest` function. We're also trying to return a reference to `result`. There is no way we can specify lifetime, and Rust won't let us create a dangling reference, and Rust won't let us create a dangling reference. The best fix would be to return an owned data type, but then the function is then responsible for cleaning up the memory.

Ultimately, lifetime syntax is about connecting the lifetimes of variables and return values of functions. Once they're connected, we can allow memory-safe operations and disallow operations that would otherwise violate memory safety.

## Lifetime Annotations in Struct Definitions

So far, we've only defined structs to hold owned data, but in that case we would need to add a lifetime parameter to the struct's definition. Listing 10-25 shows a struct `ImportantExcerpt` that holds a string slice.

Filename: src/main.rs

```

struct ImportantExcerpt<'a> {
    part: &'a str,
}

fn main() {
    let novel = String::from("Call me Ishmael");
    let first_sentence = novel.split('.')
        .next()
        .expect("Could not find a '.');
    let i = ImportantExcerpt { part: first
}

```

Listing 10-25: A struct that holds a reference, so annotation

This struct has one field, `part`, that holds a string reference. When using generic data types, we declare the name of the generic type inside the angle brackets after the name of the struct so we can use it in the body of the struct definition. This annotation `ImportantExcerpt` can't outlive the reference it

The `main` function here creates an instance of the struct that holds a reference to the first sentence of the `novel` string. The data in `novel` exists before the `ImportantExcerpt` is created. In addition, `novel` doesn't go out of scope until after the `main` function scope, so the reference in the `ImportantExcerpt` is valid.

## Lifetime Elision

You've learned that every reference has a lifetime. In this chapter, we'll learn about lifetime parameters for functions or structs that don't have lifetime annotations. In Listing 4-9, we had a function in Listing 4-9, which is shown here without lifetime annotations.

Filename: `src/lib.rs`

```
fn first_word(s: &str) -> &str {
    let bytes = s.as_bytes();

    for (i, &item) in bytes.iter().enumerate() {
        if item == b' ' {
            return &s[0..i];
        }
    }

    &s[..]
}
```

Listing 10-26: A function we defined in Listing 4-1 with lifetime annotations, even though the parameter and return type have lifetime annotations.

The reason this function compiles without lifetime annotations in Rust versions (pre-1.0) of Rust, this code wouldn't have compiled. At that time, the function would have been written like this:

```
fn first_word<'a>(s: &'a str) -> &'a str {
```

After writing a lot of Rust code, the Rust team found that the patterns for entering the same lifetime annotations over and over again in many situations were predictable and followed a few common patterns. They programmed these patterns into the compiler's inference engine to infer the lifetimes in these situations and would not require explicit annotations.

This piece of Rust history is relevant because it's the patterns that will emerge and be added to the compiler's inference engine. Annotations might be required.

The patterns programmed into Rust's analysis of lifetimes are called *elision rules*. These aren't rules for programmers to write, but cases that the compiler will consider, and if you don't write the lifetimes explicitly.

The elision rules don't provide full inference. If Rust can't infer a lifetime, but there is still ambiguity as to what lifetimes the compiler will guess what the lifetime of the remaining references is. Instead of guessing, the compiler will give you an error message. The lifetime annotations that specify how the references are used are called *input lifetimes*.

Lifetimes on function or method parameters and return values are called *output lifetimes*.

The compiler uses three rules to figure out what aren't explicit annotations. The first rule applies third rules apply to output lifetimes. If the compiler and there are still references for which it can't figure stop with an error.

The first rule is that each parameter that is a reference parameter. In other words, a function with one reference parameter: `fn foo<'a>(x: &'a i32) ;` a function with separate lifetime parameters: `fn foo<'a, 'b>(`

The second rule is if there is exactly one input lifetime assigned to all output lifetime parameters: `fn f`

The third rule is if there are multiple input lifetimes `&self` or `&mut self` because this is a method, no output lifetime parameters. This third rule makes the compiler write because fewer symbols are necessary.

Let's pretend we're the compiler. We'll apply the rules to the lifetimes of the references in the signature of the function 10-26 are. The signature starts without any lifetime

```
fn first_word(s: &str) -> &str {
```

Then the compiler applies the first rule, which specifies its own lifetime. We'll call it `'a` as usual, so now the signature is

```
fn first_word<'a>(s: &'a str) -> &str {
```

The second rule applies because there is exactly one input lifetime specifies that the lifetime of the one input parameter is the same as the output lifetime, so the signature is now this:

```
fn first_word<'a>(s: &'a str) -> &'a str {
```

Now all the references in this function signature are resolved. The compiler can continue its analysis without needing the programmer to write this function signature.

Let's look at another example, this time using the rules for lifetime parameters when we started working with

```
fn longest(x: &str, y: &str) -> &str {
```



Let's apply the first rule: each parameter gets its own lifetime instead of one, so we have two lifetimes.

```
fn longest<'a, 'b>(x: &'a str, y: &'b str)
```

You can see that the second rule doesn't apply to the return type's lifetime. The third rule doesn't apply either, because it's not a method, so none of the parameters are `&T`. By the time we get to the third rule, we still haven't figured out what the return type is, so we get an error trying to compile the code in Listing 10-10. The lifetime elision rules but still couldn't figure out what the return type is, so the signature.

Because the third rule really only applies in method signatures, let's look at it in that context next to see why the third rule means that lifetimes in method signatures very often.

## Lifetime Annotations in Method Definitions

When we implement methods on a struct with lifetimes, the lifetime of generic type parameters shown in Listing 10-11. The lifetime of parameters depends on whether they're method parameters and return values.

Lifetime names for struct fields always need to be placed before the struct's name, and then used after the struct's name, because they're part of the type.

In method signatures inside the `impl` block, references to the struct's fields, or they might be references to the struct's fields. Lifetime elision rules often make it so that lifetimes in method signatures. Let's look at some examples of the `ImportantExcerpt` that we defined in Listing 10-12.

First, we'll use a method named `level` whose parameter is a reference to the struct and whose return value is an `i32`, which is not a lifetime.

```
impl<'a> ImportantExcerpt<'a> {  
    fn level(&self) -> i32 {  
        3  
    }  
}
```

The lifetime parameter declaration after `impl` is required, but we're not required to annotate the because of the first elision rule.

Here is an example where the third lifetime elision

```
impl<'a> ImportantExcerpt<'a> {  
    fn announce_and_return_part(&self, announcement: &'a str)  
        println!("Attention please: {}", announcement);  
    self.part  
}
```

There are two input lifetimes, so Rust applies the both `&self` and `announcement` their own lifetime parameters is `&self`, the return type gets the lifetime have been accounted for.

## The Static Lifetime

One special lifetime we need to discuss is `'static` duration of the program. All string literals have to annotate as follows:

```
let s: &'static str = "I have a static lifetime";
```

The text of this string is stored directly in the binary available. Therefore, the lifetime of all string literals

You might see suggestions to use the `'static` lifetime specifying `'static` as the lifetime for a reference you have actually lives the entire lifetime of your program whether you want it to live that long, even if it comes results from attempting to create a dangling reference lifetimes. In such cases, the solution is fixing the `'static` lifetime.

## Generic Type Parameters, Trait Lifetimes Together

Let's briefly look at the syntax of specifying generic lifetimes all in one function!

```
use std::fmt::Display;

fn longest_with_an_announcement<'a, T>(x:
    &'a str
    where T: Display
{
    println!("Announcement! {}", ann);
    if x.len() > y.len() {
        x
    } else {
        y
    }
}
```

This is the `longest` function from Listing 10-22 slices. But now it has an extra parameter named `ann` that can be filled in by any type that implements the `Display` trait. This extra parameter will be printed along with the lengths of the string slices, which is why the `Display` trait is required. Because lifetimes are a type of generic, the declaration of the function and the generic type parameter `T` go in the same place as the function name.

## Summary

We covered a lot in this chapter! Now that you know how to use traits and trait bounds, and generic lifetime parameters, you can write code without repetition that works in many different ways. Traits let you apply the code to different types. Traits are useful even though the types are generic, they'll have the behavior you want. You also learned how to use lifetime annotations to ensure that there are no dangling references. And all of this analysis happens at compile time and doesn't affect runtime performance!

Believe it or not, there is much more to learn on generics in this chapter: Chapter 17 discusses trait objects, which are a different kind of trait. Chapter 19 covers more complex scenarios involving some advanced type system features. But next, we'll look at how to use the `unsafe` keyword so you can make sure your code is working the way you want.

# Writing Automated Tests

In his 1972 essay “The Humble Programmer,” Edsger Dijkstra wrote, “The only way to show the presence of a bug is to find it. The only way to show the absence of a bug is to show that the program is correct. This is a very effective way to show the presence of a bug, but it is inadequate for showing their absence.” That does as much as we can!

Correctness in our programs is the extent to which they do what we want them to do. Rust is designed with a high degree of correctness in mind, but correctness is complex and not easily automated. The compiler shoulders a huge part of this burden, but the type system still leaves a lot of room for incorrectness. As such, Rust includes support for writing tests within the language.

As an example, say we write a function called `add` that takes an integer number and returns an integer as a result. When we implement it, Rust does all the type checking and borrow checking to ensure that, for instance, we aren’t passing a `String` to this function. But Rust *can’t* check that this function actually returns the parameter plus 2 rather than the parameter minus 50! That’s where tests come in.

We can write tests that assert, for example, that the `add` function, the returned value is `5`. We can run these tests against our code to make sure any existing correct behavior is preserved.

Testing is a complex skill: although we can’t cover all the good tests in one chapter, we’ll discuss the mechanics of writing tests, talk about the annotations and macros available for testing, and the default behavior and options provided for running tests into unit tests and integration tests.

## How to Write Tests

Tests are Rust functions that verify that the code behaves in the expected manner. The bodies of test functions typically follow a three-step process:

1. Set up any needed data or state.
2. Run the code you want to test.
3. Assert the results are what you expect.

Let's look at the features Rust provides specific actions, which include the `test` attribute, a few attribute.

## The Anatomy of a Test Function

At its simplest, a test in Rust is a function that's annotated with the `test` attribute. Attributes are metadata about pieces of Rust code. The `test` attribute we used with structs in Chapter 5. To create a test function, add `#[test]` on the line before `fn`. When you run `cargo test` command, Rust builds a test runner binary that executes the `test` attribute and reports on whether each test passes or fails.

In Chapter 7, we saw that when we make a new module with a test function in it is automatically added to the `mod` block. You start writing your tests so you don't have to repeat the syntax of test functions every time you start a new module. You can add additional test functions and as many test modules as you want.

We'll explore some aspects of how tests work by creating a new library project. We'll generate code for us without actually testing any code. We'll write tests that call some code that we've written and see how it works.

Let's create a new library project called `adder`:

```
$ cargo new adder --lib
    Created library `adder` project
$ cd adder
```

The contents of the `src/lib.rs` file in your `adder` library project will look like this:

Filename: `src/lib.rs`

```
#[cfg(test)]
mod tests {
    #[test]
    fn it_works() {
        assert_eq!(2 + 2, 4);
    }
}
```

Listing 11-1: The test module and function generated by `cargo new`

For now, let's ignore the top two lines and focus on the `it_works` function. Note the `#[test]` annotation before the `fn` line.

function, so the test runner knows to treat this as a non-test function in the `tests` module to help with common operations, so we need to indicate which with the `#[test]` attribute.

The function body uses the `assert_eq!` macro from `std`. This assertion serves as an example of the format for which this test passes.

The `cargo test` command runs all tests in our project.

```
$ cargo test
   Compiling adder v0.1.0 (file:///project/target/debug/deps/adder-ce99t...)
   Finished dev [unoptimized + debuginfo] target(s) in 0.1s
   Running target/debug/deps/adder-ce99t...

running 1 test
test tests::it_works ... ok

test result: ok. 1 passed; 0 failed; 0 ignored; 0 filtered out

   Doc-tests adder

running 0 tests

test result: ok. 0 passed; 0 failed; 0 ignored; 0 filtered out
```

Listing 11-2: The output from running the `cargo test` command.

Cargo compiled and ran the test. After the `Compiling` line is the line `running 1 test`. The next line shows the test function, called `it_works`, and the result of running the test. A summary of running the tests appears next. The summary shows that all the tests passed, and the portion that reads `0 filtered out` shows the number of tests that passed or failed.

Because we don't have any tests we've marked as `ignored`. We also haven't filtered the tests because the output shows `0 filtered out`. We'll talk about ignoring tests in the next section, "Controlling How Tests Are Run."

The `0 measured` statistic is for benchmark tests. Benchmark tests are, as of this writing, only available in Rust 1.30.0. See [the Rust documentation about benchmark tests](#) to learn more.

The next part of the test output, which starts with `Doc-tests`, shows the results of any documentation tests. We don't have any documentation tests in this example.

compile any code examples that appear in our *A* us keep our docs and our code in sync! We'll discuss tests in the "Documentation Comments" section the `Doc-tests` output.

Let's change the name of our test to see how the the `it_works` function to a different name, such

Filename: `src/lib.rs`

```
#[cfg(test)]
mod tests {
    #[test]
    fn exploration() {
        assert_eq!(2 + 2, 4);
    }
}
```

Then run `cargo test` again. The output now shows `it_works`:

```
running 1 test
test tests::exploration ... ok

test result: ok. 1 passed; 0 failed; 0 ignored
```

Let's add another test, but this time we'll make a something in the test function panics. Each test main thread sees that a test thread has died, the about the simplest way to cause a panic in Chap macro. Enter the new test, `another`, so your `src`

Filename: `src/lib.rs`

```
#[cfg(test)]
mod tests {
    #[test]
    fn exploration() {
        assert_eq!(2 + 2, 4);
    }

    #[test]
    fn another() {
        panic!("Make this test fail");
    }
}
```

Listing 11-3: Adding a second test that will fail be

Run the tests again using `cargo test`. The output shows that our `exploration` test passed and a

```
running 2 tests
test tests::exploration ... ok
test tests::another ... FAILED

failures:

---- tests::another stdout ----
    thread 'tests::another' panicked at 'Make this test fail', src/lib.rs:10:8
note: Run with `RUST_BACKTRACE=1` for a backtrace

failures:
    tests::another

test result: FAILED. 1 passed; 1 failed; 0 ignored
error: test failed
```

Listing 11-4: Test results when one test passes and one fails

Instead of `ok`, the line `test tests::another ... FAILED` appears between the individual results and the summary line. The detailed reason for each test failure. In this case, it says `panicked at 'Make this test fail'`, which has the location of the failure. The next section lists just the names of all the failing tests. If there are lots of tests and lots of detailed failing test output, you can run just that test to more easily debug it. See the section “Controlling How Tests Are Run” for more details.

The summary line displays at the end: overall, one test passed and one test failed.

Now that you’ve seen what the test results look like, let’s look at some macros other than `panic!` that are useful for testing.

## Checking Results with the `assert!` Macro

The `assert!` macro, provided by the standard library, is used to ensure that some condition in a test evaluates to `true`. It takes an argument that evaluates to a Boolean. If the argument is `true`, the test passes. If the value is `false`, the `assert!` macro fails, which causes the test to fail. Using the `assert!` macro is similar to using the `panic!` macro, but it’s more specific.



code is functioning in the way we intend.

In Chapter 5, Listing 5-15, we used a `Rectangle` which are repeated here in Listing 11-5. Let's put write some tests for it using the `assert!` macro

Filename: src/lib.rs

```
#[derive(Debug)]
pub struct Rectangle {
    length: u32,
    width: u32,
}

impl Rectangle {
    pub fn can_hold(&self, other: &Rectangle) -> bool {
        self.length > other.length && self.width > other.width
    }
}
```

Listing 11-5: Using the `Rectangle` struct and its

The `can_hold` method returns a Boolean, which we can use with the `assert!` macro. In Listing 11-6, we write a test that checks whether a `Rectangle` instance that has a length of 10 can hold another `Rectangle` instance with a width of 1:

Filename: src/lib.rs

```
#[cfg(test)]
mod tests {
    use super::*;

    #[test]
    fn larger_can_hold_smaller() {
        let larger = Rectangle { length: 10, width: 5 };
        let smaller = Rectangle { length: 5, width: 1 };

        assert!(larger.can_hold(&smaller));
    }
}
```

Listing 11-6: A test for `can_hold` that checks whether a larger rectangle can hold a smaller rectangle

Note that we've added a new line inside the `tests` module. The `tests` module is a regular module that follows the same rules as any other module.

Chapter 7 in the “Privacy Rules” section. Because module, we need to bring the code under test in the inner module. We use a glob here so anything available to this `tests` module.

We’ve named our test `larger_can_hold_smaller` `Rectangle` instances that we need. Then we call the result of calling `larger.can_hold(&smaller)` return `true`, so our test should pass. Let’s find out.

```
running 1 test
test tests::larger_can_hold_smaller ... ok

test result: ok. 1 passed; 0 failed; 0 ignored
```

It does pass! Let’s add another test, this time asserting that a smaller rectangle can’t hold a larger rectangle:

Filename: `src/lib.rs`

```
#[cfg(test)]
mod tests {
    use super::*;

    #[test]
    fn larger_can_hold_smaller() {
        // --snip--
    }

    #[test]
    fn smaller_cannot_hold_larger() {
        let larger = Rectangle { length: 8, width: 7 };
        let smaller = Rectangle { length: 5, width: 1 };

        assert!(!smaller.can_hold(&larger));
    }
}
```

Because the correct result of the `can_hold` function is `false`, we negate that result before we pass it to the `assert!` macro. `assert!` will only pass if `can_hold` returns `false`:

```
running 2 tests
test tests::smaller_cannot_hold_larger ... ok
test tests::larger_can_hold_smaller ... ok

test result: ok. 2 passed; 0 failed; 0 ignored
```

Two tests that pass! Now let's see what happens introduce a bug in our code. Let's change the `can_hold` method by replacing the greater-than sign with `<` lengths:

```
// --snip--

impl Rectangle {
    pub fn can_hold(&self, other: &Rectangle) -> bool {
        self.length < other.length && self.width < other.width
    }
}
```

Running the tests now produces the following:

```
running 2 tests
test tests::smaller_cannot_hold_larger ... ok
test tests::larger_can_hold_smaller ... FAILED

failures:

---- tests::larger_can_hold_smaller stdout ----
thread 'tests::larger_can_hold_smaller' panicked at 'assertion failed: larger.can_hold(&smaller)', src/lib.rs:10:13
note: Run with `RUST_BACKTRACE=1` for a backtrace.

failures:
    tests::larger_can_hold_smaller

test result: FAILED. 1 passed; 1 failed; 0 ignored
out
```

Our tests caught the bug! Because `larger.length` comparison of the lengths in `can_hold` now returns

## Testing Equality with the `assert_eq!` macro

A common way to test functionality is to compare the value you expect the code to return to make using the `assert!` macro and passing it an expression. However, this is such a common test that the standard library provides two macros—`assert_eq!` and `assert_ne!`—to perform equality and inequality tests. These macros compare two arguments for equality. If the assertion fails, they also print the two values if the assertion fails, which is helpful for debugging. For example, the `assert_eq!` macro only takes the `==` expression, not the values that lead to the

In Listing 11-7, we write a function named `add_two` that returns the result. Then we test this function using

Filename: `src/lib.rs`

```
pub fn add_two(a: i32) -> i32 {
    a + 2
}

#[cfg(test)]
mod tests {
    use super::*;

    #[test]
    fn it_adds_two() {
        assert_eq!(4, add_two(2));
    }
}
```

Listing 11-7: Testing the function `add_two` using

Let's check that it passes!

```
running 1 test
test tests::it_adds_two ... ok

test result: ok. 1 passed; 0 failed; 0 ignored
```

The first argument we gave to the `assert_eq!` macro was the value returned by calling `add_two(2)`. The line for this test is `test` and the `ok` text indicates that our test passed!

Let's introduce a bug into our code to see what happens when `assert_eq!` fails. Change the implementation of `add_two` to add 3 instead of 2:

```
pub fn add_two(a: i32) -> i32 {
    a + 3
}
```

Run the tests again:

```

running 1 test
test tests::it_adds_two ... FAILED

failures:

---- tests::it_adds_two stdout ----
      thread 'tests::it_adds_two' panicked at
assertion failed: `(left == right)`
  left: `4`,
 right: `5`, src/lib.rs:11:8
note: Run with `RUST_BACKTRACE=1` for a backtrace

failures:
    tests::it_adds_two

test result: FAILED. 0 passed; 1 failed; 0 ignored
out

```

Our test caught the bug! The `it_adds_two` test failed with the message `assertion failed: `(left == right)` and sh` was `5`. This message is useful and helps us start debugging. The `assert_eq!` macro was `4` but the `right` value, `add_two(2)`, was `5`.

Note that in some languages and test frameworks, assertions that assert two values are equal are called `expect_eq`, which we specify the arguments matters. However, in Rust, `assert_eq!` and `assert_ne!` are used, and the order in which we specify the values doesn't matter. We could also write `assert_eq!(add_two(2), 4)`, which would result in the same error message: `assertion failed: `(left == right)` and the`

The `assert_ne!` macro will pass if the two values are not equal. This macro is most useful for cases where the value *will* be, but we know what the value definitely *won't* be. For example, if we're testing a function that takes an input and returns a value in some way, but the way in which the input is checked is different from the way in which the output is checked, then the best thing to assert is that the function is not equal to the input.

Under the surface, the `assert_eq!` and `assert_ne!` macros are implemented using `PartialEq` and `Debug` traits, respectively. When the assertions fail, they use debug formatting, which means the values are formatted using the `PartialEq` and `Debug` traits. All the primitive library types implement these traits. For structs, you need to implement `PartialEq` to assert that values are equal.

equal. You'll need to implement `Debug` to print it. Because both traits are derivable traits, as mentioned, it is usually as straightforward as adding the `#[derive(Debug)]` to your struct or enum definition. See Appendix A about these and other derivable traits.

## Adding Custom Failure Messages

You can also add a custom message to be printed as optional arguments to the `assert!`, `assert_eq!`, `assert_ne!`, and `assert_matches!` arguments specified after the one required argument. The arguments to `assert_eq!` and `assert_ne!` are strings (discussed in Chapter 8 in the “Concatenation with Macro” section), so you can pass a format string with placeholders to go in those placeholders. Custom messages for assertions means; when a test fails, you'll have a more descriptive message with the code.

For example, let's say we have a function that greets a person and we want to test that the name we pass into the function appears in the greeting.

Filename: src/lib.rs

```
pub fn greeting(name: &str) -> String {
    format!("Hello {}!", name)
}

#[cfg(test)]
mod tests {
    use super::*;

    #[test]
    fn greeting_contains_name() {
        let result = greeting("Carol");
        assert!(result.contains("Carol"));
    }
}
```

The requirements for this program haven't been clear, but we want to have to update the test when the requirements change. Instead of checking for exact equality to the value returned by `greeting`, we just assert that the output contains the text of the name.

Let's introduce a bug into this code by changing

what this test failure looks like:

```
pub fn greeting(name: &str) -> String {
    String::from("Hello!")
}
```

Running this test produces the following:

```
running 1 test
test tests::greeting_contains_name ... FAILED

failures:

---- tests::greeting_contains_name stdout
      thread 'tests::greeting_contains_name'
    failed:
    result.contains("Carol")', src/lib.rs:12:8
note: Run with `RUST_BACKTRACE=1` for a backtrace

failures:
    tests::greeting_contains_name
```

This result just indicates that the assertion failed. A more useful failure message in this case would be provided by the `greeting` function. Let's change the test function to be made from a format string with a placeholder filled in from the `greeting` function:

```
#[test]
fn greeting_contains_name() {
    let result = greeting("Carol");
    assert!(
        result.contains("Carol"),
        "Greeting did not contain name, value was '{}'",
        result
    );
}
```

Now when we run the test, we'll get a more informative failure message:

```
---- tests::greeting_contains_name stdout
      thread 'tests::greeting_contains_name'
    not
    contain name, value was `Hello!`, src/lib.rs:12:8
note: Run with `RUST_BACKTRACE=1` for a backtrace
```

We can see the value we actually got in the test (the actual value of the `greeting` function) and what happened instead of what we were expecting.

## Checking for Panics with `should_panic`

In addition to checking that our code returns the important to check that our code handles error. We consider the `Guess` type that we created in Chapter 10. The `new` function that `Guess` depends on the guarantee that `new` returns a `Guess` between 1 and 100. We can write a test that ensures that `new` returns a `Guess` instance with a value outside that range.

We do this by adding another attribute, `should_panic`. The `should_panic` attribute makes a test pass if the code inside the function panics. In other words, the code inside the function doesn't panic.

Listing 11-8 shows a test that checks that the error occurs when we expect them to:

Filename: `src/lib.rs`

```
pub struct Guess {
    value: i32,
}

impl Guess {
    pub fn new(value: i32) -> Guess {
        if value < 1 || value > 100 {
            panic!("Guess value must be between 1 and 100: {} value);
        }
    }

    pub fn value(&self) -> i32 {
        self.value
    }
}

#[cfg(test)]
mod tests {
    use super::*;

    #[test]
    #[should_panic]
    fn greater_than_100() {
        Guess::new(200);
    }
}
```

Listing 11-8: Testing that a condition will cause a panic

We place the `#[should_panic]` attribute after the test function.



test function it applies to. Let's look at the result

```
running 1 test
test tests::greater_than_100 ... ok

test result: ok. 1 passed; 0 failed; 0 ignored
```

Looks good! Now let's introduce a bug in our code. The `new` function will panic if the value is greater than 100.

```
// --snip--

impl Guess {
    pub fn new(value: i32) -> Guess {
        if value < 1 {
            panic!("Guess value must be between 1 and 100")
        }
        Guess {
            value
        }
    }
}
```

When we run the test in Listing 11-8, it will fail:

```
running 1 test
test tests::greater_than_100 ... FAILED

failures:

failures:
    tests::greater_than_100

test result: FAILED. 0 passed; 1 failed; 0 ignored
     out
```

We don't get a very helpful message in this case. In the `new` function, we see that it's annotated with `#[should_panic]`, which means that the code in the test function did not cause a panic.

Tests that use `should_panic` can be imprecise because they only check if some code has caused some panic. A `should_panic` test can fail for a different reason than the one we were expecting. To make `should_panic` tests more precise, we can add a `panic_message` attribute to the `should_panic` attribute. The test harness will then check if the provided text contains the provided text. For example, consider

Listing 11-9 where the `new` function panics with whether the value is too small or too large:

Filename: `src/lib.rs`

```
// --snip--

impl Guess {
    pub fn new(value: i32) -> Guess {
        if value < 1 {
            panic!("Guess value must be greater than or equal to 1");
        } else if value > 100 {
            panic!("Guess value must be less than or equal to 100");
        }
    }

    Guess {
        value
    }
}

#[cfg(test)]
mod tests {
    use super::*;

    #[test]
    #[should_panic(expected = "Guess value must be less than or equal to 100")]
    fn greater_than_100() {
        Guess::new(200);
    }
}
```

Listing 11-9: Testing that a condition will cause a message

This test will pass because the value we put in the `expected` parameter is a substring of the message that the `new` function panics with. We could have specified the entire panic message, but in this case would be `Guess value must be less than or equal to 100`. What you choose to specify in the `expected` parameter depends on how much of the panic message is unique or distinctive to the test to be. In this case, a substring of the panic message is used. The code in the test function executes the `else if` branch, which causes the `panic!` function to be called, which then panics with the message `Guess value must be less than or equal to 100`.

To see what happens when a `should_panic` test let's again introduce a bug into our code by swapping the `if` and the `else if value > 100` blocks:

```
if value < 1 {
    panic!("Guess value must be less than
value);
} else if value > 100 {
    panic!("Guess value must be greater th
value);
}
```

This time when we run the `should_panic` test, it

```
running 1 test
test tests::greater_than_100 ... FAILED

failures:

---- tests::greater_than_100 stdout ----
thread 'tests::greater_than_100' panicked at 'assertion failed:
greater than or equal to 1, got 200.', src/lib.rs:10:13
note: Run with `RUST_BACKTRACE=1` for a backtrace
note: Panic did not include expected string 'Guess value must be
less than or equal to 1'
equal to 100'

failures:
    tests::greater_than_100

test result: FAILED. 0 passed; 1 failed; 0 ignored
out
```

The failure message indicates that this test did indeed panic but the panic message did not include the expected string `'Guess value must be less than or equal to 1'`. The string that it did get in this case was `'Guess value must be greater than or equal to 100'`. We're now figuring out where our bug is!

## Using `Result<T, E>` in tests

So far, we've written tests that panic when they fail. But we can also use `Result<T, E>` too! Here's that first example, but

```
#[cfg(test)]
mod tests {
    #[test]
    fn it_works() -> Result<(), String> {
        if 2 + 2 == 4 {
            Ok(())
        } else {
            Err(String::from("two plus two"))
        }
    }
}
```

Here, we've changed the `it_works` function to return a `Result` instead of `bool`. Instead of `assert_eq!`, we return `Ok(())` for the success case and `Err(String::from("two plus two"))` for the failure case. As before, this test will fail based on panics, it will use the `Result<T, E>` to propagate errors. For this, you can't use `#[should_panic]` with one of the functions, you'll be returning an `Err` instead!

Now that you know several ways to write tests, let's run our tests and explore the different options.

## Controlling How Tests Are Run

Just as `cargo run` compiles your code and then runs it, `cargo test` compiles your code in test mode and runs the tests. You can specify command line options to change the behavior. For example, the default behavior of the binary produced by `cargo test` is to run tests in parallel and capture output generated during the tests. You can change this by using `--no-capture` to prevent output from being displayed and making it easier to read the results.

Some command line options go to `cargo test`, and some go to the binary. To separate these two types of arguments, you use `--`. For example, `cargo test -- --help` runs `cargo test` followed by the separator `--` and then `--help`. Running `cargo test --help` displays the help for `cargo test`, and running `cargo test -- --help` displays the help for the binary after the separator `--`.

## Running Tests in Parallel or Consecutively

When you run multiple tests, by default they run in parallel. The tests will finish running faster so you can get your code working. Because the tests are run in parallel, the tests don't depend on each other or on any shared environment, such as the current working directory.

For example, say each of your tests runs some code that writes some data to a file called `test-output.txt` and asserts that the file contains a particular value. Because the tests run at the same time, one test might write to the file when another test reads the file. The code is incorrect but because the tests have been running in parallel. One solution is to make sure the tests run one at a time.

If you don't want to run the tests in parallel or if you want to limit the number of threads used, you can send the `--test-threads` flag to the test binary. For example:

```
$ cargo test -- --test-threads=1
```

We set the number of test threads to `1`, telling the test binary to run the tests in parallel. Running the tests using one thread instead of multiple threads in parallel, but the tests won't interfere with each other.

## Showing Function Output

By default, if a test passes, Rust's test library captures the output. For example, if we call `println!` in a test, the `println!` output in the terminal; we'll see the output of the test passed. If a test fails, we'll see whatever was printed before the failure message.

As an example, Listing 11-10 has a silly function that returns 10, as well as a test that passes and prints the output.

Filename: `src/lib.rs`

```

fn prints_and_returns_10(a: i32) -> i32 {
    println!("I got the value {}", a);
    10
}

#[cfg(test)]
mod tests {
    use super::*;

    #[test]
    fn this_test_will_pass() {
        let value = prints_and_returns_10(10);
        assert_eq!(10, value);
    }

    #[test]
    fn this_test_will_fail() {
        let value = prints_and_returns_10(5);
        assert_eq!(5, value);
    }
}

```

Listing 11-10: Tests for a function that calls `println!`

When we run these tests with `cargo test`, we'll

```

running 2 tests
test tests::this_test_will_pass ... ok
test tests::this_test_will_fail ... FAILED

failures:

---- tests::this_test_will_fail stdout ----
      I got the value 8
thread 'tests::this_test_will_fail' panicked at 'assertion failed:
    left: `5`,
    right: `10`, src/lib.rs:19:8
note: Run with `RUST_BACKTRACE=1` for a backtrace
failures:
    tests::this_test_will_fail

test result: FAILED. 1 passed; 1 failed; 0 ignored

```

Note that nowhere in this output do we see `I got the value 8` printed when the test that passes runs. That output is only printed from the test that failed, `I got the value 8`, as

summary output, which also shows the cause of

If we want to see printed values for passing tests, we can capture behavior by using the `--nocapture` flag

```
$ cargo test -- --nocapture
```

When we run the tests in Listing 11-10 again with the following output:

```
running 2 tests
I got the value 4
I got the value 8
test tests::this_test_will_pass ... ok
thread 'tests::this_test_will_fail' panicked at 'assertion failed:
  left: `5`,
  right: `10`', src/lib.rs:19:8
note: Run with `RUST_BACKTRACE=1` for a backtrace
test tests::this_test_will_fail ... FAILED

failures:

failures:
    tests::this_test_will_fail

test result: FAILED. 1 passed; 1 failed; 0 ignored
out
```

Note that the output for the tests and the test result is interleaved, which indicates that the tests are running in parallel, as we talked about earlier. We can control this using the `--test-threads=1` option and the `--nocapture` option. The output looks like then!

## Running a Subset of Tests by Name

Sometimes, running a full test suite can take a long time. If you're interested in a particular area, you might want to run only the tests in that area. You can choose which tests to run by passing `cargo test` the name of the test(s) you want to run as an argument.

To demonstrate how to run a subset of tests, we'll use the `test_name` function, as shown in Listing 11-11, and choose to run only the tests that use this function.

Filename: src/lib.rs

```

pub fn add_two(a: i32) -> i32 {
    a + 2
}

#[cfg(test)]
mod tests {
    use super::*;

    #[test]
    fn add_two_and_two() {
        assert_eq!(4, add_two(2));
    }

    #[test]
    fn add_three_and_two() {
        assert_eq!(5, add_two(3));
    }

    #[test]
    fn one_hundred() {
        assert_eq!(102, add_two(100));
    }
}

```

Listing 11-11: Three tests with three different names

If we run the tests without passing any arguments, they will all run in parallel:

```

running 3 tests
test tests::add_two_and_two ... ok
test tests::add_three_and_two ... ok
test tests::one_hundred ... ok

test result: ok. 3 passed; 0 failed; 0 ignored

```

## Running Single Tests

We can pass the name of any test function to `cargo test` to run only that test:

```

$ cargo test one_hundred
    Finished dev [unoptimized + debuginfo] target(s) in 0.1s
    Running target/debug/deps/adder-06a75e1b1e1e1e1e

running 1 test
test tests::one_hundred ... ok

test result: ok. 1 passed; 0 failed; 0 ignored

```



Only the test with the name `one_hundred` ran; the test `one_hundred` passed. The test output lets us know we had more tests by displaying `2 filtered out` at the end of the output.

We can't specify the names of multiple tests in the `cargo test` command. But there is a way to run multiple tests.

## Filtering to Run Multiple Tests

We can specify part of a test name, and any tests with that name will be run. For example, because two of our tests' names contain `add`, we can run two by running `cargo test add`:

```
$ cargo test add
    Finished dev [unoptimized + debuginfo] target(s) in 0.1s
     Running target/debug/deps/adder-06a75b1d4c1b1b1b

running 2 tests
test tests::add_two_and_two ... ok
test tests::add_three_and_two ... ok

test result: ok. 2 passed; 0 failed; 0 ignored
```

This command ran all tests with `add` in the name. Note that the module in which the tests are defined is `one_hundred`. Also note that the module in which the tests are defined is `one_hundred`, so we can run all the tests in a module by running `cargo test`.

## Ignoring Some Tests Unless Specifically

Sometimes a few specific tests can be very time-consuming to run. If you want to exclude them during most runs of `cargo test`, you can use the `ignore` attribute to exclude them. For example, in the file `src/lib.rs`:

Filename: src/lib.rs

```
#[test]
fn it_works() {
    assert_eq!(2 + 2, 4);
}

#[test]
#[ignore]
fn expensive_test() {
    // code that takes an hour to run
}
```

After `#[test]` we add the `#[ignore]` line to the `expensive_test` function. When we run our tests, `it_works` runs, but `expensive_test` is ignored.

```
$ cargo test
   Compiling adder v0.1.0 (file:///project)
   Finished dev [unoptimized + debuginfo] target(s) in 0.1s
   Running target/debug/deps/adder-ce99k...

running 2 tests
test expensive_test ... ignored
test it_works ... ok

test result: ok. 1 passed; 0 failed; 1 ignored
```

The `expensive_test` function is listed as `ignored`. If we want to run ignored tests, we can use `cargo test -- --ignored`:

```
$ cargo test -- --ignored
   Finished dev [unoptimized + debuginfo] target(s) in 0.1s
   Running target/debug/deps/adder-ce99k...

running 1 test
test expensive_test ... ok

test result: ok. 1 passed; 0 failed; 0 ignored
```

By controlling which tests run, you can make sure your tests are fast. When you're at a point where it makes sense to run a few tests and you have time to wait for the results, you can use `cargo test -- --ignored` instead.

## Test Organization

As mentioned at the start of the chapter, testing

people use different terminology and organization of tests in terms of two main categories: *unit tests* are small and more focused, testing one module in isolation and private interfaces. Integration tests are entirely external code in the same way any other external code would be, and potentially exercising multiple modules per test.

Writing both kinds of tests is important to ensure you're doing what you expect them to separately and together.

## Unit Tests

The purpose of unit tests is to test each unit of code to quickly pinpoint where code is and isn't working. We'll write tests in the `src` directory in each file with the code we want to test. We'll create a module named `tests` in each file and annotate the module with `cfg(test)`.

### The Tests Module and `#[cfg(test)]`

The `#[cfg(test)]` annotation on the tests module tells Rust to compile test code only when you run `cargo test`, not when you compile the library at compile time when you only want to build the library. Because unit tests and integration tests go in a different directory, they don't need the `#[cfg(test)]` annotation. However, because unit tests go in the `src` directory, we need `#[cfg(test)]` to specify that they shouldn't be compiled by default.

Recall that when we generated the new `adder` package in chapter 1, Cargo generated this code for us:

Filename: `src/lib.rs`

```
#[cfg(test)]
mod tests {
    #[test]
    fn it_works() {
        assert_eq!(2 + 2, 4);
    }
}
```

This code is the automatically generated test module. The `#[cfg(test)]` is a *configuration* and tells Rust that the following item

certain configuration option. In this case, the compiler provided by Rust for compiling and running tests only compiles our test code only if we actively run the tests or any helper functions that might be within this module annotated with `#[test]`.

## Testing Private Functions

There's debate within the testing community about whether private functions should be tested directly, and other languages don't allow testing private functions. Regardless of which testing idiom rules do allow you to test private functions. Consider the private function `internal_adder`:

Filename: `src/lib.rs`

```
pub fn add_two(a: i32) -> i32 {
    internal_adder(a, 2)
}

fn internal_adder(a: i32, b: i32) -> i32 {
    a + b
}

#[cfg(test)]
mod tests {
    use super::*;

    #[test]
    fn internal() {
        assert_eq!(4, internal_adder(2, 2))
    }
}
```

Listing 11-12: Testing a private function

Note that the `internal_adder` function is not modified. It's just Rust code and the `tests` module is just another module. Testing `internal_adder` in a test just fine. If you don't test it, there's nothing in Rust that will compel you to.

## Integration Tests

In Rust, integration tests are entirely external to the library being tested.

the same way any other code would, which means they are part of your library's public API. Their purpose is to ensure your library work together correctly. Units of code could have problems when integrated, so tests are as important as well. To create integration tests, you

## The *tests* Directory

We create a *tests* directory at the top level of our crate. Cargo knows to look for integration test files in this directory. We name test files as we want to in this directory, and Cargo runs each individual crate.

Let's create an integration test. With the code in Listing 11-13, make a *tests* directory, create a new file named *integration\_test.rs* with the code in Listing 11-13:

Filename: *tests/integration\_test.rs*

```
extern crate adder;

#[test]
fn it_adds_two() {
    assert_eq!(4, adder::add_two(2));
}
```

Listing 11-13: An integration test of a function in a crate

We've added `extern crate adder` at the top of *integration\_test.rs*. The reason is that each test in the *tests* directory needs to import our library into each of them.

We don't need to annotate any code in *tests/integration\_test.rs*. Cargo treats the *tests* directory specially and runs the tests when we run `cargo test`. Run `cargo test` now

```

$ cargo test
  Compiling adder v0.1.0 (file:///project
  Finished dev [unoptimized + debuginfo]
  Running target/debug/deps/adder-abcat

running 1 test
test tests::internal ... ok

test result: ok. 1 passed; 0 failed; 0 igr

    Running target/debug/deps/integration_test

running 1 test
test it_adds_two ... ok

test result: ok. 1 passed; 0 failed; 0 igr

    Doc-tests adder

running 0 tests

test result: ok. 0 passed; 0 failed; 0 igr

```

The three sections of output include the unit tests. The first section for the unit tests is the same for each unit test (one named `internal` that we added). The summary line for the unit tests.

The integration tests section starts with the line `Running target/debug/deps/integration_test` (the end of your output will be different). Next, there is one integration test and a summary line for the results. The `Doc-tests adder` section starts.

Similarly to how adding more unit test functions to the `tests` section, adding more test functions to the `integration_test.rs` file. Each integration test is a function that takes a string and returns a `Result`. So if we add more files in the `tests` directory, they will be added to the `tests` section.

We can still run a particular integration test function by passing its name as an argument to `cargo test`. To run all integration tests, use the `--test` argument of `cargo test`.

```
$ cargo test --test integration_test
    Finished dev [unoptimized + debuginfo]
    Running target/debug/integration_test
```

```
running 1 test
test it_adds_two ... ok
```

```
test result: ok. 1 passed; 0 failed; 0 ignored
```

This command runs only the tests in the *tests/integration\_test.rs* file.

## Submodules in Integration Tests

As you add more integration tests, you might want to use the *tests* directory to help organize them; for example, you might want to group tests that test the same functionality they're testing. As mentioned earlier, each test file is compiled as its own separate crate.

Treating each integration test file as its own crate is more like the way end users will be using your library. Files in the *tests* directory don't share the same behavior as files in the *src* directory learned in Chapter 7 regarding how to separate modules.

The different behavior of files in the *tests* directory means you can create a set of helper functions that would be useful in multiple tests. In this section, we try to follow the steps in the “Moving Modules to Their Own Files” section to extract them into a common module. For example, if we place a function named `setup` in it, we can add a `use` statement and call from multiple test functions in multiple test files.

Filename: *tests/common.rs*

```
pub fn setup() {
    // setup code specific to your library
}
```

When we run the tests again, we'll see a new section in the output for the *common.rs* file, even though this file doesn't contain any tests. The `setup` function from anywhere:

```

running 1 test
test tests::internal ... ok

test result: ok. 1 passed; 0 failed; 0 ignored
Running target/debug/deps/common-b8b6...

running 0 tests

test result: ok. 0 passed; 0 failed; 0 ignored
Running target/debug/deps/integration...

running 1 test
test it_adds_two ... ok

test result: ok. 1 passed; 0 failed; 0 ignored
Doc-tests adder

running 0 tests

test result: ok. 0 passed; 0 failed; 0 ignored

```

Having `common` appear in the test results with `r` is not what we wanted. We just wanted to share some files.

To avoid having `common` appear in the test output, in `tests/common.rs`, we'll create `tests/common/mod.rs`. In the "Filesystems" section of Chapter 7, we used the name `module_name/mod.rs` for files of modules that have submodules for `common` here, but naming the file `common` module as an integration test file. When we move `tests/common.rs` into `tests/common/mod.rs` and delete the `tests/common.rs` file, the output will no longer appear. Files in subdirectories will be compiled as separate crates or have sections in

After we've created `tests/common/mod.rs`, we can call test files as a module. Here's an example of calling the `it_adds_two` test in `tests/integration_test.rs`:

Filename: tests/integration\_test.rs



```
extern crate adder;

mod common;

#[test]
fn it_adds_two() {
    common::setup();
    assert_eq!(4, adder::add_two(2));
}
```

Note that the `mod common;` declaration is the same as demonstrated in Listing 7-4. Then in the test function `common::setup()` function.

## Integration Tests for Binary Crates

If our project is a binary crate that only contains a `src/lib.rs` file, we can't create integration tests in the `src` directory. We can use `extern crate` to import functions defined in the library crate and expose functions that other crates can call and use on their own.

This is one of the reasons Rust projects that provide a `src/main.rs` file that calls logic that lives in the `src/lib.rs` file. Integration tests *can* test the library crate by using the `extern crate` to import the library's important functionality. If the important functionality in the `src/main.rs` file will work as well, and the library's functionality need to be tested.

## Summary

Rust's testing features provide a way to specify how the code continues to work as you expect, even as you modify it. You can test different parts of a library separately and can test the library as a whole. Integration tests check that many parts of the library work together as they use the library's public API to test the code that uses the library. Even though Rust's type system and ownership system help prevent bugs, tests are still important to reduce logic errors and ensure the code is expected to behave.

Let's combine the knowledge you learned in this chapter and work on a project!

# An I/O Project: Building a Command Line Program

This chapter is a recap of the many skills you’ve learned, plus a few more standard library features. We’ll build a program that uses file and command line input/output to practice what you now have under your belt.

Rust’s speed, safety, single binary output, and cross-platform language for creating command line tools, so for this chapter we’ll build a version of the classic command line tool `grep` (`g` for `get` and `print`). In the simplest use case, `grep` searches a file for a string. To do so, `grep` takes as its arguments a file name and a string. The file, finds lines in that file that contain the string.

Along the way, we’ll show how to make our command line program a terminal that many command line tools use. We’ll use an environment variable to allow the user to configure the behavior of the program. The standard error console stream ( `stderr` ) instead of the standard output stream. For example, the user can redirect successful output messages onscreen.

One Rust community member, Andrew Gallant, has written a very fast version of `grep`, called `ripgrep`. By comparison, our program is fairly simple, but this chapter will give you some insight into the need to understand a real-world project such as this.

Our `grep` project will combine a number of concepts you’ve learned:

- Organizing code (using what you learned in Chapter 8)
- Using vectors and strings (collections, Chapter 8)
- Handling errors (Chapter 9)
- Using traits and lifetimes where appropriate (Chapter 10)
- Writing tests (Chapter 11)

We’ll also briefly introduce closures, iterators, and `VecDeque`, and Chapter 17 will cover in detail.

## Accepting Command Line Arguments

Let’s create a new project with, as always, `cargo` and `rust`.

to distinguish it from the `grep` tool that you might

```
$ cargo new minigrep
    Created binary (application) `minigrep`
$ cd minigrep
```

The first task is to make `minigrep` accept its two filename and a string to search for. That is, we want to accept a filename, a string to search for, and a pattern.

```
$ cargo run searchstring example-filename.
```

Right now, the program generated by `cargo new` is a simple `main.rs` file. Some existing libraries on [Crates.io](https://crates.io) can help with parsing command line arguments, but because you're just getting started, we'll implement this capability ourselves.

## Reading the Argument Values

To enable `minigrep` to read the values of command line arguments, we'll need a function provided in Rust's standard library. This function returns an *iterator* of the command line arguments. We haven't discussed iterators yet (we will in Chapter 13), but for now, you only need to know two details: an iterator is a series of values, and we can call the `collect` method to collect the iterator into a collection, such as a vector, containing all the elements.

Use the code in Listing 12-1 to allow your `minigrep` program to accept command line arguments passed to it and then collect the arguments into a vector.

Filename: `src/main.rs`

```
use std::env;

fn main() {
    let args: Vec<String> = env::args().collect();
    println!("{:?}", args);
}
```

Listing 12-1: Collecting the command line arguments

First, we bring the `std::env` module into scope with the `use` statement. Then, we call its `args` function. Notice that the `std::env::args` function returns an iterator. As we discussed in Chapter 7, in cases

in more than one module, it's conventional to bring the function into scope rather than the function. By doing so, we can easily refer to `std::env`. It's also less ambiguous than adding the function with just `args`, because `args` might be a function that's defined in the current module.

---

## The `args` Function and Invalid Unicode

Note that `std::env::args` will panic if any argument in your program needs to accept arguments corresponding to `std::env::args_os` instead. That function returns `OsString` values instead of `String` values. We use `std::env::args` here for simplicity, because `OsString` values are more complex to work with.

---

On the first line of `main`, we call `env::args`, and then we turn the iterator into a vector containing all the arguments. We can use the `collect` function to create many kinds of collections. We annotate the type of `args` to specify that we want a `Vec`. We very rarely need to annotate types in Rust, but we need to annotate because Rust isn't able to infer the type of `args`.

Finally, we print the vector using the debug formatter. We first print with no arguments and then with two arguments.

```
$ cargo run
--snip--
["target/debug/minigrep"]

$ cargo run needle haystack
--snip--
["target/debug/minigrep", "needle", "haystack"]
```

Notice that the first value in the vector is `"target/debug/minigrep"`, the name of our binary. This matches the behavior of other programs that use the name by which they were invoked to access the program name or change behavior of the program based on the name used to invoke the program. But for the purpose of this example, we save only the two arguments we need.

## Saving the Argument Values in Variables

Printing the value of the vector of arguments illustrates how we can access the values specified as command line arguments. We can store the values of the two arguments in variables so we can use them for the rest of the program. We do that in Listing 12-2:

Filename: src/main.rs

```
use std::env;

fn main() {
    let args: Vec<String> = env::args().collect();

    let query = &args[1];
    let filename = &args[2];

    println!("Searching for {}", query);
    println!("In file {}", filename);
}
```

Listing 12-2: Creating variables to hold the query and filename

As we saw when we printed the vector, the program starts at index 0, so we're starting at index 1. The first argument that the program takes is the string we're searching for, so we put that into the variable `query`. The second argument will be the filename, so we put that into the variable `filename`.

We temporarily print the values of these variables to make sure we've stored them as we intend. Let's run this program again with the same arguments:

```
$ cargo run test sample.txt
   Compiling minigrep v0.1.0 (file:///project/minigrep)
   Finished dev [unoptimized + debuginfo] target(s) in 0.15s
  Running `target/debug/minigrep test sample.txt`
Searching for test
In file sample.txt
```

Great, the program is working! The values of the arguments have been stored into the right variables. Later we'll add some error handling for potential erroneous situations, such as when the user provides more than two arguments. For now, we'll ignore that situation and work on adding more functionality.

## Reading a File

Now we'll add functionality to read the file that is the command-line argument. First, we need a sample file to test. To make sure `minigrep` is working is one with a `src/main.rs` with some repeated words. Listing 12-3 has an example. Well! Create a file called `poem.txt` at the root level of the project. "I'm Nobody! Who are you?"

Filename: poem.txt

```
I'm nobody! Who are you?  
Are you nobody, too?  
Then there's a pair of us – don't tell!  
They'd banish us, you know.
```

```
How dreary to be somebody!  
How public, like a frog  
To tell your name the livelong day  
To an admiring bog!
```

Listing 12-3: A poem by Emily Dickinson makes a good test file

With the text in place, edit `src/main.rs` and add code to read the file. Listing 12-4:

Filename: `src/main.rs`

```
use std::env;  
use std::fs;  
  
fn main() {  
    // --snip--  
    println!("In file {}", filename);  
  
    let contents = fs::read_to_string(filename).expect("Something went wrong reading the file");  
  
    println!("With text:\n{}", contents);  
}
```

Listing 12-4: Reading the contents of the file specified by the command-line argument

First, we add another `use` statement to bring in the `std::fs` library: we need `std::fs` to handle files.

In `main`, we've added a new statement: `fs::read_to_string` to open that file, and then produce a new `String` containing the contents of the file.

After that line, we've again added a temporary `contents` value of `contents` after the file is read, so we can use it so far.

Let's run this code with any string as the first command-line argument (we haven't implemented the searching part yet) and see the output:

```
$ cargo run the poem.txt
   Compiling minigrep v0.1.0 (file:///project/minigrep)
   Finished dev [unoptimized + debuginfo] target(s) in 0.15s
   Running `target/debug/minigrep the poem.txt`
Searching for the
In file poem.txt
With text:
I'm nobody! Who are you?
Are you nobody, too?
Then there's a pair of us – don't tell!
They'd banish us, you know.

How dreary to be somebody!
How public, like a frog
To tell your name the livelong day
To an admiring bog!
```

Great! The code read and then printed the contents of the file. However, there are several flaws. The `main` function has multiple responsibilities, which makes it harder to understand and easier to maintain if each function is responsible for a single task. One problem is that we're not handling errors as well as we could, so these flaws aren't a big problem, but as we grow our program, it's good practice to begin refactoring to fix them cleanly. It's good practice to begin refactoring a program, because it's much easier to refactor something that is already working than to refactor something that is not working next.

## Refactoring to Improve Module Structure and Error Handling

To improve our program, we'll fix four problems: the module structure and how it's handling potential errors.

First, our `main` function now performs two tasks: reading files and searching for text. For such a small function, this isn't a major problem, but as we grow our program inside `main`, the number of responsibilities it handles will increase. As a function gains responsibilities,

reason about, harder to test, and harder to char  
It's best to separate functionality so each func

This issue also ties into the second problem: alth  
configuration variables to our program, variable  
the program's logic. The longer `main` becomes,  
into scope; the more variables we have in scope  
the purpose of each. It's best to group the config  
to make their purpose clear.

The third problem is that we've used `expect` to  
the file fails, but the error message just prints `s`  
can fail in a number of ways: for example, the fil  
have permission to open it. Right now, regardless  
`something went wrong` error message, which w

Fourth, we use `expect` repeatedly to handle dif  
program without specifying enough arguments,  
error from Rust that doesn't clearly explain the p  
error-handling code were in one place so future  
consult in the code if the error-handling logic ne  
handling code in one place will also ensure that  
meaningful to our end users.

Let's address these four problems by refactoring

## Separation of Concerns for Binary Proj

The organizational problem of allocating respon  
function is common to many binary projects. As  
developed a process to use as a guideline for sp  
binary program when `main` starts getting large.

- Split your program into a *main.rs* and a *lib.rs*.
- As long as your command line parsing logic
- When the command line parsing logic start  
*main.rs* and move it to *lib.rs*.
- The responsibilities that remain in the `mai`  
be limited to the following:



- Calling the command line parsing logic
- Setting up any other configuration
- Calling a `run` function in `lib.rs`
- Handling the error if `run` returns an

This pattern is about separating concerns: `main.rs` handles all the logic of the task at hand. `lib.rs` handles all the logic of the task at hand. By putting functions directly, this structure lets you test all of them into functions in `lib.rs`. The only code that remains in `main.rs` is to verify its correctness by reading it. Let's rework it.

## Extracting the Argument Parser

We'll extract the functionality for parsing arguments to prepare for moving the command line parsing logic to the new start of `main` that calls a new function `src/main.rs` for the moment.

Filename: `src/main.rs`

```
fn main() {
    let args: Vec<String> = env::args().collect();

    let (query, filename) = parse_config(&args);

    // --snip--
}

fn parse_config(args: &[String]) -> (&str, &str) {
    let query = &args[1];
    let filename = &args[2];

    (query, filename)
}
```

Listing 12-5: Extracting a `parse_config` function

We're still collecting the command line arguments and assigning the argument value at index `1` to the variable `query` and the value at index `2` to the variable `filename` with the whole vector to the `parse_config` function. The logic that determines which argument goes to which variable is now in `parse_config`. The values back to `main`. We still create the `query` and `filename` variables in `main` but `main` no longer has the responsibility of determining which arguments and variables correspond.

This rework may seem like overkill for our small, incremental steps. After making this change, we can ensure that the argument parsing still works. It's good to identify the cause of problems when they occur.

## Grouping Configuration Values

We can take another small step to improve the code. At the moment, we're returning a tuple, but then we have to access individual parts again. This is a sign that perhaps we should do better yet.

Another indicator that shows there's room for improvement is the `parse_config` function, which implies that the two values are part of one configuration value. We're not currently using the structure of the data other than by grouping the two values into one struct and give each of them a name. Doing so will make it easier for future maintainers to understand what different values relate to each other and what they do.

---

Note: Some people call this anti-pattern of using a tuple to return a complex type would be more appropriate *prima facie*.

---

Listing 12-6 shows the addition of a struct name named `Query` and `Filename`. We've also changed the `parse_config` function to return an instance of the `Config` struct and update the `main` function rather than having separate variables:

Filename: src/main.rs

```

fn main() {
    let args: Vec<String> = env::args().collect();

    let config = parse_config(&args);

    println!("Searching for {}", config.query);
    println!("In file {}", config.filename);

    let contents = fs::read_to_string(config.filename)
        .expect("Something went wrong reading the file");

    // --snip--
}

struct Config {
    query: String,
    filename: String,
}

fn parse_config(args: &[String]) -> Config {
    let query = args[1].clone();
    let filename = args[2].clone();

    Config { query, filename }
}

```

Listing 12-6: Refactoring `parse_config` to return a `Config` struct

The signature of `parse_config` now indicates that the body of `parse_config`, where we used to return values in `args`, we now define `Config` to contain the data. The variable in `main` is the owner of the argument values, so the `parse_config` function borrows them, which means it can't take ownership of the values. If `Config` tried to take ownership of the values, it would be a compile-time error.

We could manage the `String` data in a number of ways, though somewhat inefficient, one route is to call the `clone` method on the `String` values. This will make a full copy of the data for the `Config` struct. This takes more time and memory than storing a reference to the data, but it also makes our code very straightforward to read and understand. The lifetimes of the references; in this circumstance, simplicity is a worthwhile trade-off.

---

## The Trade-Offs of Using `clone`

There's a tendency among many Rustaceans to avoid ownership problems because of its runtime cost. It's tempting to use more efficient methods in this type of situation: copy a few strings to continue making progress, make copies only once and your filename and query are ready. You have a working program that's a bit inefficient on your first pass. As you become more experienced, you start with the most efficient solution, but for now, let's use `clone`.

---

We've updated `main` so it places the instance of `Config` into a variable named `config`, and we updated `parse_config` to use separate `query` and `filename` variables so it no longer returns a `struct` instead.

Now our code more clearly conveys that `query` and `filename` values know their purpose is to configure how the program works. The `config` instance knows values to find them in the `config` instance's purpose.

## Creating a Constructor for `Config`

So far, we've extracted the logic responsible for parsing from `main` and placed it in the `parse_config` function. We then added that the `query` and `filename` values were related to the `Config` struct. We then added a `Config` struct of `query` and `filename` and to be able to return names from the `parse_config` function.

So now that the purpose of the `parse_config` function is to create a `Config` instance, we can change `parse_config` from a function to a `new` that is associated with the `Config` struct. This is more idiomatic. We can create instances of type `String`, by calling `String::new`. Similarly, by calling `Config::new`, we'll be able to create a `Config` instance. Listing 12-7 shows the change.

Filename: src/main.rs

```

fn main() {
    let args: Vec<String> = env::args().collect();

    let config = Config::new(&args);

    // --snip--
}

// --snip--

impl Config {
    fn new(args: &[String]) -> Config {
        let query = args[1].clone();
        let filename = args[2].clone();

        Config { query, filename }
    }
}

```

Listing 12-7: Changing `parse_config` into `Config`

We've updated `main` where we were calling `parse_config` to `Config::new`. We've changed the name of `parse_config` to `new` and added an `impl` block, which associates the `new` function with the `Config` struct again to make sure it works.

## Fixing the Error Handling

Now we'll work on fixing our error handling. Recall that we were checking for values in the `args` vector at index `1` or index `2`, but the vector contains fewer than three items. Try running the program with no arguments; it will look like this:

```

$ cargo run
   Compiling minigrep v0.1.0 (file:///project/minigrep)
   Finished dev [unoptimized + debuginfo] target(s) in 0.15s
  Running `target/debug/minigrep`
thread 'main' panicked at 'index out of bounds: the len is 1 but the index is 1', src/main.rs:29:21
note: Run with `RUST_BACKTRACE=1` for a backtrace.

```

The line `index out of bounds: the len is 1 but the index is 1` is the error message intended for programmers. It won't help you understand what happened and what they should do instead. Let's improve the error message.

## Improving the Error Message

In Listing 12-8, we add a check in the `new` function enough before accessing index `1` and `2`. If the function panics and displays a better error message than

Filename: `src/main.rs`

```
// --snip--
fn new(args: &[String]) -> Config {
    if args.len() < 3 {
        panic!("not enough arguments");
    }
    // --snip--
```

Listing 12-8: Adding a check for the number of arguments

This code is similar to the `Guess::new` function we saw in Listing 9-9. It is called `panic!` when the `value` argument was called. Instead of checking for a range of values here, we check if the number of arguments is at least `3` and the rest of the function can operate. If the condition has been met. If `args` has fewer than `3` arguments, the condition is not true, and we call the `panic!` macro to end the program.

With these extra few lines of code in `new`, let's run the program with arguments again to see what the error looks like.

```
$ cargo run
   Compiling minigrep v0.1.0 (file:///project/minigrep)
   Finished dev [unoptimized + debuginfo] target(s) in 0.15s
   Running `target/debug/minigrep`
thread 'main' panicked at 'not enough arguments', src/main.rs:12:14
note: Run with `RUST_BACKTRACE=1` for a backtrace
```

This output is better: we now have a reasonable error message. However, we have extraneous information we don't want to get. The technique we used in Listing 9-9 isn't the best to handle errors. It's not appropriate for a programming problem rather than a system problem. In Chapter 9. Instead, we can use the other technique for handling errors—returning a `Result` that indicates either success or failure.

## Returning a `Result` from `new` Instead of Calling `panic!`

We can instead return a `Result` value that will contain the value in the successful case and will describe the problem in the error case. When communicating to `main`, we can use the `Result` type. Then we can change `main` to convert an `Err` value into a panic.

our users without the surrounding text about `tl` that a call to `panic!` causes.

Listing 12-9 shows the changes we need to make and the body of the function needed to return a `Result` until we update `main` as well, which we'll do in the next chapter.

Filename: `src/main.rs`

```
impl Config {
    fn new(args: &[String]) -> Result<Config> {
        if args.len() < 3 {
            return Err("not enough arguments")
        }

        let query = args[1].clone();
        let filename = args[2].clone();

        Ok(Config { query, filename })
    }
}
```

Listing 12-9: Returning a `Result` from `Config::new`

Our `new` function now returns a `Result` with a `Config` and a `&'static str` in the error case. Recall from Chapter 10 that `&'static str` is the type of string message type for now.

We've made two changes in the body of the `new` function when the user doesn't pass enough arguments, we've wrapped the `Config` return value in an `Ok` to conform to its new type signature.

Returning an `Err` value from `Config::new` allows a `Result` value returned from the `new` function to handle the error case.

## Calling `Config::new` and Handling Errors

To handle the error case and print a user-friendly message, we'll need to handle the `Result` being returned by `Config::new`. We'll also take the responsibility of exiting the process with an error code from `panic!` and implement it by handling the error state.

Filename: src/main.rs

```
use std::process;

fn main() {
    let args: Vec<String> = env::args().collect();

    let config = Config::new(&args).unwrap_or_else(|err| {
        println!("Problem parsing arguments: {}", err);
        process::exit(1);
    });

    // --snip--
}
```

Listing 12-10: Exiting with an error code if creating a configuration fails

In this listing, we've used a method we haven't covered yet, `unwrap_or_else`, which is defined on `Result<T, E>` by the standard library. This method allows us to define some custom, non-`panic!` behavior. If the value is an `Err` value, this method's behavior is similar to `unwrap_or`, but instead of returning a default value, this method's behavior is similar to `unwrap` and `unwrap_or` wrapping. However, if the value is an `Err` value, this method calls the `closure`, which is an anonymous function we define in the argument `closure`. We'll cover closures in more detail in chapter 13. We need to know that `unwrap_or_else` will pass the error value to the closure. The case is the static string `not enough arguments` to the closure in the argument `err` that appears between the `unwrap_or_else` and the closure. The closure can then use the `err` value when it runs.

We've added a new `use` line to import `process` from the `std` crate. In the closure that will be run in the error case is `process::exit(1)`. The `process::exit(1)` program immediately and return the number 1. This is similar to the `panic!`-based handling we saw in the previous listing. Let's get all the extra output. Let's try it:

```
$ cargo run
   Compiling minigrep v0.1.0 (file:///project/target/debug)
   Finished dev [unoptimized + debuginfo] target(s) in 0.15s
  Running `target/debug/minigrep`
Problem parsing arguments: not enough arguments
```

Great! This output is much friendlier for our use.

## Extracting Logic from `main`



Now that we've finished refactoring the configuration program's logic. As we stated in "Separation of Concerns", we can extract a function named `run` that will hold all the logic for the function that isn't involved with setting up configuration. Once we're done, `main` will be concise and easy to verify and write tests for all the other logic.

Listing 12-11 shows the extracted `run` function. This is an incremental improvement of extracting the function from `src/main.rs`.

Filename: `src/main.rs`

```
fn main() {  
    // --snip--  
  
    println!("Searching for {}", config.query);  
    println!("In file {}", config.filename);  
  
    run(config);  
}  
  
fn run(config: Config) {  
    let contents = fs::read_to_string(config.filename)  
        .expect("something went wrong reading the file");  
  
    println!("With text:\n{}", contents);  
}  
  
// --snip--
```

Listing 12-11: Extracting a `run` function containing the remaining logic

The `run` function now contains all the remaining logic for reading the file. The `run` function takes the `Config` struct as an argument.

## Returning Errors from the `run` Function

With the remaining program logic separated into its own function, we can now focus on the error handling, as we did with `Config::new`. In the original program, the program would panic by calling `expect`, the `run` function would panic when something goes wrong. This will let us figure out how to handle errors in a user-friendly way. Let's look at the changes we need to make to the signature and body of `run`.

Filename: `src/main.rs`

```

use std::error::Error;

// --snip--

fn run(config: Config) -> Result<(), Box<dyn Error>> {
    let contents = fs::read_to_string(config.path)?;

    println!("With text:\n{}", contents);

    Ok(())
}

```

Listing 12-12: Changing the `run` function to return a `Result`

We've made three significant changes here. First, we changed the `run` function to `Result<(), Box<dyn Error>>`. The `Result` type is a unit type, `()`, and we keep that as the value returned when the function succeeds.

For the error type, we used the *trait object* `Box<dyn Error>`. We brought `std::error::Error` into scope with a `use` statement. `Box<dyn Error>` is a *trait object* type, which means it will return a type that implements the `Error` trait. We use `dyn` to indicate that the particular type the return value will be. This gives us a "dynamic" error type that may be of different types in different error cases. `dyn` is short for "dynamic."

Second, we've removed the call to `expect` in favor of `?`. We learned about `expect` in Chapter 9. Rather than `panic!` on an error, `?` will return an `Err` value to the caller to handle.

Third, the `run` function now returns an `Ok` value. We wrap the unit type value in the `Ok` value. This is the idiomatic way to wrap a value in a `Result` at first, but using `Ok` like this is the idiomatic way to wrap a value in a `Result` for its side effects only; it doesn't return a value.

When you run this code, it will compile but will display a warning:

```

warning: unused `std::result::Result` which
--> src/main.rs:18:5
   |
18 |     run(config);
   |     ^^^^^^^^^^^^^
= note: #[warn(unused_must_use)] on by default

```

Rust tells us that our code ignored the `Result` value. We can use `unwrap` to indicate that an error occurred. But we're not ch

was an error, and the compiler reminds us that we need to add error handling code here! Let's rectify that problem.

## Handling Errors Returned from `run` in `main`

We'll check for errors and handle them using a `match` expression in Listing 12-10, but with a slight difference.

Filename: `src/main.rs`

```
fn main() {  
    // --snip--  
  
    println!("Searching for {}", config.query);  
    println!("In file {}", config.filename);  
  
    if let Err(e) = run(config) {  
        println!("Application error: {}", e);  
  
        process::exit(1);  
    }  
}
```

We use `if let` rather than `unwrap_or_else` to handle the `Err` value and call `process::exit(1)` if it does. The `run` function that we want to `unwrap` in the same way that `Config::new` returns a `Config` instance. Because `run` returns `Result` in the success case and `Err` in the error case, we don't need `unwrap_or_else` to return a default value; we would only be `unwrap`ing the `Ok` case.

The bodies of the `if let` and the `unwrap_or_else` cases are identical: we print the error and exit.

## Splitting Code into a Library Crate

Our `minigrep` project is looking good so far! Now we'll move some code into the `src/lib.rs` file so we can test it in fewer responsibilities.

Let's move all the code that isn't the `main` function into a new module.

- The `run` function definition
- The relevant `use` statements
- The definition of `Config`

- The `Config::new` function definition

The contents of `src/lib.rs` should have the signature (omitted the bodies of the functions for brevity).  
modify `src/main.rs` in the listing after this one.

Filename: `src/lib.rs`

```
use std::error::Error;
use std::fs;

pub struct Config {
    pub query: String,
    pub filename: String,
}

impl Config {
    pub fn new(args: &[String]) -> Result<Config, Error> {
        // --snip--
    }
}

pub fn run(config: Config) -> Result<(), Error> {
    // --snip--
}
```

Listing 12-13: Moving `Config` and `run` into `src/lib.rs`

We've made liberal use of the `pub` keyword: on the `Config` struct, on the `new` method, and on the `run` function. We now have that we can test!

Now we need to bring the code we moved to `src/lib.rs` into `src/main.rs`, as shown in Listing 12-14:

Filename: `src/main.rs`

```
extern crate minigrep;

use std::env;
use std::process;

use minigrep::Config;

fn main() {
    // --snip--
    if let Err(e) = minigrep::run(config) {
        // --snip--
    }
}
```

Listing 12-14: Bringing the `minigrep` crate into t

To bring the library crate into the binary crate, we add a `use minigrep::Config` line to bring the prefix the `run` function with our crate name. Now connected and should work. Run the program with everything works correctly.

Whew! That was a lot of work, but we've set our Now it's much easier to handle errors, and we've Almost all of our work will be done in `src/lib.rs` fr

Let's take advantage of this newfound modularity have been difficult with the old code but is easy tests!

## Developing the Library's Functionality Driven Development

Now that we've extracted the logic into `src/lib.rs` error handling in `src/main.rs`, it's much easier to of our code. We can call functions directly with v values without having to call our binary from the some tests for the functionality in the `Config::i`

In this section, we'll add the searching logic to the Test-driven development (TDD) process. This section follows these steps:

1. Write a test that fails and run it to make sure
2. Write or modify just enough code to make
3. Refactor the code you just added or changed to pass.
4. Repeat from step 1!

This process is just one of many ways to write software design as well. Writing the test before you write helps to maintain high test coverage throughout

We'll test drive the implementation of the function searching for the query string in the file content: match the query. We'll add this functionality in a

## Writing a Failing Test

Because we don't need them anymore, let's remove `src/lib.rs` and `src/main.rs` that we used to check that `src/lib.rs`, we'll add a `tests` module with a test function specifies the behavior we want the query and the text to search for the query in, and a text that contains the query. Listing 12-15 shows

Filename: `src/lib.rs`

```
#[cfg(test)]
mod tests {
    use super::*;

    #[test]
    fn one_result() {
        let query = "duct";
        let contents = "\
Rust:
safe, fast, productive.
Pick three.";

        assert_eq!(
            vec!["safe, fast, productive."],
            search(query, contents)
        );
    }
}
```

Listing 12-15: Creating a failing test for the `search`

This test searches for the string `"duct"`. The text one of which contains `"duct"`. We assert that the function contains only the line we expect.

We aren't able to run this test and watch it fail because the `search` function doesn't exist yet! So now we test to compile and run by adding a definition of `search` that returns an empty vector, as shown in Listing 12-16. This test fails because an empty vector doesn't match a vector containing `"safe, fast, productive."`

Filename: `src/lib.rs`

```
fn search<'a>(query: &str, contents: &'a str)
    -> Vec![]
}
```

Listing 12-16: Defining just enough of the `search`

Notice that we need an explicit lifetime `'a` defined used with the `contents` argument and the return type. The lifetime parameters specify which argument lifetime the return value. In this case, we indicate that the return value is a slice of the `contents` argument (the string slice that references slices of the argument `contents`).

In other words, we tell Rust that the data returned is valid as long as the data passed into the `search` function is valid. The data referenced by a slice needs to be valid; if the compiler assumes we're making string slices of `contents`, it will do its safety checking incorrectly.

If we forget the lifetime annotations and try to compile, we get an error:

```
error[E0106]: missing lifetime specifier
  --> src/lib.rs:5:51
   |
5  | fn search(query: &str, contents: &str)
   |                                ^
   |                                = help: this function's return type contains a lifetime, but no lifetime parameter is provided
   |                                = help: this function's return type contains a lifetime, but no lifetime parameter is provided
```

Rust can't possibly know which of the two arguments is the argument that contains the parts of that text that match, we know `contents` is connected to the return value using the lifetime annotation.

Other programming languages don't require you to specify lifetimes in the signature. So although this might seem a bit tedious, you might want to compare this example with the "Lifetimes" section in Chapter 10.

Now let's run the test:

```

$ cargo test
   Compiling minigrep v0.1.0 (file:///proj
--warnings--
   Finished dev [unoptimized + debuginfo]
   Running target/debug/deps/minigrep-ak

running 1 test
test tests::one_result ... FAILED

failures:

---- tests::one_result stdout ----
thread 'tests::one_result' panicked
==
right)`
left: `["safe, fast, productive."]`,
right: `[]`), src/lib.rs:48:8
note: Run with `RUST_BACKTRACE=1` for a ba

failures:
    tests::one_result

test result: FAILED. 0 passed; 1 failed; 0
out

error: test failed, to rerun pass '--lib'

```

Great, the test fails, exactly as we expected. Let's

## Writing Code to Pass the Test

Currently, our test is failing because we always return an empty vector. To pass the test, we need to implement `search`, our program needs to follow these steps:

- Iterate through each line of the contents.
- Check whether the line contains our query.
- If it does, add it to the list of values we're returning.
- If it doesn't, do nothing.
- Return the list of results that match.

Let's work through each step, starting with iterating through the lines of the file.

### Iterating Through Lines with the `lines` Method

Rust has a helpful method to handle line-by-line iteration named `lines`, that works as shown in Listing 12.



Filename: src/lib.rs

```
fn search<'a>(query: &str, contents: &'a str) {
    for line in contents.lines() {
        // do something with line
    }
}
```

Listing 12-17: Iterating through each line in `contents`

The `lines` method returns an iterator. We'll talk about iterators in chapter 13, but recall that you saw this way of using an iterator with a `for` loop with an iterator to run some code on each element.

## Searching Each Line for the Query

Next, we'll check whether the current line contains the query. Strings have a helpful method named `contains`. We'll use the `contains` method in the `search` function, as shown in Listing 12-18. It won't compile yet:

Filename: src/lib.rs

```
fn search<'a>(query: &str, contents: &'a str) {
    for line in contents.lines() {
        if line.contains(query) {
            // do something with line
        }
    }
}
```

Listing 12-18: Adding functionality to see whether a line contains the `query`

## Storing Matching Lines

We also need a way to store the lines that contain the query. We'll make a mutable vector before the `for` loop and push each `line` in the vector. After the `for` loop, we return the vector. Listing 12-19 shows the code:

Filename: src/lib.rs

```
fn search<'a>(query: &str, contents: &'a s
    let mut results = Vec::new();

    for line in contents.lines() {
        if line.contains(query) {
            results.push(line);
        }
    }

    results
}
```

Listing 12-19: Storing the lines that match so we

Now the `search` function should return only the test should pass. Let's run the test:

```
$ cargo test
--snip--
running 1 test
test tests::one_result ... ok

test result: ok. 1 passed; 0 failed; 0 ignored
```

Our test passed, so we know it works!

At this point, we could consider opportunities for the search function while keeping the tests passing functionality. The code in the search function is at the advantage of some useful features of iterators. We'll explore these in Chapter 13, where we'll explore iterators in detail.

## Using the `search` Function in the `run` Function

Now that the `search` function is working and tested, we can use it in our `run` function. We need to pass the `config.run` reads from the file to the `search` function. The `search` function returns a `Vec` of lines that match the query.

Filename: `src/lib.rs`

```
pub fn run(config: Config) -> Result<(), Error> {
    let contents = fs::read_to_string(config.filename)?;

    for line in search(&config.query, &contents) {
        println!("{}", line);
    }

    Ok(())
}
```

We're still using a `for` loop to return each line found.

Now the entire program should work! Let's try it with the poem and return exactly one line from the Emily Dickinson poem.

```
$ cargo run frog poem.txt
   Compiling minigrep v0.1.0 (file:///projects/minigrep)
   Finished dev [unoptimized + debuginfo] target(s) in 0.15s
   Running `target/debug/minigrep frog poem.txt`
How public, like a frog
```

Cool! Now let's try a word that will match multiple lines.

```
$ cargo run body poem.txt
   Finished dev [unoptimized + debuginfo] target(s) in 0.15s
   Running `target/debug/minigrep body poem.txt`
I'm nobody! Who are you?
Are you nobody, too?
How dreary to be somebody!
```

And finally, let's make sure that we don't get any matches that aren't anywhere in the poem, such as "monomorphization".

```
$ cargo run monomorphization poem.txt
   Finished dev [unoptimized + debuginfo] target(s) in 0.15s
   Running `target/debug/minigrep monomorphization poem.txt`
```

Excellent! We've built our own mini version of a command line program. We've also learned how to structure applications. We've also learned about lifetimes, testing, and command line parsing.

To round out this project, we'll briefly demonstrate how to use environment variables and how to print to standard error, both of which are useful when writing command line programs.

## Working with Environment Variables

We'll improve `minigrep` by adding an extra feature searching that the user can turn on via an environment variable. We'll add a command line option and require that it be used to apply, but instead we'll use an environment variable. We'll set the environment variable once and have all tests pass in that terminal session.

## Writing a Failing Test for the Case-Insensitive Search

We want to add a new `search_case_insensitive` environment variable. We'll continue to follow the same pattern as before. We'll again write a failing test. We'll add a new test: `search_case_insensitive` function and rename `case_sensitive` to clarify the differences between the two. We'll do this between lines 12-20:

Filename: `src/lib.rs`

```

#[cfg(test)]
mod tests {
    use super::*;

    #[test]
    fn case_sensitive() {
        let query = "duct";
        let contents = "\
Rust:
safe, fast, productive.
Pick three.
Duct tape.";

        assert_eq!(
            vec!["safe, fast, productive.",
                search(query, contents)
            );
    }

    #[test]
    fn case_insensitive() {
        let query = "rUsT";
        let contents = "\
Rust:
safe, fast, productive.
Pick three.
Trust me.";

        assert_eq!(
            vec!["Rust:", "Trust me."],
            search_case_insensitive(query,
            );
    }
}

```

Listing 12-20: Adding a new failing test for the `case_insensitive` search

Note that we've edited the old test's `contents` text `"Duct tape."` using a capital D that should be searching in a case-sensitive manner. Changing that we don't accidentally break the case-sensitive search already implemented. This test should pass now and work on the case-insensitive search.

The new test for the case-*insensitive* search uses the `search_case_insensitive` function we're about to implement to match the line containing `"Rust:"` with a capital

even though both have different casing than the will fail to compile because we haven't yet defined the `contains` function. Feel free to add a skeleton implementation for the `contains` function, similar to the way we did for the `search` function, and test compile and fail.

## Implementing the `search_case_insensitive` function

The `search_case_insensitive` function, shown in Listing 12-21, is the same as the `search` function. The only difference is that we call `line.to_lowercase()` on each `line` so whatever the case of the input string, when we check whether the line contains the query, it will be lowercase.

Filename: src/lib.rs

```
fn search_case_insensitive<'a>(query: &str, contents: &str) -> Vec<'a> {
    let query = query.to_lowercase();
    let mut results = Vec::new();

    for line in contents.lines() {
        if line.to_lowercase().contains(&query) {
            results.push(line);
        }
    }

    results
}
```

Listing 12-21: Defining the `search_case_insensitive` function and the line before comparing them

First, we lowercase the `query` string and store it in a variable with the same name. Calling `to_lowercase` on the query string creates a new `String` because the user's query is `"rust"`, `"RUST"`, `"Rust"`, or `"RuSt"`, and we want to be insensitive to the case.

Note that `query` is now a `String` rather than a `&str`. The `to_lowercase` method creates new data rather than referencing the original data, so `query` is now `"rust"`, as an example: that string slice doesn't exist anymore, so we have to allocate a new `String` for `query` as an argument to the `contains` method because the signature of `contains` is defined to take a `String`.

Next, we add a call to `to_lowercase` on each `li` that contains `query` to lowercase all characters. Now `query` to lowercase, we'll find matches no matter

Let's see if this implementation passes the tests:

```
running 2 tests
test tests::case_insensitive ... ok
test tests::case_sensitive ... ok

test result: ok. 2 passed; 0 failed; 0 ignored
```

Great! They passed. Now, let's call the new `search` function. First, we'll add a configuration option between case-sensitive and case-insensitive search to avoid compiler errors since we aren't initializing this field.

Filename: `src/lib.rs`

```
pub struct Config {
    pub query: String,
    pub filename: String,
    pub case_sensitive: bool,
}
```

Note that we added the `case_sensitive` field to the `run` function to check the `case_sensitive` field whether to call the `search` function or the `search_insensitive` shown in Listing 12-22. Note this still won't compile

Filename: `src/lib.rs`

```

pub fn run(config: Config) -> Result<(), Error> {
    let contents = fs::read_to_string(config.filename)?;

    let results = if config.case_sensitive {
        search(&config.query, &contents)
    } else {
        search_case_insensitive(&config.query, &contents)
    };

    for line in results {
        println!("{}", line);
    }

    Ok(())
}

```

Listing 12-22: Calling either `search` or `search_case_insensitive` in `config.case_sensitive`

Finally, we need to check for the environment variables. The environment variables are in the `env` module in `std`. We'll bring that module into scope with a `use std::env;` and we'll use the `var` method from the `env` module named `CASE_INSENSITIVE`, as shown in Listing 12-23.

Filename: `src/lib.rs`

```

use std::env;

// --snip--

impl Config {
    pub fn new(args: &[String]) -> Result<Config, Error> {
        if args.len() < 3 {
            return Err("not enough arguments");
        }

        let query = args[1].clone();
        let filename = args[2].clone();

        let case_sensitive = env::var("CASE_INSENSITIVE").map_err(|_| Error::new("CASE_INSENSITIVE is not a valid environment variable"))?;

        Ok(Config { query, filename, case_sensitive })
    }
}

```

Listing 12-23: Checking for an environment variable



Here, we create a new variable `case_sensitive` `env::var` function and pass it the name of the variable. The `env::var` method returns a `Result` variant that contains the value of the environment variable is set. It will return the `Err` variant if th

We're using the `is_err` method on the `Result` therefore `unset`, which means it *should* do a case `CASE_INSENSITIVE` environment variable is set t and the program will perform a case-insensitive of the environment variable, just whether it's set rather than using `unwrap`, `expect`, or any of the `Result`.

We pass the value in the `case_sensitive` variab function can read that value and decide whether `search_case_insensitive`, as we implemented

Let's give it a try! First, we'll run our program wit and with the query `to`, which should match any lowercase:

```
$ cargo run to poem.txt
   Compiling minigrep v0.1.0 (file:///proj
   Finished dev [unoptimized + debuginfo]
   Running `target/debug/minigrep to poe
Are you nobody, too?
How dreary to be somebody!
```

Looks like that still works! Now, let's run the prog `1` but with the same query `to`.

If you're using PowerShell, you will need to set th program in two commands rather than one:

```
$ $env:CASE_INSENSITIVE=1
$ cargo run to poem.txt
```

We should get lines that contain "to" that might

```
$ CASE_INSENSITIVE=1 cargo run to poem.txt
    Finished dev [unoptimized + debuginfo]
    Running `target/debug/minigrep to poem.txt`
Are you nobody, too?
How dreary to be somebody!
To tell your name the livelong day
To an admiring bog!
```

Excellent, we also got lines containing “To”! Our case-insensitive searching controlled by an environment variable. We can manage options set using either command line arguments or environment variables.

Some programs allow arguments *and* environment variables for configuration. In those cases, the programs decide the precedence. For another exercise on your own, try writing a program that takes a line of text and prints out the words that start with a given letter, through either a command line argument or an environment variable. Decide whether the command line argument or the environment variable has precedence if the program is run with one set to a value and the other is insensitive.

The `std::env` module contains many more useful environment variables: check out its documentation.

## Writing Error Messages to Standard Error

At the moment, we’re writing all of our output to `stdout`. Most terminals provide two kinds of output: `stdout` for general information and `stderr` for error messages. The `std::env::var` enables users to choose to direct the successful output and error messages to the screen.

The `println!` function is only capable of printing to `stdout`. We can use something else to print to standard error.

## Checking Where Errors Are Written

First, let’s observe how the content printed by `println!` is directed to `stdout` or `stderr`. We’ll do that by redirecting the standard output to a file and intentionally causing an error. We won’t redirect

content sent to standard error will continue to d

Command line programs are expected to send e stream so we can still see error messages on the standard output stream to a file. Our program is about to see that it saves the error message out

The way to demonstrate this behavior is by runn filename, *output.txt*, that we want to redirect the pass any arguments, which should cause an error

```
$ cargo run > output.txt
```

The `>` syntax tells the shell to write the content: instead of the screen. We didn't see the error m the screen, so that means it must have ended up contains:

```
Problem parsing arguments: not enough argu
```

Yup, our error message is being printed to stand error messages like this to be printed to standar run ends up in the file. We'll change that.

## Printing Errors to Standard Error

We'll use the code in Listing 12-24 to change how Because of the refactoring we did earlier in this messages is in one function, `main`. The standard macro that prints to the standard error stream, were calling `println!` to print errors to use `epi`

Filename: src/main.rs

```
fn main() {
    let args: Vec<String> = env::args().collect();

    let config = Config::new(&args).unwrap();
    eprintln!("Problem parsing arguments");
    process::exit(1);
});

if let Err(e) = minigrep::run(config) {
    eprintln!("Application error: {}", e);
    process::exit(1);
}
}
```

Listing 12-24: Writing error messages to standard error using `eprintln!`

After changing `println!` to `eprintln!`, let's run the program without any arguments and redirecting standard error to a file:

```
$ cargo run > output.txt
Problem parsing arguments: not enough arguments
```

Now we see the error onscreen and *output.txt* contains what we expect of command line programs.

Let's run the program again with arguments that redirect standard output to a file, like so:

```
$ cargo run to poem.txt > output.txt
```

We won't see any output to the terminal, and *output.txt* contains:

Filename: output.txt

```
Are you nobody, too?
How dreary to be somebody!
```

This demonstrates that we're now using standard error for error output as appropriate.

## Summary

This chapter recapped some of the major concepts from the previous chapter.

how to perform common I/O operations in Rust files, environment variables, and the `eprintln!` prepared to write command line applications. By chapters, your code will be well organized, store data structures, handle errors nicely, and be well

Next, we'll explore some Rust features that were closures and iterators.

## Functional Language F and Closures

Rust's design has taken inspiration from many e one significant influence is *functional programmi* often includes using functions as values by pass them from other functions, assigning them to va forth.

In this chapter, we won't debate the issue of wha but will instead discuss some features of Rust th languages often referred to as functional.

More specifically, we'll cover:

- *Closures*, a function-like construct you can s
- *Iterators*, a way of processing a series of ele
- How to use these two features to improve
- The performance of these two features (Sp might think!)

Other Rust features, such as pattern matching a other chapters, are influenced by the functional iterators is an important part of writing idiomatic entire chapter to them.

### Closures: Anonymous Functions Their Environment

Rust's closures are anonymous functions you ca arguments to other functions. You can create th

the closure to evaluate it in a different context. It returns values from the scope in which they're called. With these features allow for code reuse and behavior customization.

## Creating an Abstraction of Behavior with Closures

Let's work on an example of a situation in which closures are executed later. Along the way, we'll talk about threads and traits.

Consider this hypothetical situation: we work at a gym and generate custom exercise workout plans. The basic algorithm that generates the workout plan takes the app user's age, body mass index, exercise preference, and intensity number they specify. The actual algorithm is a bit more complex; what's important is that this calculation is expensive and takes a long time to run. We'll simulate this algorithm only when we need to and only calculate it when we need to and only calculate it when we need to wait more than necessary.

We'll simulate calling this hypothetical algorithm with a function `simulated_expensive_calculation` shown in Listing 13-1. The function `calculating slowly...`, wait for two seconds, and then return the value passed in:

Filename: src/main.rs

```
use std::thread;
use std::time::Duration;

fn simulated_expensive_calculation(intensity: i32) {
    println!("calculating slowly...");
    thread::sleep(Duration::from_secs(2));
    intensity
}
```

Listing 13-1: A function to stand in for a hypothetical expensive calculation that takes two seconds to run

Next is the `main` function, which contains the `main` function of this example. This function represents the code that runs the program for a workout plan. Because the interaction with the user is simulated with the use of closures, we'll hardcode values representing user input and print the outputs.

The required inputs are these:

- An intensity number from the user, which indicates whether they want a low intensity workout
- A random number that will generate some

The output will be the recommended workout plan function we'll use:

Filename: src/main.rs

```
fn main() {  
    let simulated_user_specified_value = 1  
    let simulated_random_number = 7;  
  
    generate_workout(  
        simulated_user_specified_value,  
        simulated_random_number  
    );  
}
```

Listing 13-2: A `main` function with hardcoded values for random number generation

We've hardcoded the variable `simulated_user_specified_value` as 1 and the variable `simulated_random_number` as 7 for simulation. In a real app, we'd get the intensity number from the app for the user and generate a random number, as we did in the `generate_workout` function. The `main` function calls a `generate_workout` function with these values.

Now that we have the context, let's get to the algorithm. The `generate_workout` in Listing 13-3 contains the logic that's most concerned with in this example. The rest of the code will be made to this function.

Filename: src/main.rs

```

fn generate_workout(intensity: u32, random_number: u32) {
    if intensity < 25 {
        println!(
            "Today, do {} pushups!",
            simulated_expensive_calculation(intensity)
        );
        println!(
            "Next, do {} situps!",
            simulated_expensive_calculation(intensity)
        );
    } else {
        if random_number == 3 {
            println!("Take a break today!");
        } else {
            println!(
                "Today, run for {} minutes",
                simulated_expensive_calculation(intensity)
            );
        }
    }
}

```

Listing 13-3: The business logic that prints the workout plan. The code calls to the `simulated_expensive_calculation` function.

The code in Listing 13-3 has multiple calls to the `if` block calls `simulated_expensive_calculation` and the `else` doesn't call it at all, and the code inside the `if` block calls `simulated_expensive_calculation` multiple times.

The desired behavior of the `generate_workout` function is that if the user wants a low-intensity workout (indicated by an intensity of 25 or less), the function should call `simulated_expensive_calculation` multiple times. If the user wants a high-intensity workout (a number of 25 or greater), the function should call `simulated_expensive_calculation` once or twice.

Low-intensity workout plans will recommend a run on the complex algorithm we're simulating.

If the user wants a high-intensity workout, there's a 33% chance the random number generated by the app happens to be 3, which means a break and hydration. If not, the user will get a run on the complex algorithm.

This code works the way the business wants it to work. However, the team decides that we need to make some changes to the `simulated_expensive_calculation` function in `src/lib.rs`. When those changes happen, we want to refactor the `simulated_expensive_calculation` function on `src/lib.rs` where we're currently unnecessarily calling the function multiple times.



other calls to that function in the process. That isn't needed, and we still want to call it only once

## Refactoring Using Functions

We could restructure the workout program in m duplicated call to the `simulated_expensive_cal` shown in Listing 13-4:

Filename: src/main.rs

```
fn generate_workout(intensity: u32, random_number: u8) {
    let expensive_result =
        simulated_expensive_calculation(intensity);

    if intensity < 25 {
        println!(
            "Today, do {} pushups!",
            expensive_result
        );
        println!(
            "Next, do {} situps!",
            expensive_result
        );
    } else {
        if random_number == 3 {
            println!("Take a break today!");
        } else {
            println!(
                "Today, run for {} minutes",
                expensive_result
            );
        }
    }
}
```

Listing 13-4: Extracting the calls to `simulated_expensive_calculation` and storing the result in the `expensive_result` variable

This change unifies all the calls to `simulated_expensive_calculation`. The problem of the first `if` block unnecessarily calling `simulated_expensive_calculation` twice is solved. Now we're calling this function and waiting for the result before the inner `if` block that doesn't use the result variable.

We want to define code in one place in our program where we actually need the result. This is a use case for

## Refactoring with Closures to Store Code

Instead of always calling the `simulated_expensive_calculation` `if` blocks, we can define a closure and store the result of the function call, as shown in the whole body of `simulated_expensive_calculation` introducing here:

Filename: src/main.rs

```
let expensive_closure = |num| {  
    println!("calculating slowly...");  
    thread::sleep(Duration::from_secs(2));  
    num  
};
```

Listing 13-5: Defining a closure and storing it in a variable

The closure definition comes after the `=` to assign `expensive_closure`. To define a closure, we start with `|` inside which we specify the parameters to the closure, followed by `{` because of its similarity to closure definitions in other languages. In this case, one parameter named `num`: if we had more than one parameter, we would use commas, like `|param1, param2|`.

After the parameters, we place curly brackets that enclose the closure body—these are optional if the closure body is a single line. After the curly brackets, we need a semicolon to close the closure. The closure returns the value returned from the last line in the closure body (the `num` variable) when it's called, because that line does the calculation. This is similar to function bodies.

Note that this `let` statement means `expensive_closure` is a variable that holds an anonymous function, not the *resulting value* of calling the function. That we're using a closure because we want to define a function and store that code, and call it at a later point; the closure is stored in `expensive_closure`.

With the closure defined, we can change the code in `main` to execute the code and get the resulting value. In the `main` function: we specify the variable name that holds the closure, followed by parentheses containing the argument value. Listing 13-6:

Filename: src/main.rs

```

fn generate_workout(intensity: u32, random_number: u32) {
    let expensive_closure = |num| {
        println!("calculating slowly...");
        thread::sleep(Duration::from_secs(num));
    };

    if intensity < 25 {
        println!(
            "Today, do {} pushups!",
            expensive_closure(intensity)
        );
        println!(
            "Next, do {} situps!",
            expensive_closure(intensity)
        );
    } else {
        if random_number == 3 {
            println!("Take a break today!");
        } else {
            println!(
                "Today, run for {} minutes!",
                expensive_closure(intensity)
            );
        }
    }
}

```

Listing 13-6: Calling the `expensive_closure` we defined

Now the expensive calculation is called in only one place, in the code where we need the results.

However, we’ve reintroduced one of the problems we had with the closure twice in the first `if` block, which will make the user wait twice as long as they need to wait. To avoid creating a variable local to that `if` block to hold the closure, closures provide us with another solution. We’ll see how in a moment. First let’s talk about why there aren’t type annotations for closures involved with closures.

## Closure Type Inference and Annotation

Closures don’t require you to annotate the types of the values like `fn` functions do. Type annotations are optional because they’re part of an explicit interface exposed to you

rigidly is important for ensuring that everyone a function uses and returns. But closures aren't us they're stored in variables and used without nan users of our library.

Closures are usually short and relevant only with any arbitrary scenario. Within these limited cont infer the types of the parameters and the return the types of most variables.

Making programmers annotate the types in these be annoying and largely redundant with the info available.

As with variables, we can add type annotations i clarity at the cost of being more verbose than is types for the closure we defined in Listing 13-5 v Listing 13-7:

Filename: src/main.rs

```
let expensive_closure = |num: u32| -> u32
    println!("calculating slowly...");
    thread::sleep(Duration::from_secs(2));
    num
};
```

Listing 13-7: Adding optional type annotations o types in the closure

With type annotations added, the syntax of closi of functions. The following is a vertical comparis function that adds 1 to its parameter and a closi added some spaces to line up the relevant parts similar to function syntax except for the use of p optional:

```
fn add_one_v1 (x: u32) -> u32 { x + 1 }
let add_one_v2 = |x: u32| -> u32 { x + 1 }
let add_one_v3 = |x|           { x + 1 }
let add_one_v4 = |x|           x + 1
```

The first line shows a function definition, and the closure definition. The third line removes the typ definition, and the fourth line removes the brack

closure body has only one expression. These are the same behavior when they're called.

Closure definitions will have one concrete type in and for their return value. For instance, Listing 13-8 shows a closure that just returns the value it receives as an argument. This is not very useful except for the purposes of this example. I've added type annotations to the definition: if we then try to call the closure as an argument the first time and a `u32` the second time, we'll get a compiler error.

Filename: src/main.rs

```
let example_closure = |x| x;

let s = example_closure(String::from("hello"));
let n = example_closure(5);
```

Listing 13-8: Attempting to call a closure whose type is locked in to the closure in `example_closure`, and then trying to call it with a different type

The compiler gives us this error:

```
error[E0308]: mismatched types
--> src/main.rs
   |
   | let n = example_closure(5);
   |                               ^ expected string, found integral variable
   |
   = note: expected type `std::string::String`
            found type `{integer}`
```

The first time we call `example_closure` with the type of `x` and the return type of the closure locked in to the closure in `example_closure`, and then trying to call it with a different type with the same closure.

## Storing Closures Using Generic Parameters

Let's return to our workout generation app. In Listing 13-9, we use the expensive calculation closure more times than in Listing 13-8. One way to solve this issue is to save the result of the expensive calculation in a variable. In each place we need the result, instead of calling the closure, we can use the variable. However, this method could result in a lot of repetition.

Fortunately, another solution is available to us. \ the closure and the resulting value of calling the closure only if we need the resulting value, and i rest of our code doesn't have to be responsible t may know this pattern as *memoization* or *lazy ev*

To make a struct that holds a closure, we need t because a struct definition needs to know the ty instance has its own unique anonymous type: th same signature, their types are still considered c function parameters that use closures, we use g discussed in Chapter 10.

The `Fn` traits are provided by the standard libra one of the traits: `Fn` , `FnMut` , or `FnOnce` . We'll d traits in the "Capturing the Environment with Clc can use the `Fn` trait.

We add types to the `Fn` trait bound to represen return values the closures must have to match t closure has a parameter of type `u32` and return specify is `Fn(u32) -> u32` .

Listing 13-9 shows the definition of the `Cacher` : optional result value:

Filename: src/main.rs

```
struct Cacher<T>
    where T: Fn(u32) -> u32
{
    calculation: T,
    value: Option<u32>,
}
```

Listing 13-9: Defining a `Cacher` struct that holds optional result in `value`

The `Cacher` struct has a `calculation` field of th on `T` specify that it's a closure by using the `Fn` 1 the `calculation` field must have one `u32` para parentheses after `Fn` ) and must return a `u32` (s

---

Note: Functions can implement all three of th

do doesn't require capturing a value from the function rather than a closure where we need `Fn` trait.

---

The `value` field is of type `Option<u32>`. Before `None`. When code using a `Cacher` asks for the result, it executes the closure at that time and stores the result in the `value` field. Then if the code asks for the result again, the `Cacher` will return the cached value.

The logic around the `value` field we've just described.

Filename: src/main.rs

```
impl<T> Cacher<T>
where T: Fn(u32) -> u32
{
    fn new(calculation: T) -> Cacher<T> {
        Cacher {
            calculation,
            value: None,
        }
    }

    fn value(&mut self, arg: u32) -> u32 {
        match self.value {
            Some(v) => v,
            None => {
                let v = (self.calculation)(arg);
                self.value = Some(v);
                v
            },
        }
    }
}
```

Listing 13-10: The caching logic of `Cacher`

We want `Cacher` to manage the struct fields' values. The code potentially changes the values in these fields.

The `Cacher::new` function takes a generic parameter having the same trait bound as the `Cacher` struct. The `Cacher` instance that holds the closure specifies a `None` value in the `value` field, because we have

When the calling code needs the result of evaluating the closure directly, it will call the `value` method. The closure already has a resulting value in `self.value` in `Some` without executing the closure again.

If `self.value` is `None`, the code calls the closure to calculate the result in `self.value` for future use, and returns it.

Listing 13-11 shows how we can use this `Cacher` to implement `generate_workout` from Listing 13-6:

Filename: `src/main.rs`

```
fn generate_workout(intensity: u32, random_number: u32) {
    let mut expensive_result = Cacher::new(|| {
        println!("calculating slowly...");
        thread::sleep(Duration::from_secs(2));
        num
    });

    if intensity < 25 {
        println!(
            "Today, do {} pushups!",
            expensive_result.value(intensity)
        );
        println!(
            "Next, do {} situps!",
            expensive_result.value(intensity)
        );
    } else {
        if random_number == 3 {
            println!("Take a break today!");
        } else {
            println!(
                "Today, run for {} minutes!",
                expensive_result.value(intensity)
            );
        }
    }
}
```

Listing 13-11: Using `Cacher` in the `generate_workout` function to implement caching logic

Instead of saving the closure in a variable directly, we use the `Cacher` struct that holds the closure. Then, in each place we want the result, we call `value` method on the `Cacher` instance. We can call the `value` method to get the result we want, or not call it at all, and the expensive calculation



Try running this program with the `main` function in the `simulated_user_specified_value` and `s` verify that in all the cases in the various `if` and `calculating slowly...` appears only once and takes care of the logic necessary to ensure we are more than we need to so `generate_workout` call

## Limitations of the `Cacher` Implementation

Caching values is a generally useful behavior that we can use in our code with different closures. However, the current implementation of `Cacher` that would make reusable

The first problem is that a `Cacher` instance assumes that the parameter `arg` to the `value` method. This is not

```
#[test]
fn call_with_different_values() {
    let mut c = Cacher::new(|a| a);

    let v1 = c.value(1);
    let v2 = c.value(2);

    assert_eq!(v2, 2);
}
```

This test creates a new `Cacher` instance with a closure that returns the argument into it. We call the `value` method on this `Cacher` with an `arg` value of 1, and then an `arg` value of 2, and we expect the second call should return 2.

Run this test with the `Cacher` implementation in `src/main.rs` and the test will fail on the `assert_eq!` with this message:

```
thread 'call_with_different_values' panicked at 'assertion failed:
    left: `1`,
    right: `2`, src/main.rs:10:5'
```

The problem is that the first time we called `c.value(1)`, it saved `Some(1)` in `self.value`. Thereafter, no matter what `value` method, it will always return 1.

Try modifying `Cacher` to hold a hash map rather than a single value.

hash map will be the `arg` values that are passed. The value for each key will be the result of calling the closure on that key. If `self.value` directly has a `Some` or a `None` value, return that value. If `arg` in the hash map and return the value if it's `Some`. `Cacher` will call the closure and save the result in the hash map with its `arg` value.

The second problem with the current `Cacher` is that it only caches closures that take one parameter of type `u32`. To make it more flexible, we can cache the results of closures that take a string slice as well. For example, to fix this issue, try introducing more flexibility to the `Cacher` functionality.

## Capturing the Environment with Closures

In the workout generator example, we only use functions. However, closures have an additional feature: they can capture their environment and access variables that were in scope when they're defined.

Listing 13-12 has an example of a closure stored in a variable that captures the `x` variable from the closure's surrounding environment.

Filename: src/main.rs

```
fn main() {  
    let x = 4;  
  
    let equal_to_x = |z| z == x;  
  
    let y = 4;  
  
    assert!(equal_to_x(y));  
}
```

Listing 13-12: Example of a closure that refers to variables in its environment

Here, even though `x` is not one of the parameters of the closure, the closure is allowed to use the `x` variable that's defined in the environment where `equal_to_x` is defined in.

We can't do the same with functions; if we try with a function, it won't compile:

Filename: src/main.rs

```
fn main() {
    let x = 4;

    fn equal_to_x(z: i32) -> bool { z == x }

    let y = 4;

    assert!(equal_to_x(y));
}
```

We get an error:

```
error[E0434]: can't capture dynamic environment in a
...
} closure form instead
--> src/main.rs
4 |         fn equal_to_x(z: i32) -> bool { z == x }
|
```

The compiler even reminds us that this only works for `let` bindings.

When a closure captures a value from its environment, it captures the value for use in the closure body. This use of memory is not ideal in more common cases where we want to capture a reference to the environment. Because functions are never allowed to mutate the environment, defining and using functions will never incur this cost.

Closures can capture values from their environment in three ways a function can take a parameter: by value, by mutable reference, and borrowing immutably. These are explained as follows:

- **FnOnce** consumes the variables it captures from the closure's *environment*. To consume the environment, it takes ownership of these variables and moves them out of the environment. The **Once** part of the name represents that it can take ownership of the same variables more than once.
- **FnMut** can change the environment because it takes mutable references to the variables.
- **Fn** borrows values from the environment.

When you create a closure, Rust infers which trait it implements based on how it uses the values from the environment. All closures that can be called at least once implement **FnOnce**. Closures that don't need to mutate the environment implement **FnMut**, and closures that don't need to borrow values also implement **Fn**. In Listing 13-12, the

immutably (so `equal_to_x` has the `Fn` trait) but needs to read the value in `x`.

If you want to force the closure to take ownership environment, you can use the `move` keyword but this technique is mostly useful when passing a closure so it's owned by the new thread.

We'll have more examples of `move` closures in C concurrency. For now, here's the code from Listing 17.1 added to the closure definition and using vector integers can be copied rather than moved; note

Filename: src/main.rs

```
fn main() {
    let x = vec![1, 2, 3];

    let equal_to_x = move |z| z == x;

    println!("can't use x here: {:?}", x);

    let y = vec![1, 2, 3];

    assert!(equal_to_x(y));
}
```

We receive the following error:

```
error[E0382]: use of moved value: `x`
--> src/main.rs:6:40
   |
4  |         let equal_to_x = move |z| z == x;
   |                                ----- value moved to
5  |                                here
6  |         println!("can't use x here: {:?}",
   |         ^
   |         = note: move occurs because `x` has type `Vec<i32>`
   |         does not
   |         implement the `Copy` trait
```

The `x` value is moved into the closure when the `move` keyword is used. The closure then has ownership of `x` and cannot use `x` anymore in the `println!` statement. Rerun the example.

Most of the time when specifying one of the `Fn`

and the compiler will tell you if you need `FnMut` the closure body.

To illustrate situations where closures that can be used as function parameters, let's move on to our next topic.

## Processing a Series of Items with Iterators

The iterator pattern allows you to perform some operation on each item in a sequence. An iterator is responsible for the logic of iterating over the sequence. When the sequence has finished, the iterator returns a special value. When you use an iterator, you can reimplement that logic yourself.

In Rust, iterators are *lazy*, meaning they have no state and don't consume the iterator to use it up. For example, to create an iterator over the items in the vector `v1` by calling `v1.iter()`. This code by itself doesn't do anything useful.

```
let v1 = vec![1, 2, 3];  
  
let v1_iter = v1.iter();
```

### Listing 13-13: Creating an iterator

Once we've created an iterator, we can use it in a `for` loop. In Chapter 3, we used iterators with `for` loops to iterate over a collection, although we glossed over what the call to `iter` does.

The example in Listing 13-14 separates the creation of the iterator from the `for` loop. The iterator is stored in a variable, and the iteration takes place at that time. When the `for` loop iterates over `v1_iter`, each element in the iterator is used in the loop body, and `println!` prints out each value.

```
let v1 = vec![1, 2, 3];  
  
let v1_iter = v1.iter();  
  
for val in v1_iter {  
    println!("Got: {}", val);  
}
```

Listing 13-14: Using an iterator in a `for` loop

In languages that don't have iterators provided by the compiler, you would likely write this same functionality by starting a variable at 0, incrementing it to index into the vector to get a value, and incrementing it until it reached the total number of items in the vector.

Iterators handle all that logic for you, cutting down on the chance you will potentially mess up. Iterators give you more flexibility by supporting many different kinds of sequences, not just data structures like vectors. Let's examine how iterators do that.

## The `Iterator` Trait and the `next` Method

All iterators implement a trait named `Iterator` in the `std::iter` library. The definition of the trait looks like this:

```
trait Iterator {
    type Item;

    fn next(&mut self) -> Option<Self::Item>;

    // methods with default implementation
}
```

Notice this definition uses some new syntax: `type` for defining an *associated type* with this trait. We'll talk about that in Chapter 19. For now, all you need to know is that the `Iterator` trait requires that you also define an `Item` type, which is used in the return type of the `next` method. In other words, `Item` is the type returned from the iterator.

The `Iterator` trait only requires implementors to define the `next` method, which returns one item of the iterator at a time. When an iteration is over, it returns `None`.

We can call the `next` method on iterators directly. The values returned from repeated calls to `next` on a vector are:

Filename: `src/lib.rs`

```
#[test]
fn iterator_demonstration() {
    let v1 = vec![1, 2, 3];

    let mut v1_iter = v1.iter();

    assert_eq!(v1_iter.next(), Some(&1));
    assert_eq!(v1_iter.next(), Some(&2));
    assert_eq!(v1_iter.next(), Some(&3));
    assert_eq!(v1_iter.next(), None);
}
```

Listing 13-15: Calling the `next` method on an iterator

Note that we needed to make `v1_iter` mutable because the `next` method changes internal state that the iterator uses to keep track of the sequence. In other words, this code *consumes*, or *eats*, items from the sequence. In other words, the `next` method eats up an item from the iterator. We didn't need to make `v1` mutable when we used a `for` loop because the loop took care of making `v1_iter` mutable behind the scenes.

Also note that the values we get from the calls to `next` are references to the values in the vector. The `iter` method produces references. If we want to create an iterator that produces owned values, we can call `into_iter` instead of `iter`. If we want to iterate over mutable references, we can call `iter_mut`.

## Methods that Consume the Iterator

The `Iterator` trait has a number of different methods provided by the standard library; you can find out more about the standard library API documentation for the `Iterator` trait. The methods that call the `next` method in their definition implement the `next` method when implementing the trait.

Methods that call `next` are called *consuming adapters*. One example is the `sum` method, which iterates through the items by repeatedly calling `next` on the iterator. As it iterates through, it adds each item to a total. When iteration is complete, it returns the total. Listing 13-16 has the code for the `sum` method:

Filename: `src/lib.rs`

```
#[test]
fn iterator_sum() {
    let v1 = vec![1, 2, 3];

    let v1_iter = v1.iter();

    let total: i32 = v1_iter.sum();

    assert_eq!(total, 6);
}
```

Listing 13-16: Calling the `sum` method to get the

We aren't allowed to use `v1_iter` after the call of the iterator we call it on.

## Methods that Produce Other Iterators

Other methods defined on the `Iterator` trait, like `map`, to change iterators into different kinds of iterators. Iterator adaptors to perform complex actions in iterators are lazy, you have to call one of the `collect` methods to get the results from calls to iterator adaptors.

Listing 13-17 shows an example of calling the `map` method. It takes a closure to call on each item to produce a new iterator in which each item from the original iterator has been transformed. However, this code produces a warning:

Filename: src/main.rs

```
let v1: Vec<i32> = vec![1, 2, 3];

v1.iter().map(|x| x + 1);
```

Listing 13-17: Calling the iterator adaptor `map` to

The warning we get is this:



```

warning: unused `std::iter::Map` which must be
lazy
and do nothing unless consumed
--> src/main.rs:4:5
   |
4 |         v1.iter().map(|x| x + 1);
   |         ^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^
   |
   = note: #[warn(unused_must_use)] on by default

```

The code in Listing 13-17 doesn't do anything; the iterator is never consumed. The warning reminds us why: iterator adapters don't consume the iterator here.

To fix this and consume the iterator, we'll use the `collect` method from Chapter 12 with `env::args` in Listing 12-1. This method collects the resulting values into a collection data structure.

In Listing 13-18, we collect the results of iterating over the vector `v1` from the call to `map` into a vector. This vector will contain the original vector incremented by 1.

Filename: src/main.rs

```

let v1: Vec<i32> = vec![1, 2, 3];

let v2: Vec<_> = v1.iter().map(|x| x + 1).collect();

assert_eq!(v2, vec![2, 3, 4]);

```

Listing 13-18: Calling the `map` method to create a new iterator and the `collect` method to consume the new iterator and collect the results into a vector.

Because `map` takes a closure, we can specify any logic to apply to each item. This is a great example of how closures are useful while reusing the iteration behavior that the `Iterator` trait provides.

## Using Closures that Capture Their Environment

Now that we've introduced iterators, we can derive a new iterator that captures their environment by using the `filter_map` method. The `filter_map` method on an iterator takes a closure that takes an item and returns a Boolean. If the closure returns `true`, the original item is returned by the iterator produced by `filter_map`. If the closure returns `false`, the item is filtered out.

included in the resulting iterator.

In Listing 13-19, we use `filter` with a closure taken from its environment to iterate over a collection and return only shoes that are the specified size.

Filename: src/lib.rs

```
#[derive(PartialEq, Debug)]
struct Shoe {
    size: u32,
    style: String,
}

fn shoes_in_my_size(shoes: Vec<Shoe>, shoe_size: u32)
    -> Vec<Shoe> {
    shoes.into_iter()
        .filter(|s| s.size == shoe_size)
        .collect()
}

#[test]
fn filters_by_size() {
    let shoes = vec![
        Shoe { size: 10, style: String::from("sneakers") },
        Shoe { size: 13, style: String::from("sneakers") },
        Shoe { size: 10, style: String::from("boots") },
    ];

    let in_my_size = shoes_in_my_size(shoes, 10);

    assert_eq!(
        in_my_size,
        vec![
            Shoe { size: 10, style: String::from("sneakers") },
            Shoe { size: 10, style: String::from("boots") },
        ]
    );
}
```

Listing 13-19: Using the `filter` method with a closure

The `shoes_in_my_size` function takes ownership of the `shoes` vector as a parameter. It returns a vector containing only shoes of the specified size.

In the body of `shoes_in_my_size`, we call `into_iter` to take ownership of the vector. Then we call `filter` to create an iterator that only contains elements for which the closure returns `true`.

The closure captures the `shoe_size` parameter from its environment.

the value with each shoe's size, keeping only shoe sizes that are greater than the value we specified. The `collect` gathers the values returned by the `add_size` function and returns the values returned by the function.

The test shows that when we call `shoes_in_my_size`, the values returned have the same size as the value we specified.

## Creating Our Own Iterators with the `Iterator` Trait

We've shown that you can create an iterator by creating an `iter_mut` on a vector. You can create iterators from the Rust standard library, such as hash map. You can also create your own iterator by implementing the `Iterator` trait on your own type. As mentioned, the only method you're required to implement is the `next` method. Once you've done that, you can use all the methods and implementations provided by the `Iterator` trait.

To demonstrate, let's create an iterator that will iterate over a range of values. We'll create a struct to hold some values. Then we'll implement the `Iterator` trait and use the `next` method to get the next value.

Listing 13-20 has the definition of the `Counter` struct and the `new` function to create instances of `Counter`:

Filename: src/lib.rs

```
struct Counter {
    count: u32,
}

impl Counter {
    fn new() -> Counter {
        Counter { count: 0 }
    }
}
```

Listing 13-20: Defining the `Counter` struct and a `new` function to create instances of `Counter` with an initial value of 0 for `count`

The `Counter` struct has one field named `count` that keeps track of where we are in the process of iterating. The `new` function enforces the behavior of always starting at 0 in the `count` field.

Next, we'll implement the `Iterator` trait for our `Counter` struct. We'll implement the `next` method to specify what we want to return, as shown in Listing 13-21:

Filename: `src/lib.rs`

```
impl Iterator for Counter {
    type Item = u32;

    fn next(&mut self) -> Option<Self::Item> {
        self.count += 1;

        if self.count < 6 {
            Some(self.count)
        } else {
            None
        }
    }
}
```

Listing 13-21: Implementing the `Iterator` trait for `Counter`

We set the associated `Item` type for our iterator to `u32`. The `next` method returns `u32` values. Again, don't worry about associating a type with `next` until Chapter 19.

We want our iterator to add 1 to the current state of `count` before returning. The first time `next` is called, it would return 1 first. If the value of `count` is less than 6, the value is wrapped in `Some`, but if `count` is 6 or higher, it returns `None`.

## Using Our `Counter` Iterator's `next` Method

Once we've implemented the `Iterator` trait, we can use it in a test demonstrating that we can use the iterator to iterate over a range of values by calling the `next` method on it directly, just as we did with a vector in Listing 13-15.

Filename: `src/lib.rs`

```
#[test]
fn calling_next_directly() {
    let mut counter = Counter::new();

    assert_eq!(counter.next(), Some(1));
    assert_eq!(counter.next(), Some(2));
    assert_eq!(counter.next(), Some(3));
    assert_eq!(counter.next(), Some(4));
    assert_eq!(counter.next(), Some(5));
    assert_eq!(counter.next(), None);
}
```

Listing 13-22: Testing the functionality of the `next`

This test creates a new `Counter` instance in the `next` repeatedly, verifying that we have implemented the `next` method correctly: returning the values from 1 to 5 and then `None`.

## Using Other `Iterator` Trait Methods

We implemented the `Iterator` trait by defining the `next` method. We can use any `Iterator` trait method's default implementation from the `std::iter` module, because they all use the `next` method's implementation.

For example, if for some reason we wanted to test the `Counter` trait, we could create an instance of `Counter`, pair them with values produced by `1..4` after skipping the first value, multiply each pair together, filter the results to only those that are divisible by 3, and add all the resulting values. This is what we do in the test in Listing 13-23:

Filename: `src/lib.rs`

```
#[test]
fn using_other_iterator_trait_methods() {
    let sum: u32 = Counter::new().zip(1..4)
        .map(|(a, b)| a * b)
        .filter(|&a| a % 3 == 0)
        .sum();

    assert_eq!(18, sum);
}
```

Listing 13-23: Using a variety of `Iterator` trait methods

Note that `zip` produces only four pairs; the reason is that `zip` returns `None` when either of the iterators returns `None`.

All of these method calls are possible because `works`, and the standard library provides default methods that call `next`.

## Improving Our I/O Project

With this new knowledge about iterators, we can improve Listing 12 by using iterators to make places in the code at how iterators can improve our implementation of the `search` function.

### Removing a `clone` Using an Iterator

In Listing 12-6, we added code that took a slice of an instance of the `Config` struct by indexing into it, allowing the `Config` struct to own those values. This is the implementation of the `Config::new` function.

Filename: `src/lib.rs`

```
impl Config {
    pub fn new(args: &[String]) -> Result<Config> {
        if args.len() < 3 {
            return Err("not enough arguments");
        }

        let query = args[1].clone();
        let filename = args[2].clone();

        let case_sensitive = env::var("CASE_SENSITIVE")
            .ok().unwrap_or(false);

        Ok(Config { query, filename, case_sensitive })
    }
}
```

Listing 13-24: Reproduction of the `Config::new` function

At the time, we said not to worry about the inefficiency of cloning. We can remove them in the future. Well, that time is now.

We needed `clone` here because we have a slice of `String` as a parameter `args`, but the `new` function doesn't take a `Config` instance, we had to clone the values from the slice.

`Config` so the `Config` instance can own its value.

With our new knowledge about iterators, we can take ownership of an iterator as its argument instead of passing an iterator function. This will clarify why we need the `iter` function because the iterator will access the values.

Once `Config::new` takes ownership of the iterator, we can move the `String` into `Config` rather than calling `clone` and making a copy.

## Using the Returned Iterator Directly

Open your I/O project's `src/main.rs` file, which should look like this:

Filename: `src/main.rs`

```
fn main() {
    let args: Vec<String> = env::args().collect();

    let config = Config::new(&args).unwrap_or_else(|| {
        eprintln!("Problem parsing arguments");
        process::exit(1);
    });

    // --snip--
}
```

We'll change the start of the `main` function that is shown in Listing 13-25. This won't compile until we update the code.

Filename: `src/main.rs`

```
fn main() {
    let config = Config::new(env::args()).unwrap_or_else(|| {
        eprintln!("Problem parsing arguments");
        process::exit(1);
    });

    // --snip--
}
```

Listing 13-25: Passing the return value of `env::args()` directly to `Config::new`

The `env::args` function returns an iterator! Rather than calling `collect` into a vector and then passing a slice to `Config::new`, we can pass the iterator directly.

of the iterator returned from `env::args` to `Con`

Next, we need to update the definition of `Config` file, let's change the signature of `Config::new` to won't compile because we need to update the file

Filename: src/lib.rs

```
impl Config {  
    pub fn new(mut args: std::env::Args) -  
        // --snip--
```

Listing 13-26: Updating the signature of `Config::new`

The standard library documentation for the `env::Args` iterator it returns is `std::env::Args`. We'll update the `Config::new` function so the parameter `args` has the type of `&[String]`. Because we're taking ownership of `args` by iterating over it, we can add the `mut` keyword to the parameter to make it mutable.

## Using `Iterator` Trait Methods Instead of `next`

Next, we'll fix the body of `Config::new`. The standard library mentions that `std::env::Args` implements the `Iterator` trait, so we can call the `next` method on it! Listing 13-27 updates the `next` method:

Filename: src/lib.rs



```

impl Config {
    pub fn new(mut args: std::env::Args) -
        args.next();

    let query = match args.next() {
        Some(arg) => arg,
        None => return Err("Didn't get
    };

    let filename = match args.next() {
        Some(arg) => arg,
        None => return Err("Didn't get
    };

    let case_sensitive = env::var("CASE_

    Ok(Config { query, filename, case_
    }
}

```

Listing 13-27: Changing the body of `Config::new`

Remember that the first value in the return value is the error value. We want to ignore that and get to the actual value. We do nothing with the return value. Second, we call `args.next()` to get the next argument. We put it in the `query` field of `Config`. If `args.next()` returns `None`, it means not enough arguments were provided. We return early with an `Err` value. We do the same

## Making Code Clearer with Iterator Adapter

We can also take advantage of iterators in the `search` function, which is reproduced here in Listing 13-28 as it will be used in the next chapter.

Filename: `src/lib.rs`

```

pub fn search<'a>(query: &str, contents: &Vec<'a>) -
    let mut results = Vec::new();

    for line in contents.lines() {
        if line.contains(query) {
            results.push(line);
        }
    }

    results
}

```

Listing 13-28: The implementation of the `search`

We can write this code in a more concise way using `filter` and `collect`, so also lets us avoid having a mutable intermediate vector. The `filter` method is an iterator that filters the elements of an iterator. The `collect` method is an iterator that collects the elements of an iterator into a new vector. The `filter` method is an iterator that filters the elements of an iterator. The `collect` method is an iterator that collects the elements of an iterator into a new vector. Listing 13-29 shows this.

Filename: `src/lib.rs`

```
pub fn search<'a>(query: &str, contents: &Vec<'a>)\n    .lines()\n    .filter(|line| line.contains(query))\n    .collect()\n}
```

Listing 13-29: Using iterator adaptor methods in a function

Recall that the purpose of the `search` function is to find all lines that contain the `query`. Similar to the `filter` example, we can use the `filter` adaptor to keep only the lines that `line.contains(query)`. We then collect the matching lines into another vector. Feel free to make the same change to use iterators in the `search_case_insensitive` function as well.

The next logical question is which style you should use: the original implementation in Listing 13-28 or the implementation in Listing 13-29. Most Rust programmers prefer to use the `filter` and `collect` style at first, but once you get a feel for it, they do, iterators can be easier to understand. In terms of looping and building new vectors, the code for the `filter` and `collect` style is more concise. This abstracts away some of the common concepts that are unique to this code, such as the `line` variable and the `collect` method. The `filter` and `collect` style is more concise.

But are the two implementations truly equivalent? What if that the more low-level loop will be faster. Let's find out.

## Comparing Performance: Loops vs. Iterators

To determine whether to use loops or iterators,

our `search` functions is faster: the version with  
with iterators.

We ran a benchmark by loading the entire contents of *Holmes* by Sir Arthur Conan Doyle into a `String` and printing its contents. Here are the results of the benchmark using the `for` loop and the version using iterators:

```
test bench_search_for ... bench: 19,620,
test bench_search_iter ... bench: 19,234,
```

The iterator version was slightly faster! We won't win because the point is not to prove that the two versions are equivalent in a general sense of how these two implementations compare.

For a more comprehensive benchmark, you should use various sizes as the **contents**, different words as the **query**, and all kinds of other variations. The point is that high-level abstraction, get compiled down to roughly the same as lower-level code yourself. Iterators are one of Rust's features we mean using the abstraction imposes no additional overhead, analogous to how Bjarne Stroustrup, the original C++ creator, defines *zero-overhead* in "Foundations of C++" (2017).

In general, C++ implementations obey the zero-overhead principle: you don't use, you don't pay for. And further: What's better, you write it or let the compiler code any better.

As another example, the following code is taken from the `linear` crate. This algorithm uses the linear prediction mathematics based on a linear function of the previous sample to do some math on three variables in scope: a vector of `coefficients`, and an amount by which to shift the variables within this example but not given to the caller. This variable doesn't have much meaning outside of its context, but it's an example of how Rust translates high-level ideas into low-level code.

```

let buffer: &mut [i32];
let coefficients: [i64; 12];
let qlp_shift: i16;

for i in 12..buffer.len() {
    let prediction = coefficients.iter()
        .zip(&buffer[..i])
        .map(|(&c, &b)| &c * &b)
        .sum::<i64>();

    let delta = buffer[i];
    buffer[i] = prediction as i32 + delta;
}

```

To calculate the value of `prediction`, this code iterates over the values in `coefficients` and uses the `zip` method to pair each coefficient with its corresponding value in `buffer`. Then, for each pair, it calculates the product of the coefficient and the value, and sums all the results, and shift the bits in the sum.

Calculations in applications like audio decoders are often done at a high level. Here, we're creating an iterator, using two values, and then using the `zip` method to pair them. What assembly code would this Rust code compile down to? The same assembly you'd write if you were doing the calculations manually. There are 12 iterations, so it "unrolls" the loop. `Unroll` removes the overhead of the loop controlling code for each iteration of the loop.

All of the coefficients get stored in registers, which is fast. There are no bounds checks on the array access. The optimizations that Rust is able to apply make the code run faster. Now that you know this, you can use iterators and closures. They seem like it's higher level but don't impose a performance cost for doing so.

## Summary

Closures and iterators are Rust features inspired by other language ideas. They contribute to Rust's capability to achieve high performance at low-level performance. The implementations of these features ensure that runtime performance is not affected. This is possible because of zero-cost abstractions.

Now that we've improved the expressiveness of Rust, we can use more features of `cargo` that will help us share the code.

# More About Cargo and

So far we've used only the most basic features of Cargo, but it can do a lot more. In this chapter, we'll explore some advanced features to show you how to do the following:

- Customize your build through release profiles
- Publish libraries on [crates.io](https://crates.io)
- Organize large projects with workspaces
- Install binaries from [crates.io](https://crates.io)
- Extend Cargo using custom commands

Cargo can do even more than what we cover in this chapter. For more information about all its features, see [its documentation](#).

## Customizing Builds with Release Profiles

In Rust, *release profiles* are predefined and customizable configurations that allow a programmer to have different settings for compiling code. Each profile is configured in the `Cargo.toml` file.

Cargo has two main profiles: the `dev` profile Cargo uses when you run `cargo build` and the `release` profile Cargo uses when you run `cargo build --release`. Both profiles are defined with good defaults for development and release builds.

These profile names might be familiar from the Rust documentation.

```
$ cargo build
   Finished dev [unoptimized + debuginfo] target(s) in 0.1s
$ cargo build --release
   Finished release [optimized] target(s) in 0.1s
```

The `dev` and `release` shown in this build output correspond to different profiles.

Cargo has default settings for each of the profiles. You can customize these settings in the `[profile.*]` sections in the project's `Cargo.toml` file. For any profile you want to customize, you can specify the profile name and its default settings. For example, here are the default settings for the `dev` and `release` profiles:

Filename: Cargo.toml

```
[profile.dev]
opt-level = 0

[profile.release]
opt-level = 3
```

The `opt-level` setting controls the number of compiler optimizations applied to the code, with a range of 0 to 3. Applying more optimizations is useful if you're in development and compiling your code often, as it allows you to see the results of your code changes even if the resulting code runs slower. That is the default for the `dev` profile. When you're ready to release your code, you can compile in release mode, which will compile your program many times, so release mode trades speed for faster compilation. That is why the default `opt-level` for the `release` profile is 3.

You can override any default setting by adding a `[profile.<name>]` section to your `Cargo.toml`. For example, if we want to use optimization level 1 for the `dev` profile, we can add these two lines to our project's `Cargo.toml` file:

Filename: Cargo.toml

```
[profile.dev]
opt-level = 1
```

This code overrides the default setting of 0. Now, when you run `cargo build`, it will use the defaults for the `dev` profile plus our custom setting. Because we set `opt-level` to 1, Cargo will apply fewer optimizations than the default, but not as many as in a release build.

For the full list of configuration options and default values, see the [Cargo configuration documentation](#).

## Publishing a Crate to Crates.io

We've used packages from [crates.io](#) as dependencies in our projects. To share your code with other people by publishing your crate, you need to register your crate with the crates.io registry. The crates.io registry distributes the source code of your crate to the crates.io website, which is open source.

Rust and Cargo have features that help make your code easy to use and to find in the first place. We'll cover these features in the next chapter.

next and then explain how to publish a package

## Making Useful Documentation Comments

Accurately documenting your packages will help people use them, so it's worth investing the time to write. We've discussed how to comment Rust code using two particular kinds of comment for documentation, *documentation comment*, that will generate HTML documentation from the contents of documentation comments for programmers interested in knowing how to use a crate is *implemented*.

Documentation comments use three slashes, `///`, and Markdown notation for formatting the text. Place them before the item they're documenting. Listing 14-1 shows an `add_one` function in a crate named `my_crate`.

Filename: `src/lib.rs`

```
/// Adds one to the number given.
///
/// # Examples
///
/// ```
/// let five = 5;
///
/// assert_eq!(6, my_crate::add_one(5));
/// ```
pub fn add_one(x: i32) -> i32 {
    x + 1
}
```

Listing 14-1: A documentation comment for a function

Here, we give a description of what the `add_one` function does, the heading `Examples`, and then provide code that demonstrates the `add_one` function. We can generate the HTML documentation from the documentation comment by running `cargo doc`, a tool distributed with Rust that puts the generated documentation in the `target/doc` directory.

For convenience, running `cargo doc --open` will open the crate's documentation (as well as the documentation of its dependencies) and open the result in a web browser.

function and you'll see how the text in the document is shown in Figure 14-1:

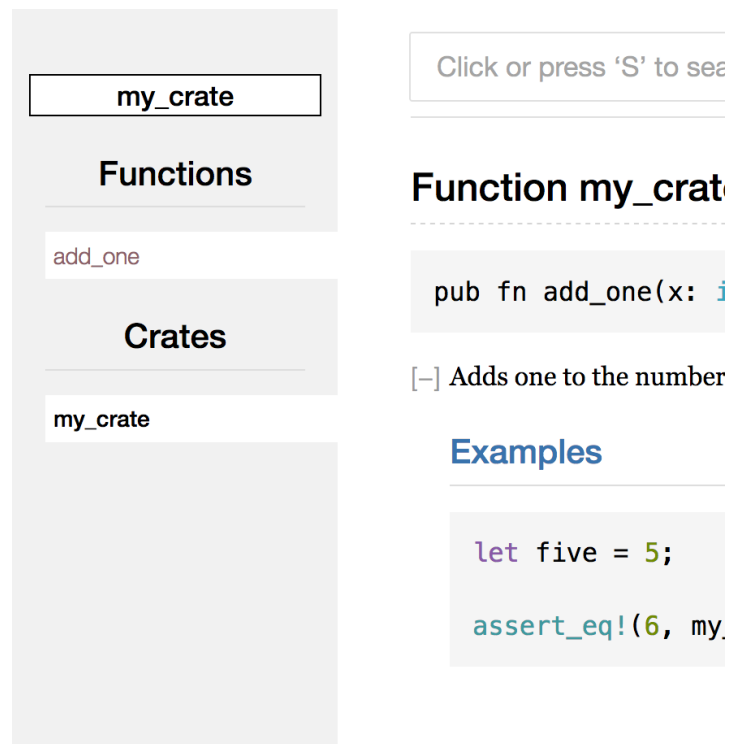


Figure 14-1: HTML documentation for the `add_one` function

## Commonly Used Sections

We used the `# Examples` Markdown heading in the documentation to create the "Examples" section in the HTML with the title "Examples." Here are some sections that are commonly used in their documentation:

- **Panics:** The scenarios in which the function will panic. Callers of the function who don't want their program to panic should not call the function in these situations.
- **Errors:** If the function returns a `Result`, documentation should explain what errors might occur and what conditions might cause them. This is helpful to callers so they can write code that handles errors in different ways.
- **Safety:** If the function is `unsafe` to call (we'll discuss this later), there should be a section explaining why it's unsafe and what invariants that the function expects callers to maintain.

Most documentation comments don't need all of these sections. Use the checklist to remind you of the aspects of your code that callers might be interested in knowing about.



## Documentation Comments as Tests

Adding example code blocks in your documentation shows how to use your library, and doing so has an advantage: `cargo test` will run the code examples in your documentation as tests. But nothing is wrong with documentation with examples. But nothing is wrong with documentation with examples because the code has changed since the documentation was written. `cargo test` with the documentation for the `add_one` function will see a section in the test results like this:

```
Doc-tests my_crate

running 1 test
test src/lib.rs - add_one (line 5) ... ok

test result: ok. 1 passed; 0 failed; 0 ignored
```

Now if we change either the function or the example code, the test will fail. If the example panics and run `cargo test` again, we'll see a failure. The example and the code are out of sync with each other.

## Commenting Contained Items

Another style of doc comment, `//!`, adds documentation to the items it precedes rather than adding documentation to the module. We typically use these doc comments (the `//!` convention) or inside a module to document the items it contains.

For example, if we want to add documentation to the `my_crate` crate that contains the `add_one` function, we can add comments that start with `//!` to the beginning of the file (Listing 14-2):

Filename: src/lib.rs

```
//! # My Crate
//!
//! `my_crate` is a collection of utilities
//! for performing calculations more convenient.

/// Adds one to the number given.
// --snip--
```

Listing 14-2: Documentation for the `my_crate` crate

Notice there isn't any code after the last line that

started the comments with `///!` instead of `///`, contains this comment rather than an item that the item that contains this comment is the `src/lib` comments describe the entire crate.

When we run `cargo doc --open`, these comments generate the documentation for `my_crate` above the list in Figure 14-2:

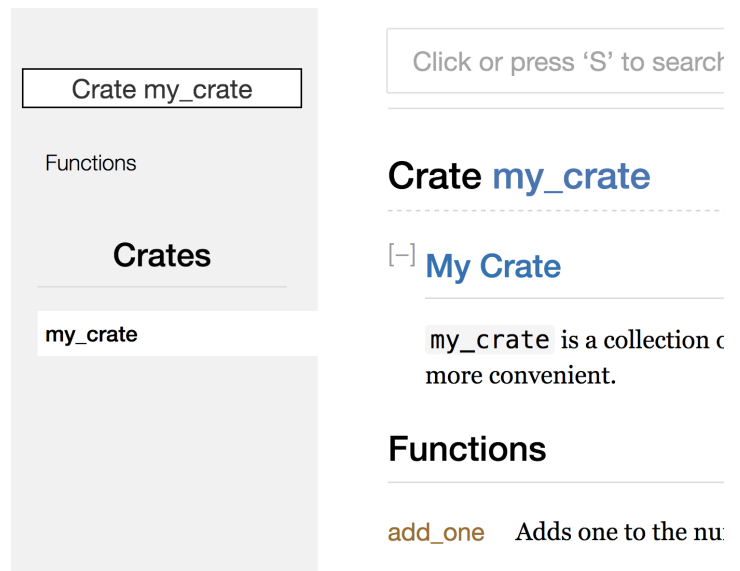


Figure 14-2: Rendered documentation for `my_cr` describing the crate as a whole

Documentation comments within items are useful especially. Use them to explain the overall purpose of the crate so that users understand the crate's organization.

## Exporting a Convenient Public API with

In Chapter 7, we covered how to organize our code using the `pub` keyword, how to make items public using the `pub` keyword, how to make items public using the `pub` keyword into a scope with the `use` keyword. However, the `use` keyword while you're developing a crate might not be very convenient. You might want to organize your structs in a hierarchy so that people who want to use a type you've defined don't have to go through a lot of trouble finding out that type exists. They might also want to use the `use` keyword to make the code more concise. For example, you can use the `use` keyword to make the code more concise:

```
use my_crate::some_module::another_module::UsefulType; .
```

The structure of your public API is a major consideration when designing a crate.

People who use your crate are less familiar with have difficulty finding the pieces they want to use hierarchy.

The good news is that if the structure *isn't* convenient library, you don't have to rearrange your internal export items to make a public structure that's different using `pub use`. Re-exporting takes a public item to another location, as if it were defined in the other

For example, say we made a library named `art` and this library has two modules: a `kinds` module containing `PrimaryColor` and `SecondaryColor` and a `utils` module named `mix`, as shown in Listing 14-3:

Filename: `src/lib.rs`

```
///! # Art
///!
///! A library for modeling artistic concepts

pub mod kinds {
    /// The primary colors according to traditional art
    pub enum PrimaryColor {
        Red,
        Yellow,
        Blue,
    }

    /// The secondary colors according to traditional art
    pub enum SecondaryColor {
        Orange,
        Green,
        Purple,
    }
}

pub mod utils {
    use kinds::*;

    /// Combines two primary colors in equal parts to create
    /// a secondary color.
    pub fn mix(c1: PrimaryColor, c2: PrimaryColor) -> SecondaryColor {
        // --snip--
    }
}
```

Listing 14-3: An `art` library with items organized by module

Figure 14-3 shows what the front page of the documentation for the `art` crate by `cargo doc` would look like:

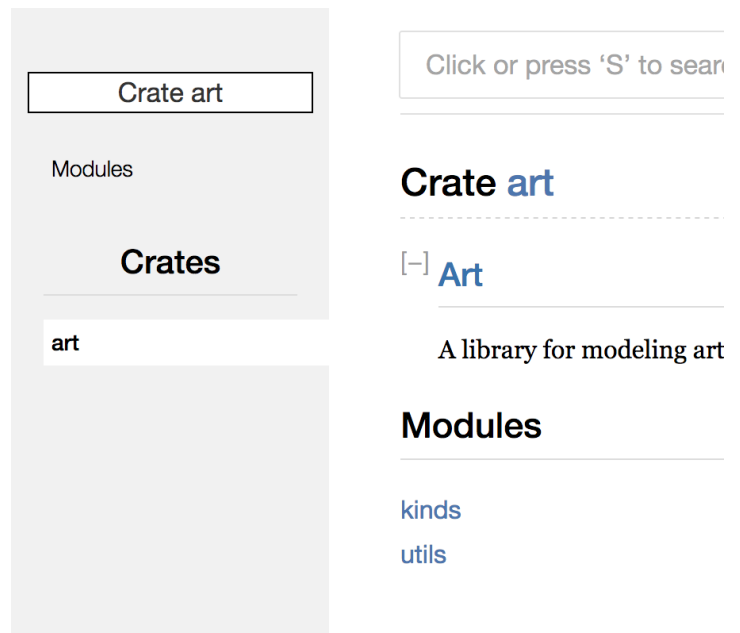


Figure 14-3: Front page of the documentation for the `art` crate

Note that the `PrimaryColor` and `SecondaryColor` are not on the `art` crate's page, nor is the `mix` function. We have to click

Another crate that depends on this library would use the items from `art`, specifying the module structure. Listing 14-4 shows an example of a crate that uses the `art` crate:

Filename: `src/main.rs`

```
extern crate art;

use art::kinds::PrimaryColor;
use art::utils::mix;

fn main() {
    let red = PrimaryColor::Red;
    let yellow = PrimaryColor::Yellow;
    mix(red, yellow);
}
```

Listing 14-4: A crate using the `art` crate's items

The author of the code in Listing 14-4, which uses the `art` crate, uses `PrimaryColor` from the `kinds` module and

module structure of the `art` crate is more relevant to developers using the `art` crate than to developers using the `art` crate. The `art` crate organizes parts of the crate into the `kinds` module, which may contain any useful information for someone trying to use the crate. Instead, the `art` crate's module structure forces developers to figure out where to look, and because developers must specify the module name, it is not clear where to look.

To remove the internal organization from the public API, we can add `pub use` statements to the `lib.rs` file, as shown in Listing 14-5:

Filename: `src/lib.rs`

```
///! # Art
///!
///! A library for modeling artistic concepts

pub use kinds::PrimaryColor;
pub use kinds::SecondaryColor;
pub use utils::mix;

pub mod kinds {
    // --snip--
}

pub mod utils {
    // --snip--
}
```

Listing 14-5: Adding `pub use` statements to re-export the crate's public API

The API documentation that `cargo doc` generates for the crate now shows the `SecondaryColor` types and the `mix` function on the front page, as shown in Figure 14-5.

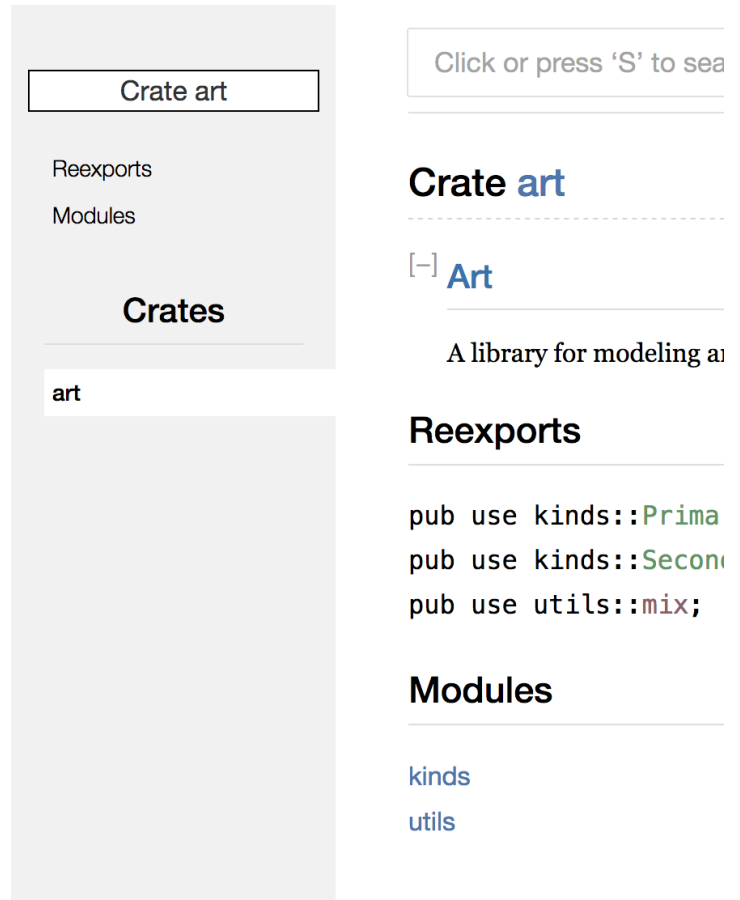


Figure 14-4: The front page of the documentation

The `art` crate users can still see and use the interface demonstrated in Listing 14-4, or they can use the interface in Listing 14-5, as shown in Listing 14-6:

Filename: `src/main.rs`

```
extern crate art;

use art::PrimaryColor;
use art::mix;

fn main() {
    // --snip--
}
```

Listing 14-6: A program using the re-exported interface

In cases where there are many nested modules, the `pub use` level with `pub use` can make a significant difference in how you use the crate.

Creating a useful public API structure is more of an iterative process to find the API that works best for your use case.

flexibility in how you structure your crate internal structure from what you present to your users. If you've installed to see if their internal structure

## Setting Up a Crates.io Account

Before you can publish any crates, you need to create an API token. To do so, visit the home page at [crates.io](https://crates.io) and create an account. (The GitHub account is currently a requirement, but there are other ways of creating an account in the future.) Once you have an account, visit your account settings at <https://crates.io/me/> and retrieve your API key. Then run the `cargo login` command with your API key, like this:

```
$ cargo login abcdefghijklmnopqrstuvwxyz01
```

This command will inform Cargo of your API token. The token is stored in `~/.cargo/credentials`. Note that this token is a *secret*: do not share it with anyone for any reason, you should keep it secret on [crates.io](https://crates.io).

## Adding Metadata to a New Crate

Now that you have an account, let's say you have a crate you want to publish, you'll need to add some metadata to the `[package]` section of the crate's `Cargo.toml` file.

Your crate will need a unique name. While you're creating a crate, you can name a crate whatever you'd like. However, crates are published on a first-come, first-served basis. Once a crate name is taken, you cannot create a crate with that name. Search for the name you want to use to see whether it has been used. If it hasn't, edit the `[package]` section to use the name for publishing, like this:

Filename: Cargo.toml

```
[package]
name = "guessing_game"
```

Even if you've chosen a unique name, when you run `cargo publish` at this point, you'll get a warning and then

```
$ cargo publish
    Updating registry `https://github.com/
warning: manifest has no description, license,
documentation,
homepage or repository.
--snip--
error: api errors: missing or empty metadata
```

The reason is that you're missing some crucial information that is required so people will know what your crate can use it. To rectify this error, you need to include a `description` field in your `Cargo.toml` file.

Add a description that is just a sentence or two, which will appear in search results. For the `license` field, you need to specify a valid license identifier. You can find a list of valid identifiers on the [Linux Foundation's Software Package Data Exchange \(SPDX\) License List](https://spdx.org/licenses/). For example, to specify the MIT License, add the `MIT` identifier:

Filename: Cargo.toml

```
[package]
name = "guessing_game"
license = "MIT"
```

If you want to use a license that doesn't appear in the list, you can include the full text of that license in a file, and then specify the name of that file in the `license-file` field.

Guidance on which license is appropriate for your project can be found in the [Rust License Guide](https://www.rust-lang.org/en-US/about/licenses). Many people in the Rust community license their projects by using a dual license of `MIT OR Apache-2.0`. You can also specify multiple license identifiers separated by `OR` to specify multiple licenses for your project.

With a unique name, the version, the author details, a description, and a license, your project that is ready to publish might look like this:

Filename: Cargo.toml



```
[package]
name = "guessing_game"
version = "0.1.0"
authors = ["Your Name <you@example.com>"]
description = "A fun game where you guess
chosen."
license = "MIT OR Apache-2.0"

[dependencies]
```

[Cargo's documentation](#) describes other metadata you can discover and use your crate more easily.

## Publishing to Crates.io

Now that you've created an account, saved your crate, and specified the required metadata, you'll upload a specific version to [crates.io](#) for others

Be careful when publishing a crate because a package can never be overwritten, and the code cannot be deleted. It acts as a permanent archive of code so that builds from [crates.io](#) will continue to work. Allow for the possibility of fulfilling that goal impossible. However, there is a way to publish versions you can publish.

Run the `cargo publish` command again. It should

```
$ cargo publish
Updating registry `https://github.com/rust-lang/crates.io`
Packaging guessing_game v0.1.0 (file:///projects/guessing_game)
Verifying guessing_game v0.1.0 (file:///projects/guessing_game)
Compiling guessing_game v0.1.0
    Finished dev [unoptimized + debuginfo] target(s) in 0.1s
Uploading guessing_game v0.1.0 (file:///projects/guessing_game)
```

Congratulations! You've now shared your code with the world. You can easily add your crate as a dependency of the

## Publishing a New Version of an Existing Crate

When you've made changes to your crate and are ready to release, change the `version` value specified in your `Cargo.toml`.

[Semantic Versioning rules](#) to decide what an app based on the kinds of changes you've made. The new version.

## Removing Versions from Crates.io with

Although you can't remove previous versions of projects from adding them as a new dependency is broken for one reason or another. In such situations, you can yank a version.

Yanking a version prevents new projects from starting to depend on that version while allowing all existing projects that depend on that version. Essentially, a yank means that the version will not break, and any future *Cargo.lock* files generated will not break.

To yank a version of a crate, run `cargo yank` and specify the version to yank:

```
$ cargo yank --vers 1.0.1
```

By adding `--undo` to the command, you can also undo the yank and start depending on a version again:

```
$ cargo yank --vers 1.0.1 --undo
```

A yank *does not* delete any code. For example, if you accidentally uploaded secrets. If that happens, you should delete them immediately.

## Cargo Workspaces

In Chapter 12, we built a package that included a library and a binary. As your project develops, you might find that the library and binary are doing different things and you want to split up your package further in some way. In this situation, Cargo offers a feature called *workspaces* that allows you to develop related packages that are developed in tandem.

### Creating a Workspace

A *workspace* is a set of packages that share the same dependencies. Let's make a project using a workspace—we'll use the structure of the workspace. There are many ways to do this, but we're going to show one common way. We'll have a binary crate and two libraries. The binary, which will provide the functionality, will use the two libraries. One library will provide an `add_one` function, and the other will provide an `add_two` function. These three crates will be managed by Cargo. We'll start by creating a new directory for the workspace:

```
$ mkdir add
$ cd add
```

Next, in the *add* directory, we create the *Cargo.toml* file for the workspace. This file won't have a `[package]` section like other *Cargo.toml* files. Instead, it will start with a `[workspace]` section to tell Cargo that we're creating a workspace. We'll use the `members` field to add members to the workspace by specifying their relative paths. In this case, that path is *adder*:

Filename: Cargo.toml

```
[workspace]

members = [
    "adder",
]
```

Next, we'll create the *adder* binary crate by running `cargo new adder` in the *add* directory:

```
$ cargo new adder
    Created binary (application) `adder`
```

At this point, we can build the workspace by running `cargo build` in the *add* directory. The directory structure should look like this:

```
├── Cargo.lock
├── Cargo.toml
├── adder
│   ├── Cargo.toml
│   └── src
│       └── main.rs
└── target
```

The workspace has one *target* directory at the top level where binaries will be placed into; the *adder* crate doesn't have its own *target* directory. To build the workspace, we'll run `cargo build` from inside the *add* directory:

end up in *add/target* rather than *add/adder/target* directory in a workspace like this because the crates depend on each other. If each crate had its own *target* directory, we would have to recompile each of the other crates in the workspace every time we rebuild *add*. By sharing one *target* directory, we can avoid this rebuilding.

## Creating the Second Crate in the Workspace

Next, let's create another member crate in the workspace. We'll add the *add-one* crate to the top-level *Cargo.toml* to specify the *add-one* package.

Filename: Cargo.toml

```
[workspace]

members = [
    "adder",
    "add-one",
]
```

Then generate a new library crate named *add-one* using the `cargo new` command:

```
$ cargo new add-one --lib
    Created library `add-one` project
```

Your *add* directory should now have these directories and files:

```
├── Cargo.lock
├── Cargo.toml
├── add-one
│   ├── Cargo.toml
│   └── src
│       └── lib.rs
├── adder
│   ├── Cargo.toml
│   └── src
│       └── main.rs
└── target
```

In the *add-one/src/lib.rs* file, let's add an *add\_one* function that takes an integer and returns the integer plus one.

Filename: add-one/src/lib.rs

```
pub fn add_one(x: i32) -> i32 {
    x + 1
}
```

Now that we have a library crate in the workspace, the `adder` crate depends on the library crate `add-one`. First, we add a dependency on `add-one` to `adder/Cargo.toml`.

Filename: `adder/Cargo.toml`

```
[dependencies]

add-one = { path = "../add-one" }
```

Cargo doesn't assume that crates in a workspace need to be explicit about the dependency relationship.

Next, let's use the `add_one` function from the `add-one` crate. Open the `adder/src/main.rs` file and add an `extern crate` to bring the `add-one` library crate into scope. Then call the `add_one` function, as in Listing 14-7:

Filename: `adder/src/main.rs`

```
extern crate add_one;

fn main() {
    let num = 10;
    println!("Hello, world! {} plus one is",
        add_one::add_one(num));
}
```

Listing 14-7: Using the `add-one` library crate from the workspace

Let's build the workspace by running `cargo build`:

```
$ cargo build
   Compiling add-one v0.1.0 (file:///project/add-one)
   Compiling adder v0.1.0 (file:///project/adder)
    Finished dev [unoptimized + debuginfo] target(s) in 0.1s
```

To run the binary crate from the `adder` directory, we use the `cargo run` command to run the workspace we want to use by using the `-p` flag:

```
$ cargo run -p adder
    Finished dev [unoptimized + debuginfo]
    Running `target/debug/adder`
Hello, world! 10 plus one is 11!
```

This runs the code in *adder/src/main.rs*, which de

## Depending on an External Crate in a Workspace

Notice that the workspace has only one *Cargo.lock* workspace rather than having a *Cargo.lock* in each crate. This means that all crates are using the same version of all dependencies. In the *adder/Cargo.toml* and *add-one/Cargo.toml* files, we specify one version of `rand` and record that in the one *Cargo.lock* file. This means that all crates in the workspace use the same dependencies means that they will be compatible with each other. Let's add the `rand` section in the *add-one/Cargo.toml* file to be able to use the `rand` crate:

Filename: add-one/Cargo.toml

```
[dependencies]
rand = "0.3.14"
```

We can now add `extern crate rand;` to the *ad* whole workspace by running `cargo build` in the `ad` directory to compile the `rand` crate:

```
$ cargo build
  Updating registry `https://github.com/
Downloading rand v0.3.14
--snip--
Compiling rand v0.3.14
Compiling add-one v0.1.0 (file:///proje
Compiling adder v0.1.0 (file:///project
Finished dev [unoptimized + debuginfo]
```

The top-level *Cargo.lock* now contains information on `rand`. However, even though `rand` is used to use it in other crates in the workspace unless we well. For example, if we add `extern crate rand` `adder` crate, we'll get an error:

```
$ cargo build
   Compiling adder v0.1.0 (file:///project)
error: use of unstable library feature 'rand'
(see
issue #27703)
--> adder/src/main.rs:1:1
   |
1 | extern crate rand;
```

To fix this, edit the *Cargo.toml* file for the `adder` dependency for that crate as well. Building the set of dependencies for `adder` in *Cargo.lock*, but no downloaded. Cargo has ensured that every crate will be using the same version. Using the workspace saves space because we won't have multiple crates in the workspace will be compatible with

## Adding a Test to a Workspace

For another enhancement, let's add a test of the the `add_one` crate:

Filename: add-one/src/lib.rs

```
pub fn add_one(x: i32) -> i32 {
    x + 1
}

#[cfg(test)]
mod tests {
    use super::*;

    #[test]
    fn it_works() {
        assert_eq!(3, add_one(2));
    }
}
```

Now run `cargo test` in the top-level *add* directory

```

$ cargo test
  Compiling add-one v0.1.0 (file:///project/crates/add-one)
  Compiling adder v0.1.0 (file:///project/crates/adder)
  Finished dev [unoptimized + debuginfo] target(s) in 0.1s
  Running target/debug/deps/add_one-f02...

running 1 test
test tests::it_works ... ok

test result: ok. 1 passed; 0 failed; 0 ignored; 0 measured; 0 filtered out
     Running target/debug/deps/adder-f88a1...

running 0 tests

test result: ok. 0 passed; 0 failed; 0 ignored; 0 measured; 0 filtered out
     Doc-tests add-one

running 0 tests

test result: ok. 0 passed; 0 failed; 0 ignored; 0 measured; 0 filtered out

```

The first section of the output shows that the `it_works` test passed. The next section shows that zero tests were run for the `adder` crate, and then the last section shows zero documentation tests for the `add-one` crate. Running `cargo test` in a workspace structure runs tests for all the crates in the workspace.

We can also run tests for one particular crate in a workspace directory by using the `-p` flag and specifying the crate name.

```

$ cargo test -p add-one
  Finished dev [unoptimized + debuginfo] target(s) in 0.1s
  Running target/debug/deps/add_one-b32...

running 1 test
test tests::it_works ... ok

test result: ok. 1 passed; 0 failed; 0 ignored; 0 measured; 0 filtered out
     Doc-tests add-one

running 0 tests

test result: ok. 0 passed; 0 failed; 0 ignored; 0 measured; 0 filtered out

```

This output shows `cargo test` only ran the tests for the `add-one` crate.



If you publish the crates in the workspace to *http* workspace will need to be published separately. not have an `--all` flag or a `-p` flag, so you must run `cargo publish` on each crate in the workspace.

For additional practice, add an `add-two` crate to the `add-one` crate!

As your project grows, consider using a workspace instead of individual components than one big blob of code. A workspace can make coordination between them at the same time.

## Installing Binaries from Crates

The `cargo install` command allows you to install binaries. It isn't intended to replace system packages; it's more for developers to install tools that others have shared. It will only install packages that have binary targets. A binary target is created if the crate has a `src/main.rs` file or a `bin/` directory, as opposed to a library target that isn't runnable or is used within other programs. Usually, crates have information in their `Cargo.toml` about whether a crate is a library, has a binary target, or both.

All binaries installed with `cargo install` are stored in a `bin/` folder. If you installed Rust using `rustup.rs` and did not specify a directory, this directory will be `$HOME/.cargo/bin`. Ensure that your `PATH` is able to run programs you've installed with `cargo install`.

For example, in Chapter 12 we mentioned that the `grep` tool called `ripgrep` for searching files. If you want to install it, run the following:

```
$ cargo install ripgrep
Updating registry `https://github.com/rust-lang/cargo`
Downloading ripgrep v0.3.2
--snip--
Compiling ripgrep v0.3.2
Finished release [optimized + debuginfo] target(s) in 0.5s
Installing ~/.cargo/bin/rg
```

The last line of the output shows the location and name of the binary, which in the case of `ripgrep` is `rg`. As long as the `PATH` is set correctly, you can run `rg` from anywhere.

`$PATH`, as mentioned previously, you can then run `rustier` faster, rustier tool for searching files!

## Extending Cargo with Custom

Cargo is designed so you can extend it with new subcommands. If a binary in your `$PATH` is named `cargo-foo`, it can be used as a Cargo subcommand by running `cargo foo`. This binary is also listed when you run `cargo --list`. To install extensions and then run them just like the standard Cargo commands, this is a convenient benefit of Cargo's design!

## Summary

Sharing code with Cargo and [crates.io](https://crates.io) is part of the Rust ecosystem. It's useful for many different tasks. Rust's standard library and crates are easy to share, use, and improve on a timeline. Don't be shy about sharing code that's useful to be useful to someone else as well!

## Smart Pointers

A *pointer* is a general concept for a variable that stores an address that refers to, or "points at," some other data. In Rust, a reference is a pointer, which you learned about in chapter 10. It's created by the `&` symbol and borrows the value they point to. References have capabilities other than referring to data. Also, there are other kinds of pointers we use most often.

*Smart pointers*, on the other hand, are data structures that point to data but also have additional metadata and capabilities. This capability isn't unique to Rust: smart pointers originated in C++ and are used well. In Rust, the different smart pointers define different capabilities beyond that provided by reference. The first smart pointer in this chapter is the *reference counting* smart pointer, which allows you to have multiple owners of data by keeping track of how many owners remain, cleaning up the data.

In Rust, which uses the concept of ownership and

between references and smart pointers is that references borrow data; in contrast, in many cases, smart pointers

We've already encountered a few smart pointers: `Vec<T>` in Chapter 8, although we didn't call these types smart pointers because they don't have the ability to manipulate it. They also have metadata (such as length) and capabilities or guarantees (such as with `String` being UTF-8).

Smart pointers are usually implemented using `Deref` and `Drop` traits. `Deref` distinguishes a smart pointer from an ordinary pointer. `Drop` implements the `Deref` and `Drop` traits. The `Deref` trait makes a smart pointer struct to behave like a reference struct. The `Drop` trait is either references or smart pointers. The `Drop` trait is that is run when an instance of the smart pointer is dropped. We'll discuss both traits and demonstrate why they're needed.

Given that the smart pointer pattern is a general pattern in Rust, this chapter won't cover every existing smart pointer, and you can even write your own smart pointers, and you can even write your own smart pointers in the standard library:

- `Box<T>` for allocating values on the heap
- `Rc<T>`, a reference counting type that enables shared ownership
- `Ref<T>` and `RefMut<T>`, accessed through `RefCell` and `RefMutCell`, which enforce the borrowing rules at runtime instead of compile time

In addition, we'll cover the *interior mutability* pattern, which exposes an API for mutating an interior value. We'll also discuss how they can leak memory and how to prevent them.

Let's dive in!

## Using `Box<T>` to Point to Data

The most straightforward smart pointer is a *box*. It allows you to store data on the heap rather than on the stack. The pointer to the heap data. Refer to Chapter 15 for more on the stack and the heap.

Boxes don't have performance overhead, other than the cost of allocating on the heap. But they don't have many other features.

them most often in these situations:

- When you have a type whose size can't be used to use a value of that type in a context that
- When you have a large amount of data and ensure the data won't be copied when you
- When you want to own a value and you can't have a particular trait rather than being of a specific

We'll demonstrate the first situation in the "Enums" section. In the second case, transferring ownership takes a long time because the data is copied around. In this situation, we can store the large amount of data on the heap; only the small amount of pointer data is copied. The third case, references, stays in one place on the heap. The third and Chapter 17 devotes an entire section, "Using References of Different Types," just to that topic. So what you'll see in Chapter 17!

## Using a `Box<T>` to Store Data on the Heap

Before we discuss this use case for `Box<T>`, we'll look at how to use it with values stored within a `Box<T>`.

Listing 15-1 shows how to use a box to store an `i32` value on the heap.

Filename: `src/main.rs`

```
fn main() {  
    let b = Box::new(5);  
    println!("b = {}", b);  
}
```

Listing 15-1: Storing an `i32` value on the heap using `Box`

We define the variable `b` to have the value of a `i32`. This value is allocated on the heap. This program will print the data in the box similar to how we would if it were an owned value, when a box goes out of scope, as the box is deallocated. The deallocation happens for the box, not the data it points to (stored on the heap).

Putting a single value on the heap isn't very useful by themselves in this way very often. Having values

they're stored by default, is more appropriate in at a case where boxes allow us to define types that didn't have boxes.

## Enabling Recursive Types with Boxes

At compile time, Rust needs to know how much memory whose size can't be known at compile time is a *recursive type* as part of itself another value of the same type. In theory, it could theoretically continue infinitely, Rust doesn't know how to handle recursive type needs. However, boxes have a known size. Once you have a recursive type definition, you can have recursive types.

Let's explore the *cons list*, which is a data type common in many programming languages, as an example of a recursive type. It's straightforward except for the recursion; therefore, it will be useful any time you get into more recursive types.

### More Information About the Cons List

A *cons list* is a data structure that comes from the Lisp family of dialects. In Lisp, the `cons` function (short for "construct pair" from its two arguments, which usually are a value and a list) forms a list.

The `cons` function concept has made its way into programming jargon: "to cons *x* onto *y*" means to create a new container instance by putting the element *x* at the front, followed by the container *y*.

Each item in a cons list contains two elements: the value and the next item. The last item in the list contains only the value. A cons list is produced by recursively calling the `cons` function. The name `Nil` is used to denote the base case of the recursion, as the "null" or "nil" concept in Chapter 6, which is the end of the list.

Although functional programming languages use cons lists, it isn't a commonly used data structure in Rust. For linked lists of items in Rust, `Vec<T>` is a better choice to use. Cons lists *are* useful in various situations, but by studying how boxes let us define a recursive data type with

Listing 15-2 contains an enum definition for a `cons` list that won't compile yet because the `List` type doesn't have a `Nil` value to demonstrate.

Filename: `src/main.rs`

```
enum List {  
    Cons(i32, List),  
    Nil,  
}
```

Listing 15-2: The first attempt at defining an enum structure of `i32` values

---

Note: We're implementing a `cons` list that holds `i32` values for the purposes of this example. We could have implemented a `Vec` type, as discussed in Chapter 10, to define a `cons` list that holds `T` values for any type `T`.

---

Using the `List` type to store the list `1, 2, 3` with `Nil` at the end

Filename: `src/main.rs`

```
use List::{Cons, Nil};  
  
fn main() {  
    let list = Cons(1, Cons(2, Cons(3, Nil)))  
}
```

Listing 15-3: Using the `List` enum to store the list `1, 2, 3`

The first `Cons` value holds `1` and another `List` value. The second `Cons` value holds `2` and another `List` value. The third `Cons` value holds `3` and a `List` value, which is `Nil`, that signals the end of the list.

If we try to compile the code in Listing 15-3, we get the following error:

```

error[E0072]: recursive type `List` has infinite size
--> src/main.rs:1:1
  |
1 | enum List {
  | ^^^^^^^^^ recursive type has infinite size
2 |     Cons(i32, List),
  |               ----- recursive without indirection
  |
  = help: insert indirection (e.g., a `Box`) into the type
to make `List` representable

```

Listing 15-4: The error we get when attempting to compile a recursive type

The error shows this type “has infinite size.” The problem is with a variant that is recursive: it holds another `List`, so the compiler can’t figure out how much space it needs to store it. To fix this, we get this error a bit. First, let’s look at how Rust stores a value of a non-recursive type.

## Computing the Size of a Non-Recursive Type

Recall the `Message` enum we defined in Listing 6-1, which has the definitions in Chapter 6:

```

enum Message {
    Quit,
    Move { x: i32, y: i32 },
    Write(String),
    ChangeColor(i32, i32, i32),
}

```

To determine how much space to allocate for a `Message`, we look at each of the variants to see which variant needs the most space. `Message::Quit` doesn’t need any space, `Message::Move` stores two `i32` values, and so forth. Because only one variant is stored, the space a `Message` value will need is the space it needs for its largest variant.

Contrast this with what happens when Rust tries to compute the size of a recursive type like the `List` enum in Listing 15-4. Looking at the `Cons` variant, which holds a value of type `List`. Therefore, `Cons` needs an amount of space equal to the size of a `List`. To figure out how much memory the compiler looks at the variants, starting with the

a value of type `i32` and a value of type `List`, as shown in Figure 15-1.

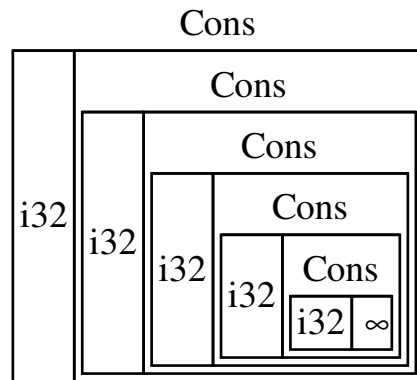


Figure 15-1: An infinite `List` consisting of infinite

### Using `Box<T>` to Get a Recursive Type with a

Rust can't figure out how much space to allocate, so the compiler gives the error in Listing 15-4. But the compiler's suggestion:

```
= help: insert indirection (e.g., a `Box`  
to  
make `List` representable
```

In this suggestion, “indirection” means that instead of changing the data structure to store the value directly, we change it to store a pointer to the value instead.

Because a `Box<T>` is a pointer, Rust always knows how much space to allocate for it. This means we can put a `Box<T>` inside the `Cons` variant directly. The `Box<T>` will point to the next `List` node rather than inside the `Cons` variant. Conceptually, we’re “holding” other lists, but this implementation is simpler because each list node points to one another rather than inside one another.

We can change the definition of the `List` enum in Listing 15-3 to the code in Listing 15-5, which uses `Box` to store the next list node.

Filename: src/main.rs



```

enum List {
    Cons(i32, Box<List>),
    Nil,
}

use List::{Cons, Nil};

fn main() {
    let list = Cons(1,
        Box::new(Cons(2,
            Box::new(Cons(3,
                Box::new(Nil))))));
}

```

Listing 15-5: Definition of `List` that uses `Box<T>`

The `Cons` variant will need the size of an `i32` pointer data. The `Nil` variant stores no values, so it doesn't need any size. We now know that any `List` value will take the size of a box's pointer data. By using a box, we've solved the problem so the compiler can figure out the size it needs to allocate. Figure 15-2 shows what the `Cons` variant looks like now.

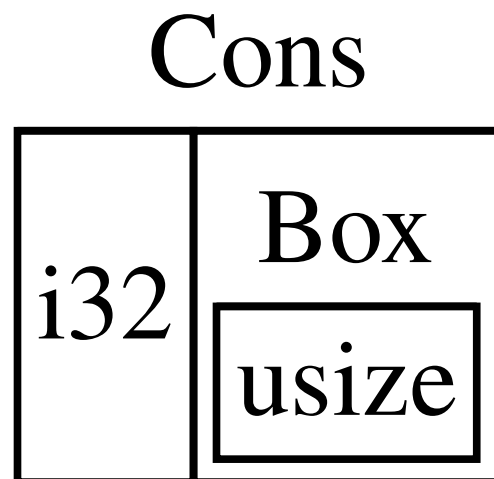


Figure 15-2: A `List` that is not infinitely sized because it uses boxes

Boxes provide only the indirection and heap allocation. They don't have any special capabilities, like those we'll see with the `Cell` and `RefCell` types. Boxes don't have any performance overhead that these other types do. They can be useful in cases like the cons list where they're needed to store pointers. We'll look at more use cases for boxes in Chapter 16.

The `Box<T>` type is a smart pointer because it allows `Box<T>` values to be treated like references. In this scope, the heap data that the box is pointing to is managed by the `Drop` trait implementation. Let's explore these traits; they will be even more important to the function pointer types we'll discuss in the rest of this chapter.

## Treating Smart Pointers Like References with the `Deref` Trait

Implementing the `Deref` trait allows you to customize the dereference operator, `*` (as opposed to the multiplication operator `*`) in such a way that a smart pointer can be used like a reference. You can write code that operates on references and smart pointers interchangeably.

Let's first look at how the dereference operator works. Then we'll try to define a custom type that behaves like a reference. We'll see how the dereference operator doesn't work like a reference. Finally, we'll explore how implementing the `Deref` trait makes smart pointers work in a similar way as references. Then we'll look at how it lets us work with either references or smart pointers.

---

There's one big difference between the `MyBox` and the real `Box<T>`: our version will not store its data. In this example, we'll use the `Deref` trait, and so where the data is stored is less important than the pointer-like behavior.

---

## Following the Pointer to the Value with `Deref`

A regular reference is a type of pointer, and one that points to a value stored somewhere else. In Listing 13-2, we'll use the `Deref` trait to dereference a `i32` value and then use the dereference operator to access the value.

Filename: `src/main.rs`

```
fn main() {  
    let x = 5;  
    let y = &x;  
  
    assert_eq!(5, x);  
    assert_eq!(5, *y);  
}
```

### Listing 15-6: Using the dereference operator to f

The variable `x` holds an `i32` value, `5`. We set `y` to point to `x` and then we assert that `x` is equal to `5`. However, if we want to compare the value in `y`, we have to use `*y` to follow the reference (hence *dereference*). Once we dereference `y`, we get the value `5` pointing to that we can compare with `5`.

If we tried to write `assert_eq!(5, y);` instead,

```
error[E0277]: the trait bound `{integer}:
is
not satisfied
--> src/main.rs:6:5
|
6 |         assert_eq!(5, y);
|           ^^^^^^^^^^^^^^^^^^ can't compare `{integer}`
= help: the trait `std::cmp::PartialEq` is not implemented for `{integer}`
```

Comparing a number and a reference to a number are two different types. We must use the dereference operator to get the value it's pointing to.

## Using Box<T> Like a Reference

We can rewrite the code in Listing 15-6 to use a dereference operator will work as shown in Listi

Filename: src/main.rs

```
fn main() {
    let x = 5;
    let y = Box::new(x);

    assert_eq!(5, x);
    assert_eq!(5, *y);
}
```

Listing 15-7: Using the dereference operator on

The only difference between Listing 15-7 and Listing 15-6 is that `y` is an instance of a `Box` pointing to the value in `x` rather than the value of `x`. In the last assertion, we can use the box's pointer in the same way that we did when we explored `Box<T>` to enable dereferencing the pointer operator by defining our own box type.

## Defining Our Own Smart Pointer

Let's build a smart pointer similar to the `Box<T>` type in the `std::boxed` library to experience how smart pointers behave. We'll start with the default. Then we'll look at how to add the ability to dereference the pointer.

The `Box<T>` type is ultimately defined as a tuple `(T, usize)`. Listing 15-8 defines a `MyBox<T>` type in the same way. The `new` function matches the `new` function defined on `Box<T>`.

Filename: src/main.rs

```
struct MyBox<T>(T);

impl<T> MyBox<T> {
    fn new(x: T) -> MyBox<T> {
        MyBox(x)
    }
}
```

Listing 15-8: Defining a `MyBox<T>` type

We define a struct named `MyBox` and declare a `new` function. We want our type to hold values of any type. The `MyBox` struct has one element of type `T`. The `MyBox::new` function takes a value of type `T` and returns a `MyBox` instance that holds the value passed in.

Let's try adding the `main` function in Listing 15-7 the `MyBox<T>` type we've defined instead of `Box` compile because Rust doesn't know how to dere

Filename: src/main.rs

```
fn main() {  
    let x = 5;  
    let y = MyBox::new(x);  
  
    assert_eq!(5, x);  
    assert_eq!(5, *y);  
}
```

Listing 15-9: Attempting to use `MyBox<T>` in the `Box<T>`

Here's the resulting compilation error:

```
error[E0614]: type `MyBox<{integer}>` canr  
--> src/main.rs:14:19  
    |  
14 |         assert_eq!(5, *y);  
    |                        ^^
```

Our `MyBox<T>` type can't be dereferenced because of the lack of dereferencing ability on our type. To enable dereferencing with the `Deref` trait.

## Treating a Type Like a Reference by Implementing the Deref Trait

As discussed in Chapter 10, to implement a trait for the trait's required methods. The `Deref` trait requires us to implement one method named `deref` that returns a reference to the inner data. Listing 15-10 contains the code to add to the definition of `MyBox`:

Filename: src/main.rs

```

use std::ops::Deref;

impl<T> Deref for MyBox<T> {
    type Target = T;

    fn deref(&self) -> &T {
        &self.0
    }
}

```

Listing 15-10: Implementing `Deref` on `MyBox<T>`

The `type Target = T;` syntax defines an associated type. Associated types are a slightly different way of defining types, but you don't need to worry about them for now; we'll cover them in chapter 19.

We fill in the body of the `deref` method with `&self.0` to return a reference to the value we want to access with the `*` operator. Now that `MyBox<T>` implements `Deref`, the `*` operator that calls `*` on the `MyBox<T>` value now compiles.

Without the `Deref` trait, the compiler can only call `deref` if you call the `deref` method. The `Deref` trait method gives the compiler the ability to take a value of type `MyBox<T>` and call the `deref` method to get a `&T` reference.

When we entered `*y` in Listing 15-9, behind the scenes the compiler substituted

```
*(&y.deref())
```

Rust substitutes the `*` operator with a call to the `deref` method. This Rust feature lets us write code that works whether we have a regular reference or a type that implements `Deref`.

The reason the `deref` method returns a reference is to avoid a double dereference outside the parentheses in `*(&y.deref())`. In Rust's ownership system, if the `deref` method returned a value instead of a reference to the value, the value would be moved out of the `MyBox<T>`, and you would lose ownership of the inner value inside `MyBox<T>` if you didn't use the dereference operator.

Note that the `*` operator is replaced with a call to the `deref` method just once, each time we use a `*` operator. The substitution of the `*` operator does not recurse.

type `i32`, which matches the `5` in `assert_eq!`

## Implicit Deref Coercions with Function

*Deref coercion* is a convenience that Rust performs on function calls. Deref coercion converts a reference to a type that `Deref` can convert to the type that the function or method expects. This happens automatically when we pass a reference to a function or method that doesn't take a reference. A sequence of calls to `deref` is performed until the type we provided into the type the parameter needs.

Deref coercion was added to Rust so that programmers don't need to add as many explicit references. The deref coercion feature also lets us write more concise references or smart pointers.

To see deref coercion in action, let's use the `MyBox` as well as the implementation of `Deref` that we saw in Listing 15-10. Listing 15-11 shows the definition of a function that has a string parameter.

Filename: `src/main.rs`

```
fn hello(name: &str) {  
    println!("Hello, {}!", name);  
}
```

Listing 15-11: A `hello` function that has the parameter `&str`

We can call the `hello` function with a string slice, `hello("Rust");` for example. Deref coercion makes it possible to pass a reference to a value of type `MyBox<String>`, as shown in Listing 15-12.

Filename: `src/main.rs`

```
fn main() {  
    let m = MyBox::new(String::from("Rust"));  
    hello(&m);  
}
```

Listing 15-12: Calling `hello` with a reference to `MyBox` because of deref coercion

Here we're calling the `hello` function with the argument `&m`, which is a reference to the `MyBox` object `m`.

a `MyBox<String>` value. Because we implement Listing 15-10, Rust can turn `&MyBox<String>` into a `String` slice, and this is in the API documentation. Rust can turn the `&String` into `&str`, which matches the

If Rust didn't implement deref coercion, we would have to write Listing 15-13 instead of the code in Listing 15-12 to call `hello` on `&MyBox<String>`.

Filename: src/main.rs

```
fn main() {  
    let m = MyBox::new(String::from("Rust"  
    hello(&(*m)[..]));  
}
```

Listing 15-13: The code we would have to write if Rust didn't have deref coercion

The `(*m)` dereferences the `MyBox<String>` into a `String`. Then `[..]` takes a string slice of the `String` that is equal to the signature of `hello`. The code without deref coercion is more verbose and harder to understand with all of these symbols involved. Rust does these conversions for us automatically.

When the `Deref` trait is defined for the types in use, you can use `Deref::deref` as many times as necessary to get to the parameter's type. The number of times that `Deref` is used is resolved at compile time, so there is no runtime coercion!

## How Deref Coercion Interacts with Mutable References

Similar to how you use the `Deref` trait to override deref coercion for immutable references, you can use the `DerefMut` trait to override deref coercion for mutable references.

Rust does deref coercion when it finds types and

- From `&T` to `&U` when `T: Deref<Target=U>`
- From `&mut T` to `&mut U` when `T: DerefMut<Target=U>`
- From `&mut T` to `&U` when `T: Deref<Target=U>`

The first two cases are the same except for mutable references.



have a `&T`, and `T` implements `Deref` to some `t` transparently. The second case states that the smart pointer has mutable references.

The third case is trickier: Rust will also coerce a mutable reference to an immutable one. But the reverse is *not* possible: immutable references to mutable references. Because of the borrowing rules that require that a mutable reference must be the only reference to the data (otherwise the program wouldn't compile). Converting one mutable reference to an immutable reference will never break the borrowing rules. Converting an immutable reference to a mutable reference would require that there is only one reference to the data, and the borrowing rules don't guarantee that. So, the assumption that converting an immutable reference to a mutable reference is possible.

## Running Code on Cleanup with `Drop`

The second trait important to the smart pointer is `Drop`. It allows you to customize what happens when a value is about to be destroyed. You implement an implementation for the `Drop` trait on any type that represents a resource used to release resources like files or network connections. This is the context of smart pointers because the function `drop` is always used when implementing a smart pointer. The `Drop` trait is used to deallocate the space on the heap that the smart pointer occupies.

In some languages, the programmer must call a cleanup function every time they finish using an instance of a smart pointer. This might become overloaded and crash. In Rust, you can write code to be run whenever a value goes out of scope. This code is called a `Drop` implementation. As a result, you don't need to write cleanup code everywhere in a program that an instance of a smart pointer is used. You still won't leak resources!

Specify the code to run when a value goes out of scope by implementing the `Drop` trait. The `Drop` trait requires you to implement a function that takes a mutable reference to `self`. To see when Rust calls `Drop`, use `println!` statements for now.

Listing 15-14 shows a `CustomSmartPointer` struct that implements `Drop` so that it will print `Dropping CustomSmartPointer!` when it goes out of scope. This example demonstrates when Rust runs the `Drop` implementation.

Filename: src/main.rs

```
struct CustomSmartPointer {  
    data: String,  
}  
  
impl Drop for CustomSmartPointer {  
    fn drop(&mut self) {  
        println!("Dropping CustomSmartPoi  
self.data);  
    }  
}  
  
fn main() {  
    let c = CustomSmartPointer { data: Str  
    let d = CustomSmartPointer { data: Str  
    println!("CustomSmartPointers created.  
}
```

Listing 15-14: A `CustomSmartPointer` struct that would put our cleanup code

The `Drop` trait is included in the prelude, so we implement the `Drop` trait on `CustomSmartPoint` for the `drop` method that calls `println!`. The `drop` method is the last method you would place any logic that you wanted to run before the variable goes out of scope. We're printing some text here to demonstrate.

In `main`, we create two instances of `CustomSmartPointer`. At the end of `main`, `CustomSmartPointer` will go out of scope, and Rust will call the `drop` method, printing our final message. Note that we call `drop` method explicitly.

When we run this program, we'll see the following output:

```
CustomSmartPointers created.  
Dropping CustomSmartPointer with data `other`  
Dropping CustomSmartPointer with data `my`
```

Rust automatically calls `drop` for us when our variable goes out of scope. Variables are dropped in reverse order of creation. `d` was dropped before `c`. This example gives you a basic idea of how the `Drop` method works; usually you would specify the cleanup logic to run rather than a print message.

## Dropping a Value Early with `std::mem::`

Unfortunately, it's not straightforward to disable `Drop`. Disabling `Drop` isn't usually necessary; the whole program is taken care of automatically. Occasionally, however, you may want to drop a value early. One example is when using smart pointers. You may want to force the `drop` method that releases the memory. The same scope can acquire the lock. Rust doesn't let you call the `drop` method manually; instead you have to call the `std::mem::drop` function from the standard library if you want to force a value to be dropped from scope.

If we try to call the `Drop` trait's `drop` method manually, we get an error. Listing 15-14, as shown in Listing 15-14.

Filename: src/main.rs

```
fn main() {
    let c = CustomSmartPointer { data: String::new() };
    println!("CustomSmartPointer created.");
    c.drop();
    println!("CustomSmartPointer dropped before the end of main.");
}
```

Listing 15-15: Attempting to call the `drop` method manually to clean up early

When we try to compile this code, we'll get this error:

```
error[E0040]: explicit use of destructor not allowed
  --> src/main.rs:14:7
   |
14 |         c.drop();
   |         ^^^^^ explicit destructor calls
```

This error message states that we're not allowed to call `drop` explicitly. The error message uses the term *destructor*, which is the function that cleans up an instance. A *destructor* is a function that creates an instance. The `drop` function in Rust is a destructor.

Rust doesn't let us call `drop` explicitly because it would be trying to clean up the same value twice. The `Drop` trait's `drop` method is called automatically at the end of `main`. This would be trying to clean up the same value twice.

We can't disable the automatic insertion of `Drop`. We can't call the `drop` method explicitly. So, if we want to drop a value early, we have to use `std::mem::drop`.

up early, we can use the `std::mem::drop` function.

The `std::mem::drop` function is different than `std::mem::forget`. We call it by passing the value we want to force to be dropped. The function is in the prelude, so we can modify the `drop` function, as shown in Listing 15-16:

Filename: src/main.rs

```
fn main() {  
    let c = CustomSmartPointer { data: String::new() };  
    println!("CustomSmartPointer created.");  
    drop(c);  
    println!("CustomSmartPointer dropped before the end of main.");  
}
```

Listing 15-16: Calling `std::mem::drop` to explicitly drop a value at the end of scope

Running this code will print the following:

```
CustomSmartPointer created.  
Dropping CustomSmartPointer with data `some data`.  
CustomSmartPointer dropped before the end of main.
```

The text `Dropping CustomSmartPointer with data `some data`.` and `CustomSmartPointer dropped before the end of main.` are printed because the `drop` method code is called to drop `c` at that point.

You can use code specified in a `Drop` trait implementation to perform cleanup convenient and safe: for instance, you can use it to free memory allocated by a C memory allocator! With the `Drop` trait and Rust's garbage collection, you don't have to remember to clean up because Rust does it for you.

You also don't have to worry about problems related to dangling pointers or values still in use: the ownership system that manages memory in Rust also ensures that `drop` gets called only once when a value goes out of scope.

Now that we've examined `Box<T>` and some of its methods, let's look at a few other smart pointers defined in the `std::rc` module.

## `Rc<T>`, the Reference Counter

In the majority of cases, ownership is clear: you own a given value. However, there are cases when a single value has multiple owners. For example, in graph data structures, multiple pointers can point to the same node, and that node is conceptually owned by all of them. In such cases, a node shouldn't be cleaned up unless it doesn't have any more owners.

To enable multiple ownership, Rust has a type called `Rc` for *reference counting*. The `Rc<T>` type keeps track of the number of references to a value, which determines whether or not a value is still valid. When the number of references reaches zero, the value can be cleaned up and its memory is invalid.

Imagine `Rc<T>` as a TV in a family room. When the first person turns it on, it starts playing. Others can come into the room and watch it. When the last person leaves the room, they turn off the TV because it's not being watched. If the first person turns off the TV while others are still watching it, the TV will be turned off for the remaining TV watchers!

We use the `Rc<T>` type when we want to allocate memory for data that will be shared between parts of our program to read and we can't determine which part will finish using the data last. If we knew which part would finish using the data last, that part would be the data's owner, and the normal ownership rules would take effect.

Note that `Rc<T>` is only for use in single-threaded programs. In Chapter 16, we'll cover how to deal with concurrency in multi-threaded programs.

## Using `Rc<T>` to Share Data

Let's return to our cons list example in Listing 15-1. In this example, we'll use `Rc<T>` to share data between two lists. Conceptually, this looks similar to Figure 15-3:

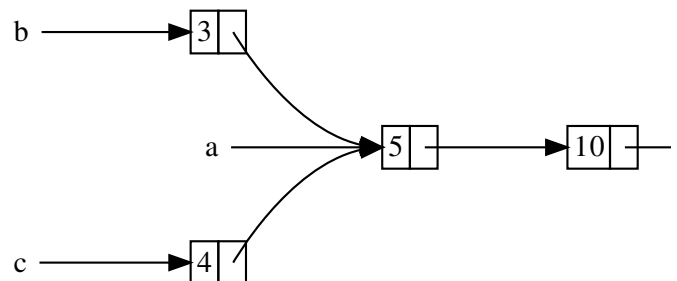


Figure 15-3: Two lists, `b` and `c`, sharing owners

We'll create list `a` that contains 5 and then 10. `T` starts with 3 and `c` that starts with 4. Both `b` and `c` are first `a` list containing 5 and 10. In other words, both `b` and `c` are containing 5 and 10.

Trying to implement this scenario using our definition, as shown in Listing 15-17:

Filename: `src/main.rs`

```
enum List {
    Cons(i32, Box<List>),
    Nil,
}

use List::{Cons, Nil};

fn main() {
    let a = Cons(5,
        Box::new(Cons(10,
            Box::new(Nil))));
    let b = Cons(3, Box::new(a));
    let c = Cons(4, Box::new(a));
}
```

Listing 15-17: Demonstrating we're not allowed to share ownership of a third list

When we compile this code, we get this error:

```
error[E0382]: use of moved value: `a`
--> src/main.rs:13:30
   |
12 |         let b = Cons(3, Box::new(a));
   |                                - value moved here
13 |         let c = Cons(4, Box::new(a));
   |                                ^ value used here
   |
   = note: move occurs because `a` has type `List`, which does not implement the `Copy` trait
```

The `Cons` variants own the data they hold, so when we move `a` into `b` and `b` owns `a`. Then, when we try to use `a` again, it's not allowed to because `a` has been moved.

We could change the definition of `Cons` to hold

would have to specify lifetime parameters. By specifying that every element in the list has a lifetime of 'static', the borrow checker wouldn't let us compile the example, because the temporary `Nil` value would have no reference to it.

Instead, we'll change our definition of `List` to use `Rc`, as shown in Listing 15-18. Each `Cons` variant will now point to a `List`. When we create `b`, instead of creating a new `List`, we'll clone the `Rc<List>` that `a` is holding, thereby increasing the reference count from one to two and letting `a` and `b` share ownership. Similarly, when we create `c`, we'll clone `a`, increasing the reference count to three. Every time we call `Rc::clone`, the reference count of the `Rc<List>` will increase, and the data won't be cloned.

Filename: src/main.rs

```
enum List {
    Cons(i32, Rc<List>),
    Nil,
}

use List::{Cons, Nil};
use std::rc::Rc;

fn main() {
    let a = Rc::new(Cons(5, Rc::new(Cons(10, Nil))));
    let b = Cons(3, Rc::clone(&a));
    let c = Cons(4, Rc::clone(&a));
}
```

Listing 15-18: A definition of `List` that uses `Rc`

We need to add a `use` statement to bring `Rc<T>` into scope. In `main`, we create the list holding 5 and 10 in `a`. Then when we create `b` and `c`, we call `Rc::clone` to clone the reference to the `Rc<List>` in `a` as an argument.

We could have called `a.clone()` rather than `Rc::clone` to use `Rc::clone` in this case. The implementation of `clone` for `Rc` only increments the reference count, which is much faster than making a deep copy of all the data like most types' `clone` implementations. Cloning copies of data can take a lot of time. By using `Rc`, we can visually distinguish between the deep-copy clones and the shallow clones that increase the reference count. When

```
Rc::clone .
```

## Cloning an `Rc<T>` Increases the Refere

Let's change our working example in Listing 15-1 changing as we create and drop references to th

In Listing 15-19, we'll change `main` so it has an `int` variable that can see how the reference count changes when

Filename: src/main.rs

```
fn main() {  
    let a = Rc::new(Cons(5, Rc::new(Cons(1,  
println!("count after creating a = {}",  
    let b = Cons(3, Rc::clone(&a));  
println!("count after creating b = {}",  
    {  
        let c = Cons(4, Rc::clone(&a));  
        println!("count after creating c = ",  
    }  
    println!("count after c goes out of scope  
Rc::strong_count(&a));  
}
```

### Listing 15-19: Printing the reference count

At each point in the program where the reference count, which we can get by calling the function is named `strong_count` rather than `count`, we have a `weak_count`; we'll see what `weak_count` is in the "Reference Cycles" section.

This code prints the following:

```
count after creating a = 1
count after creating b = 2
count after creating c = 3
count after c goes out of scope = 2
```

We can see that the `Rc<List>` in `a` has an initial count of 1. When we call `clone`, the count goes up by 1. When `a` goes out of scope, the count goes down by 1. We don't have to call a function to decrease the count; the `Rc::clone` trait has to call `Rc::clone` to increase the reference count. The `Drop` trait decreases the reference count automatically.



out of scope.

What we can't see in this example is that when | end of `main`, the count is then 0, and the `Rc<Li` point. Using `Rc<T>` allows a single value to have ensures that the value remains valid as long as a

Via immutable references, `Rc<T>` allows you to your program for reading only. If `Rc<T>` allowed references too, you might violate one of the bor multiple mutable borrows to the same place car inconsistencies. But being able to mutate data is discuss the interior mutability pattern and the `Rc` conjunction with an `Rc<T>` to work with this imr

## `RefCell<T>` and the Interior M

*Interior mutability* is a design pattern in Rust that when there are immutable references to that da by the borrowing rules. To mutate data, the patt structure to bend Rust's usual rules that govern yet covered unsafe code; we will in Chapter 19. \ mutability pattern when we can ensure that the runtime, even though the compiler can't guaran then wrapped in a safe API, and the outer type is

Let's explore this concept by looking at the `RefC` mutability pattern.

## Enforcing Borrowing Rules at Runtime

Unlike `Rc<T>`, the `RefCell<T>` type represents holds. So, what makes `RefCell<T>` different fro borrowing rules you learned in Chapter 4:

- At any given time, you can have *either* (but or any number of immutable references.
- References must always be valid.

With references and `Box<T>`, the borrowing rule time. With `RefCell<T>`, these invariants are enf

you break these rules, you'll get a compiler error. If you break these rules, your program will panic and exit.

The advantages of checking the borrowing rules can be caught sooner in the development process, and performance because all the analysis is complete at compile time. Checking the borrowing rules at compile time is the default in Rust, which is why this is Rust's default.

The advantage of checking the borrowing rules at compile time is that memory-safe scenarios are then allowed, where compile-time checks. Static analysis, like the Rust compiler, can't detect some properties of code. Some properties of code are impossible to detect at compile time. A famous example is the Halting Problem, which is an interesting topic to research.

Because some analysis is impossible, if the Rust compiler complies with the ownership rules, it might reject some programs that are conservative. If Rust accepted an incorrect program, it would break the guarantees Rust makes. However, if Rust rejects a program, the programmer will be inconvenienced, but nothing is lost. The `RefCell<T>` type is useful when you're sure you can mutate the value inside, but the compiler is unable to understand and guarantee it.

Similar to `Rc<T>`, `RefCell<T>` is only for use in mutable contexts. It will give you a compile-time error if you try using it in an immutable context. Chapter 16 shows about how to get the functionality of `RefCell<T>`.

Here is a recap of the reasons to choose `Box<T>`:

- `Rc<T>` enables multiple owners of the same data, while `Box<T>` has single owners.
- `Box<T>` allows immutable or mutable borrows, while `Rc<T>` allows only immutable borrows checked at runtime.
- Because `RefCell<T>` allows mutable borrows, you can mutate the value inside the `RefCell<T>` even if it's behind an immutable reference.

Mutating the value inside an immutable value is not allowed. We'll look at a situation in which interior mutability is needed.

## Interior Mutability: A Mutable Borrow

A consequence of the borrowing rules is that we can't borrow it mutably. For example, this code won't compile:

```
fn main() {  
    let x = 5;  
    let y = &mut x;  
}
```

If you tried to compile this code, you'd get the following error:

```
error[E0596]: cannot borrow immutable local variable `x` as mutable  
--> src/main.rs:3:18  
  |  
2 |     let x = 5;  
  |           - consider changing this to `mut x`  
3 |     let y = &mut x;  
  |               ^ cannot borrow mutable variable as immutable
```

However, there are situations in which it would be useful to have interior mutability. In Rust, we can achieve this by using the `RefCell` type. `RefCell` allows us to have interior mutability, meaning that a value can be mutable even if it's behind an immutable reference. This is useful in situations where we need to mutate a value in its methods but appear immutable to other code. For example, a counter that increments its value in its methods would not be able to mutate the value if it's behind an immutable reference. But `RefCell` allows us to have interior mutability. But `RefCell` doesn't change the borrowing rules completely: the borrow checker still enforces the borrowing rules for mutability, and the borrowing rules are checked at runtime. If you violate the borrowing rules, you'll get a `panic!` instead of a compiler error.

Let's work through a practical example where we use `RefCell` to create a mutable immutable value and see why that is useful.

## A Use Case for Interior Mutability: Mock Objects

A *test double* is the general programming concept used to replace a real object with a mock object type during testing. *Mock objects* are specific types of test doubles that are used to simulate the behavior of a real object that happens during a test so you can assert that the code behaves as expected.

Rust doesn't have objects in the same sense as C++ or Java. Rust doesn't have mock object functionality built into the language like other languages do. However, you can definitely simulate the same purposes as a mock object.

Here's the scenario we'll test: we'll create a library that takes a maximum value and sends messages based on the current value. This library could be used to keep track of the number of API calls they're allowed to make, for example.

Our library will only provide the functionality of the `Counter` trait.

value is and what the messages should be at which the library will be expected to provide the mechanism. An application could put a message in the application's message, or something else. The library doesn't know what is something that implements a trait we'll provide. Listing 15-20 shows the library code:

Filename: src/lib.rs

```
pub trait Messenger {
    fn send(&self, msg: &str);
}

pub struct LimitTracker<'a, T: 'a + Messenger> {
    messenger: &'a T,
    value: usize,
    max: usize,
}

impl<'a, T> LimitTracker<'a, T>
    where T: Messenger {
    pub fn new(messenger: &T, max: usize)
        LimitTracker {
            messenger,
            value: 0,
            max,
        }

    pub fn set_value(&mut self, value: usize) {
        self.value = value;

        let percentage_of_max = self.value as f64 / self.max as f64;

        if percentage_of_max >= 0.75 && percentage_of_max < 1.0 {
            self.messenger.send("Warning: You are at 75% of your quota!");
        } else if percentage_of_max >= 0.9 && percentage_of_max < 1.0 {
            self.messenger.send("Urgent warning: You are at 90% of your quota!");
        } else if percentage_of_max >= 1.0 {
            self.messenger.send("Error: You have exceeded your quota!");
        }
    }
}
```

Listing 15-20: A library to keep track of how close a value is to a maximum and warn when the value is at certain levels

One important part of this code is that the `Messenger` trait has a `send` method that takes an immutable reference to `Self` and a `value` parameter. This is the interface our mock object needs to have. This is important because to test the behavior of the `set_value` method, we need to know what we pass in for the `value` parameter, but we also need a way for us to make assertions on it. We want to be able to assert that the `value` is with something that implements the `Messenger` trait. When we pass different numbers for `value`, the `set_value` method will send appropriate messages.

We need a mock object that, instead of sending messages, will only keep track of the messages it receives. To create an instance of the mock object, create a `LimitTracker` struct. Implement the `set_value` method on `LimitTracker`, and let it store the messages we expect. Listing 15-21 shows an example of how to do just that, but the borrow checker won't allow it.

Filename: src/lib.rs

```

#[cfg(test)]
mod tests {
    use super::*;

    struct MockMessenger {
        sent_messages: Vec<String>,
    }

    impl MockMessenger {
        fn new() -> MockMessenger {
            MockMessenger { sent_messages:
        }

    impl Messenger for MockMessenger {
        fn send(&self, message: &str) {
            self.sent_messages.push(String

    }

    #[test]
    fn it_sends_an_over_75_percent_warning
        let mock_messenger = MockMessenger
        let mut limit_tracker = LimitTrack

        limit_tracker.set_value(80);

        assert_eq!(mock_messenger.sent_mes
    }
}

```

Listing 15-21: An attempt to implement a `MockMessenger` borrow checker

This test code defines a `MockMessenger` struct that holds a `Vec` of `String` values to keep track of the messages sent. It also has an associated function `new` to make it convenient to create `MockMessenger` values that start with an empty list of messages. Then, we implement the `Messenger` trait for `MockMessenger` so we can give a `MockMessenger` a `send` method. In the definition of the `send` method, we take the message as a parameter and store it in the `MockMessenger` list of `sent_messages`.

In the test, we're testing what happens when the limit is set to something that is more than 75 percent of the total. We create a `MockMessenger`, which will start with an empty list of messages, and a `LimitTracker` and give it a reference to the new `MockMessenger`. We call the `set_value` method on the `LimitTracker` to set the value to 80.

more than 75 percent of 100. Then we assert that `MockMessenger` is keeping track of should now be

However, there's one problem with this test, as :

```
error[E0596]: cannot borrow immutable field
mutable
--> src/lib.rs:52:13
   |
51 |         fn send(&self, message: &str)
   |               ----- use `&mut self`
52 |         self.sent_messages.push($
   |                               ^^^^^^^^^^^^^^^^^^^^^^^^^ cannot
```

We can't modify the `MockMessenger` to keep track of the state. The `send` method takes an immutable reference to `self`. If we change the error text to use `&mut self` instead, but the signature in the `Messenger` trait wouldn't match the signature in the `Messenger` trait (what error message you get).

This is a situation in which interior mutability can be used. We can wrap `sent_messages` within a `RefCell<T>`, and then the `send` message can borrow `sent_messages` to store the messages we've seen. The code looks like:

Filename: src/lib.rs

```

#[cfg(test)]
mod tests {
    use super::*;
    use std::cell::RefCell;

    struct MockMessenger {
        sent_messages: RefCell<Vec<String>
    }

    impl MockMessenger {
        fn new() -> MockMessenger {
            MockMessenger { sent_messages:
        }

    impl Messenger for MockMessenger {
        fn send(&self, message: &str) {
            self.sent_messages.borrow_mut(
        }

    #[test]
    fn it_sends_an_over_75_percent_warning
        // --snip--

        assert_eq!(mock_messenger.sent_mes
    }
}

```

Listing 15-22: Using `RefCell<T>` to mutate an ir considered immutable

The `sent_messages` field is now of type `RefCell<Vec<String>`. In the `new` function, we create a around the empty vector.

For the implementation of the `send` method, th borrow of `self`, which matches the trait definit `RefCell<Vec<String>>` in `self.sent_messages` value inside the `RefCell<Vec<String>>`, which i on the mutable reference to the vector to keep t test.

The last change we have to make is in the assert inner vector, we call `borrow` on the `RefCell<Ve` reference to the vector.

Now that you've seen how to use `RefCell<T>`, l



## Keeping Track of Borrows at Runtime with `RefCell`

When creating immutable and mutable references respectively. With `RefCell<T>`, we use the `borrow` and `borrow_mut` methods, which are part of the safe API that belongs to `RefCell`. `RefCell` is a smart pointer type `Ref<T>`, and `borrow_mut` returns a mutable reference `RefMut<T>`. Both types implement `Deref`, so we can use them to access references.

The `RefCell<T>` keeps track of how many `Ref<T>` references are currently active. Every time we call `borrow`, it increments the count of how many immutable borrows are active. When the count of immutable borrows goes down by one, it decrements the count. According to the borrowing rules, `RefCell<T>` lets us have many immutable borrows at any point in time.

If we try to violate these rules, rather than getting a compile-time error, the implementation of `RefCell<T>` will panic. Listing 15-23 shows a modification of the implementation of `RefCell` that deliberately tries to create two mutable borrows. This illustrates that `RefCell<T>` prevents us from doing so.

Filename: src/lib.rs

```
impl Messenger for MockMessenger {
    fn send(&self, message: &str) {
        let mut one_borrow = self.sent_messages.borrow_mut();
        let mut two_borrow = self.sent_messages.borrow_mut();

        one_borrow.push(String::from(message));
        two_borrow.push(String::from(message));
    }
}
```

Listing 15-23: Creating two mutable references in the same scope. `RefCell<T>` will panic

We create a variable `one_borrow` for the `RefMut<T>` mutable reference. Then we create another mutable borrow `two_borrow`. This makes two mutable references to the same data, which is not allowed. When we run the tests for our library, the tests pass without any errors, but the test will fail:

```

---- tests::it_sends_an_over_75_percent_wa
      thread 'tests::it_sends_an_over_75_per
at
'already borrowed: BorrowMutError', src/li
note: Run with `RUST_BACKTRACE=1` for a ba

```

Notice that the code panicked with the message  
. This is how `RefCell<T>` handles violations of t

Catching borrowing errors at runtime rather than  
find a mistake in your code later in the development  
your code was deployed to production. Also, you  
performance penalty as a result of keeping track  
than compile time. However, using `RefCell<T>`  
object that can modify itself to keep track of the  
using it in a context where only immutable values  
`RefCell<T>` despite its trade-offs to get more flexibility  
provide.

## Having Multiple Owners of Mutable Data `RefCell<T>`

A common way to use `RefCell<T>` is in combination  
lets you have multiple owners of some data, but  
data. If you have an `Rc<T>` that holds a `RefCell`  
have multiple owners *and* that you can mutate!

For example, recall the cons list example in Listing 15-23  
allow multiple lists to share ownership of another  
immutable values, we can't change any of the values  
them. Let's add in `RefCell<T>` to gain the ability  
Listing 15-24 shows that by using a `RefCell<T>`  
the value stored in all the lists:

Filename: src/main.rs

```

#[derive(Debug)]
enum List {
    Cons(Rc<RefCell<i32>>, Rc<List>),
    Nil,
}

use List::{Cons, Nil};
use std::rc::Rc;
use std::cell::RefCell;

fn main() {
    let value = Rc::new(RefCell::new(5));

    let a = Rc::new(Cons(Rc::clone(&value), Nil));

    let b = Cons(Rc::new(RefCell::new(6)), Nil);
    let c = Cons(Rc::new(RefCell::new(10)), Nil);

    *value.borrow_mut() += 10;

    println!("a after = {:?}", a);
    println!("b after = {:?}", b);
    println!("c after = {:?}", c);
}

```

Listing 15-24: Using `Rc<RefCell<i32>>` to create a linked list

We create a value that is an instance of `Rc<RefCell<i32>>` named `value` so we can access it directly later. We then create a `Cons` variant that holds `value`. We need to clone `value` to transfer ownership of the inner `5` value rather than transferring ownership to `a` or having `a` borrow from `value`.

We wrap the list `a` in an `Rc<List>` so when we create `a`, which is what we did in Listing 15-18.

After we've created the lists in `a`, `b`, and `c`, we modify `value` by calling `borrow_mut` on `value`, which we discussed in Chapter 5 (see the section "When to Use `RefCell`"). This returns a `RefMut<i32>` smart pointer, and we use it to change the inner value.

When we print `a`, `b`, and `c`, we can see that the value of `a` is now 15 rather than 5:

```
a after = Cons(RefCell { value: 15 }, Nil)
b after = Cons(RefCell { value: 6 }, Cons(
c after = Cons(RefCell { value: 10 }, Cons
```

This technique is pretty neat! By using `RefCell<List>` value. But we can use the methods on `RefCell` to get interior mutability so we can modify our data without violating the borrowing rules protect us from data races. It's a bit of speed for this flexibility in our data structure.

The standard library has other types that provide interior mutability, `RefCell`, which is similar except that instead of giving you a mutable reference, it's copied in and out of the `Cell<T>`. There's also `Atomic` mutability that's safe to use across threads; we'll look at that in the standard library docs for more details on these types.

## Reference Cycles Can Leak Memory

Rust's memory safety guarantees make it difficult to create memory that is never cleaned up (known as a memory leak). Entirely avoiding memory leaks is not one of Rust's guarantees. Disallowing data races at compile time is, meaning Rust ensures that references are dropped. We can see that Rust allows memory leaks by creating references where items refer to each other, creating a cycle. If the reference count of each item in the cycle never reaches 0, and the values will never be dropped.

### Creating a Reference Cycle

Let's look at how a reference cycle might happen. We'll look at the definition of the `List` enum and a `tail` method.

Filename: src/main.rs

```

use std::rc::Rc;
use std::cell::RefCell;
use List::{Cons, Nil};

#[derive(Debug)]
enum List {
    Cons(i32, RefCell<Rc<List>>),
    Nil,
}

impl List {
    fn tail(&self) -> Option<&RefCell<Rc<List>>> {
        match self {
            Cons(_, item) => Some(item),
            Nil => None,
        }
    }
}

```

Listing 15-25: A cons list definition that holds a `Cons` variant is referring to

We're using another variation of the `List` definition. The element in the `Cons` variant is now `RefCell<Rc<List>>`, having the ability to modify the `i32` value as we modify which `List` value a `Cons` variant is pointing to. This method makes it convenient for us to access the `tail` variant.

In Listing 15-26, we're adding a `main` function to Listing 15-25. This code creates a list in `a` and a list in `b`, then modifies the list in `a` to point to `b`, creating a recursive structure. The following statements along the way show what the reference process looks like.

Filename: src/main.rs

```

fn main() {
    let a = Rc::new(Cons(5, RefCell::new(
        Nil

    println!("a initial rc count = {}", Rc::count_of(&a));
    println!("a next item = {:?}", a.tail());

    let b = Rc::new(Cons(10, RefCell::new(
        Nil

    println!("a rc count after b creation = {}", Rc::count_of(&a));
    println!("b initial rc count = {}", Rc::count_of(&b));
    println!("b next item = {:?}", b.tail());

    if let Some(link) = a.tail() {
        *link.borrow_mut() = Rc::clone(&b);
    }

    println!("b rc count after changing a = {}", Rc::count_of(&b));
    println!("a rc count after changing a = {}", Rc::count_of(&a));

    // Uncomment the next line to see that this causes a stack overflow
    // println!("a next item = {:?}", a.tail());
}

```

Listing 15-26: Creating a reference cycle of two

We create an `Rc<List>` instance holding a `List` initial list of `5, Nil`. We then create an `Rc<List>` value in the variable `b` that contains the value 1

We modify `a` so it points to `b` instead of `Nil`, using the `tail` method to get a reference to the `RefCell` inside the variable `link`. Then we use the `borrow_mut` to change the value inside from an `Rc<List>` to `b` in `b`.

When we run this code, keeping the last `println!` we'll get this output:

```

a initial rc count = 1
a next item = Some(RefCell { value: Nil })
a rc count after b creation = 2
b initial rc count = 1
b next item = Some(RefCell { value: Cons(5, RefCell { value: Nil }) })
b rc count after changing a = 2
a rc count after changing a = 2

```

The reference count of the `Rc<List>` instances change the list in `a` to point to `b`. At the end of

which will decrease the count in each of the `RefCell`

However, because `a` is still referencing the `RefCell`, it has a count of 1 rather than 0, so the memory cannot be dropped. The memory will just sit there with the reference cycle, we've created a diagram in Figure 15-4.

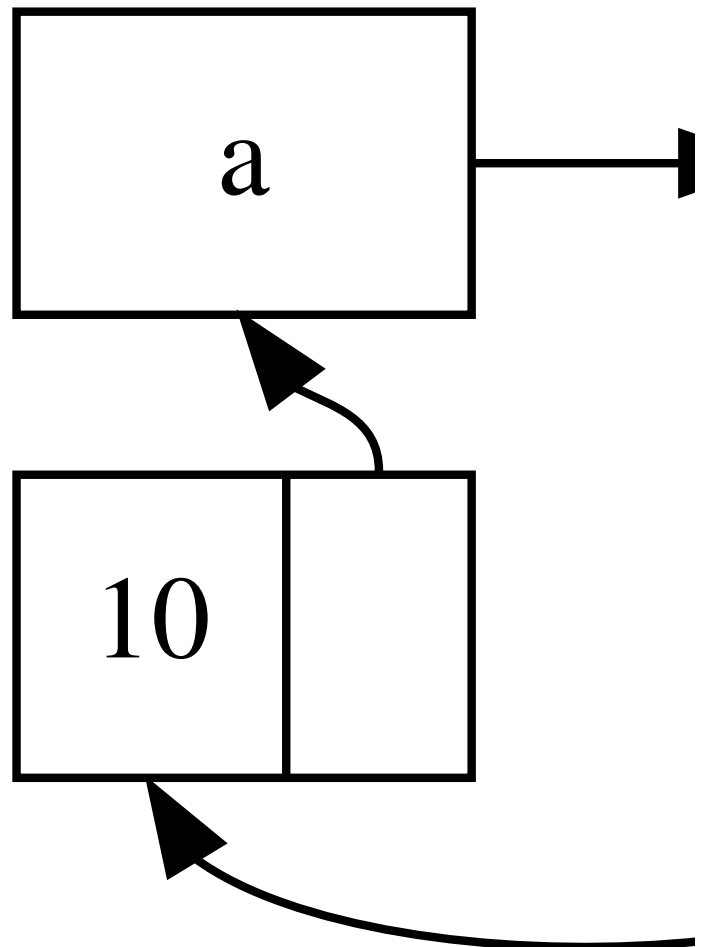


Figure 15-4: A reference cycle of lists `a` and `b`

If you uncomment the last `println!` and run the program, you will see a reference cycle with `a` pointing to `b` pointing to `a` and so on.

In this case, right after we create the reference cycle, the consequences of this cycle aren't very dire. However, if you allocate lots of memory in a cycle and hold onto it, you would use more memory than it needed and might eventually run out of available memory.

Creating reference cycles is not easily done, but `RefCell` values that contain `RefCell` values of types with interior mutability and reference counting can create such cycles.

create cycles; you can't rely on Rust to catch them. It can be a logic bug in your program that you should identify and other software development practices to minimize.

Another solution for avoiding reference cycles is that some references express ownership and so can't have cycles made up of some ownership relationships, and only the ownership relationships can be dropped. In Listing 15-25, we always want to reorganizing the data structure isn't possible. Let's make it made up of parent nodes and child nodes to see if it's an appropriate way to prevent reference cycles.

## Preventing Reference Cycles: Turning Strong References into Weak References

So far, we've demonstrated that calling `Rc::clone` on an `Rc<T>` instance, and an `Rc<T>` instance is only created when the value is cloned. You can also create a *weak reference* to the value by calling `Rc::downgrade` and passing a reference to the `Rc<T>` instance, you get a smart pointer of type `Weak<T>`. Instead of calling `Rc::clone` on the `Rc<T>` instance by 1, calling `Rc::downgrade` on the `Rc<T>` instance uses `weak_count` to keep track of how many weak references to the `Rc<T>` type uses `strong_count` to keep track of how many strong references to the `Rc<T>` type uses. The difference is the `weak_count` is the number of weak references to the `Rc<T>` instance to be cleaned up.

Strong references are how you can share ownership. Weak references don't express an ownership relationship because any cycle involving some weak references will be cleaned up. The strong reference count of values involved is 0.

Because the value that `Weak<T>` references might be cleaned up with the value that a `Weak<T>` is pointing to, you can't dereference it. Do this by calling the `upgrade` method on a `Weak<T>` instance, returning an `Option<Rc<T>>`. You'll get a result of `Some` if the value is still valid and a result of `None` if the `Rc<T>` value has been cleaned up. If you call `upgrade` on a `Weak<T>` instance, Rust will ensure that the value is handled, and there won't be an invalid pointer.

As an example, rather than using a list whose items know about their parent, we'll create a tree whose items know about their children.



## Creating a Tree Data Structure: a `Node` with `children`

To start, we'll build a tree with nodes that know their own value. We'll create a struct named `Node` that holds its own `i32` value and a `Vec` of `Node` values:

Filename: src/main.rs

```
use std::rc::Rc;
use std::cell::RefCell;

#[derive(Debug)]
struct Node {
    value: i32,
    children: RefCell<Vec<Rc<Node>>>,
}
```

We want a `Node` to own its children, and we want to store references to the child nodes in the `children` field. We'll use `Vec<T>` items to be values of type `Rc<Node>`. We want to store references to the child nodes, so we have a `RefCell<Vec<Rc<Node>>>`.

Next, we'll use our struct definition and create one instance of a `Node` with the value 3 and no children, and another instance of a `Node` with the value 5 and `leaf` as one of its children, as shown in Listing 15-27.

Filename: src/main.rs

```
fn main() {
    let leaf = Rc::new(Node {
        value: 3,
        children: RefCell::new(vec![]),
    });

    let branch = Rc::new(Node {
        value: 5,
        children: RefCell::new(vec![Rc::clone(&leaf)]),
    });
}
```

Listing 15-27: Creating a `leaf` node with no children and a `branch` node with `leaf` as one of its children

We clone the `Rc<Node>` in `leaf` and store that in the `children` field of `branch`. Now `leaf` has two owners: `leaf` and `branch`.

through `branch.children`, but there's no way to reason is that `leaf` has no reference to `branch` want `leaf` to know that `branch` is its parent. We

## Adding a Reference from a Child to Its Parent

To make the child node aware of its parent, we r `Node` struct definition. The trouble is in deciding We know it can't contain an `Rc<T>`, because tha `leaf.parent` pointing to `branch` and `branch.c` would cause their `strong_count` values to neve

Thinking about the relationships another way, a if a parent node is dropped, its child nodes shou child should not own its parent: if we drop a chil This is a case for weak references!

So instead of `Rc<T>`, we'll make the type of `par` `RefCell<Weak<Node>>`. Now our `Node` struct de

Filename: src/main.rs

```
use std::rc::{Rc, Weak};
use std::cell::RefCell;

#[derive(Debug)]
struct Node {
    value: i32,
    parent: RefCell<Weak<Node>>,
    children: RefCell<Vec<Rc<Node>>>,
}
```

A node will be able to refer to its parent node bu 15-28, we update `main` to use this new definitio refer to its parent, `branch`:

Filename: src/main.rs

```

fn main() {
    let leaf = Rc::new(Node {
        value: 3,
        parent: RefCell::new(Weak::new()),
        children: RefCell::new(vec![]),
    });

    println!("leaf parent = {:?}", leaf.parent.borrow());

    let branch = Rc::new(Node {
        value: 5,
        parent: RefCell::new(Weak::new()),
        children: RefCell::new(vec![Rc::clone(&leaf)]),
    });

    *leaf.parent.borrow_mut() = Rc::downgrade(&branch);

    println!("leaf parent = {:?}", leaf.parent.borrow());
}

```

Listing 15-28: A `leaf` node with a weak reference to its parent

Creating the `leaf` node looks similar to how we created the `branch` node in Listing 15-27 with the exception of the `parent` field: `leaf` creates a new, empty `Weak<Node>` reference instead of a `Rc<Node>`.

At this point, when we try to get a reference to the parent of `leaf` using the `upgrade` method, we get a `None` value. We see this in the `println!` statement:

```
leaf.parent = None
```

When we create the `branch` node, it will also have a `parent` field, because `branch` doesn't have any children yet. Once we have the `leaf` node, we can modify `leaf` to give it a `Weak<Node>` reference to its parent by using the `upgrade` method on the `RefCell<Weak<Node>>` in the `parent` field. We use the `Rc::downgrade` function to create a `Weak<Node>` from a `Rc<Node>` in `branch`.

When we print the parent of `leaf` again, this time we get `Some(Weak)`: now `leaf` can access its parent! When we print the parent of `branch`, we see a cycle that eventually ended in a stack overflow like in Listing 15-29: `Weak<Node>` references are printed as `(Weak)`:

```
leaf parent = Some(Node { value: 5, parent
children: RefCell { value: [Node { value:
(Weak) },
children: RefCell { value: [] } }] } })
```

The lack of infinite output indicates that this code can also tell this by looking at the values we get

`Rc::weak_count`.

## Visualizing Changes to `strong_count` and `weak_c`

Let's look at how the `strong_count` and `weak_c` instances change by creating a new inner scope into that scope. By doing so, we can see what has then dropped when it goes out of scope. The code is 15-29:

Filename: src/main.rs

```

fn main() {
    let leaf = Rc::new(Node {
        value: 3,
        parent: RefCell::new(Weak::new()),
        children: RefCell::new(vec![]),
    });

    println!(
        "leaf strong = {}, weak = {}",
        Rc::strong_count(&leaf),
        Rc::weak_count(&leaf),
    );

    {
        let branch = Rc::new(Node {
            value: 5,
            parent: RefCell::new(Weak::new()),
            children: RefCell::new(vec![Rc::clone(&leaf)]),
        });

        *leaf.parent.borrow_mut() = Rc::clone(&branch);

        println!(
            "branch strong = {}, weak = {}",
            Rc::strong_count(&branch),
            Rc::weak_count(&branch),
        );

        println!(
            "leaf strong = {}, weak = {}",
            Rc::strong_count(&leaf),
            Rc::weak_count(&leaf),
        );
    }

    println!("leaf parent = {:?}", leaf.parent.borrow().weak().unwrap());
    println!(
        "leaf strong = {}, weak = {}",
        Rc::strong_count(&leaf),
        Rc::weak_count(&leaf),
    );
}

```

Listing 15-29: Creating `branch` in an inner scope and checking reference counts

After `leaf` is created, its `Rc<Node>` has a strong reference count of 1. In the inner scope, we create `branch` and associate its `parent` with `leaf`. When we print the counts, the `Rc<Node>` in `branch` will have a count of 1 (for `leaf.parent` pointing to `branch`).

the counts in `leaf`, we'll see it will have a strong a clone of the `Rc<Node>` of `leaf` stored in `branch` with a weak count of 0.

When the inner scope ends, `branch` goes out of scope. The `Rc<Node>` count decreases to 0, so its `Node` is dropped. `leaf.parent` has no bearing on whether or not there are memory leaks!

If we try to access the parent of `leaf` after the end of the scope, we get a panic. At the end of the program, the `Rc<Node>` in `leaf` has a count of 0, because the variable `leaf` is now `None` again.

All of the logic that manages the counts and values for `Weak<T>` and their implementations of the `Drop` trait to break the relationship from a child to its parent should be in `Node`. If you're able to have parent nodes point to children, you're creating a reference cycle and memory leaks.

## Summary

This chapter covered how to use smart pointers and the trade-offs that Rust makes by default with them. `RefCell<T>` has a known size and points to data allocated on the heap and keeps track of the number of references to data on the heap owned by its owners. The `RefCell<T>` type with its interior mutability can be used when we need an immutable type but need to mutate it. It also enforces the borrowing rules at runtime if you violate them.

Also discussed were the `Deref` and `Drop` traits, which are part of the functionality of smart pointers. We explored reference cycles and how to prevent them using `Weak<T>`.

If this chapter has piqued your interest and you want to learn more about smart pointers, check out ["The Rustonomicon"](#) for more details.

Next, we'll talk about concurrency in Rust. You'll learn how to use smart pointers.

## Fearless Concurrency

Handling concurrent programming safely and efficiently are two different goals. *Concurrent programming*, where different parts of a program run independently, and *parallel programming*, where different parts of a program run at the same time, are becoming increasingly important due to the advantage of their multiple processors. Historically, concurrent programming has been difficult and error prone: Rust hopes to change that.

Initially, the Rust team thought that ensuring memory safety and solving concurrency problems were two separate challenges that required different methods. Over time, the team discovered that there was a powerful set of tools to help manage memory safety and concurrency: leveraging ownership and type checking, many concurrency errors in Rust rather than runtime errors. There's a lot of time trying to reproduce the exact circumstances under which a concurrency bug occurs, incorrect code will refuse to compile rather than explaining the problem. As a result, you can fix your code before it is shipped rather than potentially after it has been shipped. This is one of the key aspects of Rust's *fearless concurrency*. Fearless concurrency means that Rust is free of subtle bugs and is easy to refactor with confidence.

---

Note: For simplicity's sake, we'll refer to many of the concepts in this chapter rather than being more precise by saying *concurrent* or *parallel*. We were about concurrency and/or parallelism, but in this chapter, please mentally substitute *concurrent* for *parallel* and *concurrent* for *parallel*.

---

Many languages are dogmatic about the solution to concurrency problems. For example, Erlang has elegant functional programming for concurrency but has only obscure ways to share state. Rust offers only a subset of possible solutions is a reasonable compromise because a higher-level language promises benefits that outweigh the loss of performance gain abstractions. However, lower-level languages offer a solution with the best performance in any given situation. Therefore, Rust offers a variety of solutions, whatever way is appropriate for your situation and your hardware.

Here are the topics we'll cover in this chapter:

- How to create threads to run multiple pieces of code
- *Message-passing* concurrency, where channels are used to communicate
- *Shared-state* concurrency, where multiple threads share data
- The `Sync` and `Send` traits, which extend Rust's concurrency model

defined types as well as types provided by

## Using Threads to Run Code Sir

In most current operating systems, an executed and the operating system manages multiple pro you can also have independent parts that run si these independent parts are called *threads*.

Splitting the computation in your program into r performance because the program does multipl adds complexity. Because threads can run simul guarantee about the order in which parts of you This can lead to problems, such as:

- Race conditions, where threads are accessi inconsistent order
- Deadlocks, where two threads are waiting resource the other thread has, preventing
- Bugs that happen only in certain situations reliably

Rust attempts to mitigate the negative effects of multithreaded context still takes careful thought different from that in programs running in a sing

Programming languages implement threads in a systems provide an API for creating new threads the operating system APIs to create threads is so operating system thread per one language threa

Many programming languages provide their own Programming language-provided threads are kn that use these green threads will execute them i operating system threads. For this reason, the g model: there are **M** green threads per **N** operat are not necessarily the same number.

Each model has its own advantages and trade-o to Rust is runtime support. *Runtime* is a confusir meanings in different contexts.

In this context, by *runtime* we mean code that is



binary. This code can be large or small depending on the target assembly language. Rust will have some amount of runtime overhead colloquially when people say a language has “no runtime.” Smaller runtimes have fewer features, which make it easier to combine them in more contexts. Although many languages are in exchange for more features, Rust needs to have a compromise on being able to call into C to maintain

The green-threading M:N model requires a large number of threads. As such, the Rust standard library only implements green-threading. Because Rust is such a low-level language, it doesn't support M:N threading if you would rather trade overhead for lower costs of

Now that we've defined threads in Rust, let's explore the API provided by the standard library.

## Creating a New Thread with `spawn`

To create a new thread, we call the `thread::spawn` function (we talked about closures in Chapter 13) containing a closure that runs on the new thread. The example in Listing 16-1 prints some text from a new thread:

Filename: src/main.rs

```
use std::thread;
use std::time::Duration;

fn main() {
    thread::spawn(|| {
        for i in 1..10 {
            println!("hi number {} from the spawned thread!", i);
            thread::sleep(Duration::from_millis(250));
        }
    });

    for i in 1..5 {
        println!("hi number {} from the main thread!", i);
        thread::sleep(Duration::from_millis(250));
    }
}
```

Listing 16-1: Creating a new thread to print one thing and something else

Note that with this function, the new thread will end, whether or not it has finished running. This is a little different every time, but it will look similar

```
hi number 1 from the main thread!  
hi number 1 from the spawned thread!  
hi number 2 from the main thread!  
hi number 2 from the spawned thread!  
hi number 3 from the main thread!  
hi number 3 from the spawned thread!  
hi number 4 from the main thread!  
hi number 4 from the spawned thread!  
hi number 5 from the spawned thread!
```

The calls to `thread::sleep` force a thread to stop, allowing a different thread to run. The threads are not guaranteed: it depends on how your operating system runs. In this example, the main thread printed first, even though the spawned thread appears first in the code. And even though the spawned thread prints until `i` is 9, it only got to 5 before the main thread finished.

If you run this code and only see output from the spawned thread, try increasing the numbers in the range. The operating system will switch between the threads.

## Waiting for All Threads to Finish Using

The code in Listing 16-1 not only stops the spawned thread due to the main thread ending, but also guarantees that the spawned thread will get to run at all. The reason is that the spawned thread can run while the main thread is running.

We can fix the problem of the spawned thread not running completely, by saving the return value of `thread::spawn`. The return type of `thread::spawn` is `JoinHandle`. A `JoinHandle` is a handle to the spawned thread. If we call the `join` method on it, we will wait for its thread to finish. To use the `JoinHandle` of the thread we created, we need to use the `join` method. This ensures the spawned thread finishes before the `main` thread ends.

Filename: src/main.rs

```

use std::thread;
use std::time::Duration;

fn main() {
    let handle = thread::spawn(|| {
        for i in 1..10 {
            println!("hi number {} from the spawned thread!", i);
            thread::sleep(Duration::from_millis(250));
        }
    });

    for i in 1..5 {
        println!("hi number {} from the main thread!", i);
        thread::sleep(Duration::from_millis(250));
    }

    handle.join().unwrap();
}

```

Listing 16-2: Saving a `JoinHandle` from `thread::spawn` to run to completion

Calling `join` on the handle blocks the thread until the thread represented by the handle terminates. *Blocking* means the thread is prevented from performing work or exiting. Because of this, the main thread's `for` loop, running Listing 16-2, prints this:

```

hi number 1 from the main thread!
hi number 2 from the main thread!
hi number 1 from the spawned thread!
hi number 3 from the main thread!
hi number 2 from the spawned thread!
hi number 4 from the main thread!
hi number 3 from the spawned thread!
hi number 4 from the spawned thread!
hi number 5 from the spawned thread!
hi number 6 from the spawned thread!
hi number 7 from the spawned thread!
hi number 8 from the spawned thread!
hi number 9 from the spawned thread!

```

The two threads continue alternating, but the main thread's `handle.join()` and does not end until the spawned thread terminates.

But let's see what happens when we instead move the `for` loop in `main`, like this:

Filename: `src/main.rs`

```

use std::thread;
use std::time::Duration;

fn main() {
    let handle = thread::spawn(|| {
        for i in 1..10 {
            println!("hi number {} from the spawned thread!", i);
            thread::sleep(Duration::from_millis(250));
        }
    });

    handle.join().unwrap();

    for i in 1..5 {
        println!("hi number {} from the main thread!", i);
        thread::sleep(Duration::from_millis(250));
    }
}

```

The main thread will wait for the spawned thread to finish before continuing, so the output won't be interleaved anymore, as follows:

```

hi number 1 from the spawned thread!
hi number 2 from the spawned thread!
hi number 3 from the spawned thread!
hi number 4 from the spawned thread!
hi number 5 from the spawned thread!
hi number 6 from the spawned thread!
hi number 7 from the spawned thread!
hi number 8 from the spawned thread!
hi number 9 from the spawned thread!
hi number 1 from the main thread!
hi number 2 from the main thread!
hi number 3 from the main thread!
hi number 4 from the main thread!

```

Small details, such as where `join` is called, can make a big difference in how threads run at the same time.

## Using `move` Closures with Threads

The `move` closure is often used alongside `thread::spawn` to ensure that a closure can use data from one thread in another thread.

In Chapter 13, we mentioned we can use the `move` keyword with a closure to force the closure to take ownership of its environment. This technique is especially useful when you want to pass data from one thread to another.

to transfer ownership of values from one thread

Notice in Listing 16-1 that the closure we pass to `thread::spawn` we're not using any data from the main thread in the spawned thread to capture the values it needs. Listing 16-3 shows a main thread and use it in the spawned thread. I see in a moment.

Filename: src/main.rs

```
use std::thread;

fn main() {
    let v = vec![1, 2, 3];

    let handle = thread::spawn(|| {
        println!("Here's a vector: {:?}",
    });

    handle.join().unwrap();
}
```

Listing 16-3: Attempting to use a vector created

The closure uses `v`, so it will capture `v` and move it to the new environment. Because `thread::spawn` runs this thread, it will be able to access `v` inside that new thread. But it will get the following error:

```
error[E0373]: closure may outlive the current function
   |
   | `v`,
   | which is owned by the current function
   | --> src/main.rs:6:32
   |
6 |         let handle = thread::spawn(|| {
   |                                ^^ may
7 |             println!("Here's a vector: {:?}",
   |
   | help: to force the closure to take ownership of the
   |       referenced variables), use the `move` keyword
   |
6 |         let handle = thread::spawn(move || {
   |                                     ^^^^^^^^
```

Rust *infers* how to capture `v`, and because `println!` takes a reference, the closure tries to borrow `v`. However, there's

the spawned thread will run, so it doesn't know `v` is no longer valid.

Listing 16-4 provides a scenario that's more likely to be valid:

Filename: src/main.rs

```
use std::thread;

fn main() {
    let v = vec![1, 2, 3];

    let handle = thread::spawn(|| {
        println!("Here's a vector: {:?}", v);
    });

    drop(v); // oh no!

    handle.join().unwrap();
}
```

Listing 16-4: A thread with a closure that attempts to use a variable in the main thread that drops `v`

If we were allowed to run this code, there's a potential race condition: the spawned thread immediately puts `v` in the background without running the garbage collector. Meanwhile, the main thread immediately drops its reference to `v` inside, but the main thread immediately continues to execute the function we discussed in Chapter 15. Then, when the spawned thread executes, `v` is no longer valid, so a reference to it is invalid.

To fix the compiler error in Listing 16-3, we can use the `move` keyword.

```
help: to force the closure to take ownership of
      referenced
      (mutated)
      variables), use the `move` keyword
6 |     let handle = thread::spawn(move || {
  |                               ^^^^^^^
```

By adding the `move` keyword before the closure, we transfer ownership of the values it's using rather than allowing it to borrow the values. The modification to Listing 16-4 looks like this and runs as we intend:

Filename: src/main.rs

```

use std::thread;

fn main() {
    let v = vec![1, 2, 3];

    let handle = thread::spawn(move || {
        println!("Here's a vector: {:?}",
    });

    handle.join().unwrap();
}

```

Listing 16-5: Using the `move` keyword to force a values it uses

What would happen to the code in Listing 16-4 if we use a `move` closure? Would `move` fix that case? It would cause a different error because what Listing 16-4 is trying to do is move the value out of the environment. If we added `move` to the closure, we would move the value out of the environment, and we could no longer call `drop` on it. We get this compiler error instead:

```

error[E0382]: use of moved value: `v`
--> src/main.rs:10:10
   |
6  |         let handle = thread::spawn(move || {
   |         -----
here
...
10 |         drop(v); // oh no!
   |         ^ value used here after move
   |
= note: move occurs because `v` has type `Vec<u32>` which does
       not implement the `Copy` trait

```

Rust's ownership rules have saved us again! We saw in Listing 16-3 because Rust was being conservative and conserving memory, which meant the main thread could theoretically still have a reference. By telling Rust to move ownership of the vector, we're guaranteeing Rust that the main thread won't use it. In Listing 16-4 in the same way, we're then violating the ownership rules by using it in the main thread. The `move` keyword overrides borrowing; it doesn't let us violate the ownership rules.

With a basic understanding of threads and the thread pool, we can move on to threads with threads.

# Using Message Passing to Traverse Threads

One increasingly popular approach to ensuring safety is message passing, where threads or actors communicate by sending data. Here's the idea in a slogan from [the Go language](#): communicate by sharing memory; instead, share communication.

One major tool Rust has for accomplishing message passing is a *channel*, a programming concept that Rust's standard library has a good implementation of. You can imagine a channel is like a channel of water, such as a stream or a river. If you throw a rubber duck or boat into a stream, it will travel downstream to the next person.

A channel in programming has two halves: a transmitter half is the upstream location where you put the rubber duck, and the receiver half is where the rubber duck ends up. The code calls methods on the transmitter with the code on the receiver half checks the receiving end for arriving messages. If either the transmitter or receiver half is dropped, the channel is closed.

Here, we'll work up to a program that has one thread that sends values down a channel, and another thread that receives them. We'll be sending simple values between threads. This is a useful feature. Once you're familiar with the technique, you can use it to implement a chat system or a system where multiple threads do a calculation and send the parts to one thread that combines them.

First, in Listing 16-6, we'll create a channel but not use it yet. It won't compile yet because Rust can't tell what type of data the channel carries.

Filename: src/main.rs

```
use std::sync::mpsc;

fn main() {
    let (tx, rx) = mpsc::channel();
}
```

Listing 16-6: Creating a channel and assigning the halves

We create a new channel using the `mpsc::channel()` function, which stands for *multi-producer, single consumer*. In short, the way Rust implements channels means a channel can have multiple senders but only one receiver.



only one *receiving* end that consumes those values together into one big river: everything sent down one river at the end. We'll start with a single producer when we get this example working.

The `mpsc::channel` function returns a tuple, the first element is the transmitting end and the second element is the receiving end. This is traditionally used in many fields for *transmitter* and *receiver* variables as such to indicate each end. We're using a pattern that deconstructs the tuple; we'll discuss pattern matching and destructuring in Chapter 18. Using `let (tx, rx) =` is a convenient approach to extract the pieces of the tuple.

Let's move the transmitting end into a spawned thread. The spawned thread is communicating with the receiver. This is like putting a rubber duck in the river upstream from one thread to another.

Filename: src/main.rs

```
use std::thread;
use std::sync::mpsc;

fn main() {
    let (tx, rx) = mpsc::channel();

    thread::spawn(move || {
        let val = String::from("hi");
        tx.send(val).unwrap();
    });
}
```

Listing 16-7: Moving `tx` to a spawned thread and using `unwrap`

Again, we're using `thread::spawn` to create a new thread. We move `tx` into the closure so the spawned thread needs to own the transmitting end of the channel through the channel.

The transmitting end has a `send` method that takes a value. The `send` method returns a `Result<T, E>` type, so if the value is dropped and there's nowhere to send a value, then it returns an error. In this example, we're calling `unwrap` to panic in case of an error. In a real application, we would handle it properly: return proper error handling.

In Listing 16-8, we'll get the value from the receiving thread. This is like retrieving the rubber duck from like getting a chat message.

Filename: src/main.rs

```
use std::thread;
use std::sync::mpsc;

fn main() {
    let (tx, rx) = mpsc::channel();

    thread::spawn(move || {
        let val = String::from("hi");
        tx.send(val).unwrap();
    });

    let received = rx.recv().unwrap();
    println!("Got: {}", received);
}
```

Listing 16-8: Receiving the value “hi” in the main

The receiving end of a channel has two useful methods: `recv`, short for *receive*, which will block the thread until a value is sent down the channel. Once a value is received, it returns a `Result<T, E>`. When the sending end of the channel reaches an error to signal that no more values will be coming.

The `try_recv` method doesn't block, but will instead return immediately: an `Ok` value holding a message if one is available, or an `Err` if there aren't any messages this time. Using `try_recv` in a loop to do work while waiting for messages: we could check every so often, handles a message if one is available, and then check again for a little while until checking again.

We've used `recv` in this example for simplicity; in a real application, the main thread to do other than wait for messages would be more appropriate.

When we run the code in Listing 16-8, we'll see the following output from the thread:

```
Got: hi
```

Perfect!

## Channels and Ownership Transference

The ownership rules play a vital role in message safe, concurrent code. Preventing errors in conc of thinking about ownership throughout your Ru to show how channels and ownership work toge use a `val` value in the spawned thread *after* we compiling the code in Listing 16-9 to see why thi

Filename: src/main.rs

```
use std::thread;
use std::sync::mpsc;

fn main() {
    let (tx, rx) = mpsc::channel();

    thread::spawn(move || {
        let val = String::from("hi");
        tx.send(val).unwrap();
        println!("val is {}", val);
    });

    let received = rx.recv().unwrap();
    println!("Got: {}", received);
}
```

Listing 16-9: Attempting to use `val` after we've s

Here, we try to print `val` after we've sent it down this would be a bad idea: once the value has been could modify or drop it before we try to use the thread's modifications could cause errors or unexpected or nonexistent data. However, Rust gives us an Listing 16-9:

```
error[E0382]: use of moved value: `val`
--> src/main.rs:10:31
   |
9  |         tx.send(val).unwrap();
   |         --- value moved here
10 |         println!("val is {}", val);
   |                                   ^^^ val
   |
   = note: move occurs because `val` has type `String` which does
         not implement the `Copy` trait
```

Our concurrency mistake has caused a compile-time error: the ownership of its parameter, and when the value is moved, the ownership of it. This stops us from accidentally moving the ownership system checks that everything is

## Sending Multiple Values and Seeing the Results

The code in Listing 16-8 compiled and ran, but it was slow. To speed it up, we made some modifications that will prove the code works concurrently: the spawned thread will now send a value every second between each message.

Filename: src/main.rs

```
use std::thread;
use std::sync::mpsc;
use std::time::Duration;

fn main() {
    let (tx, rx) = mpsc::channel();

    thread::spawn(move || {
        let vals = vec![
            String::from("hi"),
            String::from("from"),
            String::from("the"),
            String::from("thread"),
        ];

        for val in vals {
            tx.send(val).unwrap();
            thread::sleep(Duration::from_secs(1));
        }
    });

    for received in rx {
        println!("Got: {}", received);
    }
}
```

Listing 16-10: Sending multiple messages and pausing between them

This time, the spawned thread has a vector of strings to send to the main thread. We iterate over them, sending each one by calling the `thread::sleep` function with

In the main thread, we're not calling the `recv` function, so we're treating `rx` as an iterator. For each value received, the channel is closed, iteration will end.

When running the code in Listing 16-10, you should see a 1-second pause in between each line:

```
Got: hi
Got: from
Got: the
Got: thread
```

Because we don't have any code that pauses the main thread, we can tell that the main thread is waiting for the other thread.

## Creating Multiple Producers by Cloning

Earlier we mentioned that `mpsc` was an acronym for *multiple producer, single consumer*. Let's put `mpsc` to use and expand the code to create multiple threads that all send values to the same channel, the transmitting half of the channel, as shown in Listing 16-11.

Filename: `src/main.rs`

```

// --snip--

let (tx, rx) = mpsc::channel();

let tx1 = mpsc::Sender::clone(&tx);
thread::spawn(move || {
    let vals = vec![
        String::from("hi"),
        String::from("from"),
        String::from("the"),
        String::from("thread"),
    ];

    for val in vals {
        tx1.send(val).unwrap();
        thread::sleep(Duration::from_secs(1))
    }
});

thread::spawn(move || {
    let vals = vec![
        String::from("more"),
        String::from("messages"),
        String::from("for"),
        String::from("you"),
    ];

    for val in vals {
        tx.send(val).unwrap();
        thread::sleep(Duration::from_secs(1))
    }
});

for received in rx {
    println!("Got: {}", received);
}

// --snip--

```

Listing 16-11: Sending multiple messages from r

This time, before we create the first spawned thread, we clone the original sender to the other end of the channel. This will give us a new sender to use in the first spawned thread. We pass the original sender to the second spawned thread. This gives us two threads, each with its own sending end of the channel.

When you run the code, your output should look

```
Got: hi
Got: more
Got: from
Got: messages
Got: for
Got: the
Got: thread
Got: you
```

You might see the values in another order; it doesn't make concurrency interesting as well as difficult. `thread::sleep`, giving it various values in the duration, is nondeterministic and creates different output each time.

Now that we've looked at how channels work, let's look at concurrency.

## Shared-State Concurrency

Message passing is a fine way of handling concurrency. Consider this part of the slogan from the Go language, "communicate by sharing memory."

What would communicating by sharing memory mean? Message-passing enthusiasts not use it and do it differently.

In a way, channels in any programming language are like pointers because once you transfer a value down a channel, you own the value. Shared memory concurrency is like multiple pointers accessing the same memory location at the same time. Smart pointers made multiple ownership possible, but they add complexity because these different owners need rules. Ownership rules greatly assist in getting this managed. Let's look at mutexes, one of the more common synchronization primitives for shared memory.

## Using Mutexes to Allow Access to Data

*Mutex* is an abbreviation for *mutual exclusion*, as it ensures that only one thread can access some data at any given time. To access the data, a thread signals that it wants access by asking to acquire the mutex. The mutex is a structure that is part of the mutex that keeps track of which thread has the mutex.

access to the data. Therefore, the mutex is described as the locking system.

Mutexes have a reputation for being difficult to use, but they follow two rules:

- You must attempt to acquire the lock before you access the data.
- When you're done with the data that the mutex protects, you must release the data so other threads can acquire the lock.

For a real-world metaphor for a mutex, imagine a panel discussion with only one microphone. Before a panelist can speak, they want to use the microphone. When they get the microphone, they hold it as long as they want to and then hand the microphone to the next person who requests to speak. If a panelist forgets to hand the microphone back, no one else is able to speak. If a panelist finishes with it, no one else is able to speak. If the microphone goes wrong, the panel won't work at all.

Management of mutexes can be incredibly tricky, but many people are enthusiastic about channels. However, if you don't have ownership rules, you can't get locking and unlocking.

## The API of `Mutex<T>`

As an example of how to use a mutex, let's start with a simple context, as shown in Listing 16-12:

Filename: src/main.rs

```
use std::sync::Mutex;

fn main() {
    let m = Mutex::new(5);

    {
        let mut num = m.lock().unwrap();
        *num = 6;
    }

    println!("m = {:?}", m);
}
```

Listing 16-12: Exploring the API of `Mutex<T>` in a simple context

As with many types, we create a `Mutex<T>` using `Mutex::new()`. To access the data inside the mutex, we use the `lock()` call. This call will block the current thread so it can't do anything else until the lock is released.



lock.

The call to `lock` would fail if another thread holds the lock; no one would ever be able to get the lock, so we thread panic if we're in that situation.

After we've acquired the lock, we can treat the `ref` as a mutable reference to the data inside. The type of `lock` before using the value in `m: Mutex<i32>` is `MutexGuard<i32>`. We can't use the `i32` value. We can't access the inner `i32` otherwise.

As you might suspect, `Mutex<T>` is a smart pointer. It returns a smart pointer called `MutexGuard`. This smart pointer points at our inner data; the smart pointer also holds the lock. The smart pointer releases the lock automatically when it goes out of scope at the end of the inner scope in Listing 16-12. As the lock is released, the mutex is no longer blocked. The lock release happens automatically.

After dropping the lock, we can print the mutex value and change the inner `i32` to 6.

## Sharing a `Mutex<T>` Between Multiple Threads

Now, let's try to share a value between multiple threads. We'll have 10 threads and have them each increment a counter from 0 to 10. Note that the next few examples will have those errors to learn more about using `Mutex<T>` correctly. Listing 16-13 has our starting example.

Filename: src/main.rs

```

use std::sync::Mutex;
use std::thread;

fn main() {
    let counter = Mutex::new(0);
    let mut handles = vec![];

    for _ in 0..10 {
        let handle = thread::spawn(move || {
            let mut num = counter.lock().unwrap();

            *num += 1;
        });
        handles.push(handle);
    }

    for handle in handles {
        handle.join().unwrap();
    }

    println!("Result: {}", *counter.lock().unwrap());
}

```

Listing 16-13: Ten threads each increment a counter

We create a `counter` variable to hold an `i32` in Listing 16-12. Next, we create 10 threads by iterating over `0..10` and calling `thread::spawn` and give all the threads the same reference to `counter`. Each thread acquires a lock on the `counter` and then adds 1 to the value in the mutex. When the thread finishes, `num` will go out of scope and release the lock so the next thread can acquire it.

In the main thread, we collect all the join handle calls by pushing them into the `handles` vector. Then we call `join` on each handle to make sure all the threads have finished. Finally, we print the result of the counter.

We hinted that this example wouldn't compile. But it does! The reason is that the `counter` variable is a `Mutex`, which ensures that only one thread can access it at a time.

```

error[E0382]: capture of moved value: `counter`
--> src/main.rs:10:27
   |
9  |         let handle = thread::spawn(move || {
   |                                     --
here
10 |             let mut num = counter.lock().unwrap();
   |                                     ^^^^^^^^^^ `counter` moved here
   |
   = note: move occurs because `counter` has type `Mutex<u32>`, which does not implement the `Copy` trait

error[E0382]: use of moved value: `counter`
--> src/main.rs:21:29
   |
9  |         let handle = thread::spawn(move || {
   |                                     --
here
...
21 |         println!("Result: {}", *counter.lock().unwrap());
   |                                     ^^^^^^^^^^ `counter` moved here
   |
   = note: move occurs because `counter` has type `Mutex<u32>`, which does not implement the `Copy` trait

error: aborting due to 2 previous errors

```

The error message states that the `counter` value was moved when we call `lock`. That description is not allowed!

Let's figure this out by simplifying the program. Instead of a `for` loop, let's just make two threads without a loop. We'll modify the first `for` loop in Listing 16-13 with this code

```

use std::sync::Mutex;
use std::thread;

fn main() {
    let counter = Mutex::new(0);
    let mut handles = vec![];

    let handle = thread::spawn(move || {
        let mut num = counter.lock().unwrap();

        *num += 1;
    });
    handles.push(handle);

    let handle2 = thread::spawn(move || {
        let mut num2 = counter.lock().unwrap();

        *num2 += 1;
    });
    handles.push(handle2);

    for handle in handles {
        handle.join().unwrap();
    }

    println!("Result: {}", *counter.lock())
}

```

We make two threads and change the variable `handle2` and `num2`. When we run the code this following:

```

error[E0382]: capture of moved value: `counter`
--> src/main.rs:16:24
   |
 8 |         let handle = thread::spawn(move || {
   |                                     -----
here
...
16 |             let mut num2 = counter.lock().unwrap();
   |                               ^^^^^^^^^ value captured here
   |
   = note: move occurs because `counter` has type `Mutex<u32>`, which does not implement the `Copy` trait

error[E0382]: use of moved value: `counter`
--> src/main.rs:26:29
   |
 8 |         let handle = thread::spawn(move || {
   |                                     -----
here
...
26 |         println!("Result: {}", *counter.lock().unwrap());
   |                                   ^^^^^^^^^ value captured here
   |
   = note: move occurs because `counter` has type `Mutex<u32>`, which does not implement the `Copy` trait

error: aborting due to 2 previous errors

```

Aha! The first error message indicates that `counter` is moved when we try to call `lock` on it and store the result. Rust is telling us that we can't move ownership of `counter` to a different thread associated with `handle`. This was hard to see earlier because our threads were created in different iterations of the `for` loop. In the next multiple-ownership method we discussed in Chapter 15, we'll use the `Rc` type to create a reference counted value. Let's do the same thing with `Mutex` in Listing 16-14 and wrap the `Mutex<T>` in `Rc<T>` to share ownership to the thread. Now that we've seen the `for` loop, and we'll keep the `move` key

## Multiple Ownership with Multiple Threads

In Chapter 15, we gave a value multiple owners by using `Rc` to create a reference counted value. Let's do the same thing with `Mutex` in Listing 16-14 and wrap the `Mutex<T>` in `Rc<T>` to share ownership to the thread. Now that we've seen the `for` loop, and we'll keep the `move` key

Filename: src/main.rs

```
use std::rc::Rc;
use std::sync::Mutex;
use std::thread;

fn main() {
    let counter = Rc::new(Mutex::new(0));
    let mut handles = vec![];

    for _ in 0..10 {
        let counter = Rc::clone(&counter);
        let handle = thread::spawn(move || {
            let mut num = counter.lock().unwrap();

            *num += 1;
        });
        handles.push(handle);
    }

    for handle in handles {
        handle.join().unwrap();
    }

    println!("Result: {}", *counter.lock().unwrap());
}
```

Listing 16-14: Attempting to use `Rc<T>` to allow `Mutex<T>`

Once again, we compile and get... different error

```
error[E0277]: the trait bound `std::rc::Rc<std::marker::Send>` is not satisfied in `[closure@src/main.rs:11:10: 15:10 counter:std::rc::Rc<std::sync::Mutex>]`
  --> src/main.rs:11:22
   |
11 |         let handle = thread::spawn(move || {
   |                                     ^^^^^^^^^^^^^^^^^^^
   |                                     `std::rc::Rc<std::sync::Mutex<i32>>`
   |                                     cannot be sent between threads safely
   |
   = help: within `[closure@src/main.rs:11:22: 15:10 counter:std::rc::Rc<std::sync::Mutex<i32>>]`
   `std::marker::Send` is not implemented for `std::rc::Rc<std::sync::Mutex<i32>>`
   = note: required because it appears with the type `[closure@src/main.rs:11:36: 15:10 counter:std::rc::Rc<std::sync::Mutex<i32>>]`
   = note: required by `std::thread::spawn`
```

Wow, that error message is very wordy! Here are the first inline error says

``std::rc::Rc<std::sync::Mutex<i32>>`` cannot

The reason for this is in the next important part distilled error message says `the trait bound `Send`` in the next section: it's one of the traits with threads are meant for use in concurrent situations

Unfortunately, `Rc<T>` is not safe to share across threads. It has a reference count, it adds to the count for each clone, and decrements the count when each clone is dropped. But it doesn't make sure that changes to the count can't be interleaved with other operations. This could lead to wrong counts—subtle bugs that could be avoided by using `Arc<T>` instead. `Arc<T>` is like `Rc<T>` but one that makes changes to the reference count atomically.

## Atomic Reference Counting with `Arc<T>`

Fortunately, `Arc<T>` is a type like `Rc<T>` that is safe to share across threads. The *a* stands for *atomic*, meaning it's an *atomic* operation. `Arc` is an additional kind of concurrency primitive that is part of the Rust standard library. For more documentation for `std::sync::Arc`, you just need to know that atomics work like locks and share across threads.

You might then wonder why all primitive types and standard library types aren't implemented to use `Arc<T>` by default. The answer comes with a performance penalty that you only pay when you need it. If you're just performing operations on values that don't need to be shared, they run faster if it doesn't have to enforce the guarantee of atomicity.

Let's return to our example: `Arc<T>` and `Rc<T>`. We can modify our program by changing the `use` line, the call to `new_mutex`, and the call to `new_mutex`. Listing 16-15 will finally compile and run:

Filename: src/main.rs

```

use std::sync::{Mutex, Arc};
use std::thread;

fn main() {
    let counter = Arc::new(Mutex::new(0));
    let mut handles = vec![];

    for _ in 0..10 {
        let counter = Arc::clone(&counter);
        let handle = thread::spawn(move || {
            let mut num = counter.lock().unwrap();

            *num += 1;
        });
        handles.push(handle);
    }

    for handle in handles {
        handle.join().unwrap();
    }

    println!("Result: {}", *counter.lock().unwrap());
}

```

Listing 16-15: Using an `Arc<T>` to wrap the `Mutex<T>` across multiple threads

This code will print the following:

```
Result: 10
```

We did it! We counted from 0 to 10, which may not teach us a lot about `Mutex<T>` and thread safety, but this structure to do more complicated operations than this strategy, you can divide a calculation into increments across threads, and then use a `Mutex<T>` to have each thread work with its part.

## Similarities Between `RefCell<T>` / `Rc<T>`

You might have noticed that `counter` is an immutable reference to the value inside it; this means `Mutex` is like the `Cell` family does. In the same way we used `RefCell` to mutate contents inside an `Rc<T>`, we use an `Arc<T>`.



Another detail to note is that Rust can't protect you if you use `Mutex<T>`. Recall in Chapter 15 that using `Rc<T>` can create reference cycles, where two `Rc<T>` values leak memory. Similarly, `Mutex<T>` comes with the potential to occur when an operation needs to lock two resources, but has acquired one of the locks, causing them to wait for the other to finish. If you're interested in deadlocks, try creating a Rust program that demonstrates deadlock mitigation strategies for mutexes and implementing them in Rust. The standard library documentation for `MutexGuard` offers useful information.

We'll round out this chapter by talking about the `Weak` type and how you can use them with custom types.

## Extensible Concurrency with Traits

Interestingly, the Rust language has *very* few concurrency features. The only concurrency feature we've talked about so far in this book is the `std::sync::Mutex` in the standard library, not the language. Your options are limited to the language or the standard library; you can't write your own concurrency features or use those written by others.

However, two concurrency concepts are embedded in the language: the `Sync` and `Send` traits.

### Allowing Transference of Ownership Between Threads

The `Send` marker trait indicates that ownership of a value can be transferred between threads. Almost every `Rc<T>` is an exception, including `Rc<T>`: this cannot be `Send`. If you try to transfer ownership of the clone and the original might update the reference count at the same time, it could lead to a data race. `Rc<T>` is implemented for use in single-threaded situations and has a thread-unsafe performance penalty.

Therefore, Rust's type system and trait bounds ensure that you cannot send an `Rc<T>` value across threads unsafely. When we tried to send an `Rc<T>` value across threads in Chapter 16-14, we got the error `the trait Send is not implemented for Rc<T>`. When we switched to `Arc<T>`, which is `Send`, the error went away.

Any type composed entirely of `Send` types is `Send`. Almost all primitive types are `Send`, aside from `UnsafeCell` in Chapter 19.

## Allowing Access from Multiple Threads

The `Sync` marker trait indicates that it is safe to be referenced from multiple threads. In other words, if a reference to `T` is `Send`, meaning the reference is `Send`, primitive types are `Sync`, and types that are `Send` are also `Sync`.

The smart pointer `Rc<T>` is also not `Sync` for thread safety. The `RefCell<T>` type (which we talked about in Chapter 19) is not `Sync`. The implementation of `RefCell<T>` does at runtime is not thread-safe. `RefCell<T>` and can be used to share access with multiple threads using `Mutex<T>`. See the “Between Multiple Threads” section.

## Implementing `Send` and `Sync` Manually

Because types that are made up of `Send` and `Sync` types are `Send` and `Sync`, we don't have to implement those traits for them. They don't even have any methods to implement. The only invariants related to concurrency.

Manually implementing these traits involves implementing `unsafe` Rust code in Chapter 19. It is that building new concurrent types not made with careful thought to uphold the safety guarantees and information about these guarantees and how to

## Summary

This isn't the last you'll see of concurrency in this book. We'll use the concepts in this chapter in a more realistic examples discussed here.

As mentioned earlier, because very little of how

the language, many concurrency solutions are in more quickly than the standard library, so be sure to use state-of-the-art crates to use in multithreaded situations.

The Rust standard library provides channels for concurrent types, such as `Mutex<T>` and `Arc<T>`, that are safe. The type system and the borrow checker ensure that you won't end up with data races or invalid references. On this point, you can rest assured that it will happily run on multiplatform targets. To track-down bugs common in other languages, Rust has a longer a concept to be afraid of: go forth and make things work fearlessly!

Next, we'll talk about idiomatic ways to model programs. As your Rust programs get bigger. In addition, we'll discuss those you might be familiar with from object-oriented programming.

## Object Oriented Programming of Rust

Object-oriented programming (OOP) is a way of programming that came from Simula in the 1960s. Those objects influenced the design of the architecture in which objects pass messages to each other. In 1967, *object-oriented programming* was used to describe this architecture. Some definitions describe what OOP is; some definitions describe what is not. In this chapter, we'll look at characteristics that are commonly considered object-oriented. We'll see how these characteristics translate to idiomatic Rust. We'll look at the object-oriented design pattern in Rust and discuss implementing a solution using some of Rust's standard library.

### Characteristics of Object-Oriented Programming

There is no consensus in the programming community as to what a language must have to be considered object-oriented. Some programming paradigms, including OOP; for example, came from functional programming in Chapter 7. However, certain common characteristics, namely objects, are shared. We'll look at what each of those characteristics means.

## Objects Contain Data and Behavior

The book *Design Patterns: Elements of Reusable Object-Oriented Software* (Addison-Wesley Professional, 1994) colloquially referred to as *The Gang of Four* (GoF) defines object-oriented design patterns. It defines OOP as:

---

Object-oriented programs are made up of objects and the procedures that operate on that data called *methods* or *operations*.

---

Using this definition, Rust is object oriented: structs provide methods on structs and enums. If methods aren't *called* objects, they provide the same as the Gang of Four's definition of objects.

## Encapsulation that Hides Implementation

Another aspect commonly associated with OOP means that the implementation details of an object are hidden from the code that uses that object. Therefore, the only way to interact with an object is through its public interface. Code using the object shouldn't be able to reach its internal data or behavior directly. This enables the programmer to change the object's internals without needing to change the code that uses it.

We discussed how to control encapsulation in C++ using the `public` keyword to decide which modules, types, functions, and variables are public, and by default everything else is private. In Rust, the `pub` keyword is used to mark public items. For example, the struct `AveragedCollection` that has a field containing a list of numbers can also have a field that contains the average of those numbers. This means the average doesn't have to be computed every time it is needed. In other words, `AveragedCollection` is more efficient than `Vec`. Listing 17-1 has the definition of the `AveragedCollection` struct.

Filename: `src/lib.rs`

```
pub struct AveragedCollection {  
    list: Vec<i32>,  
    average: f64,  
}
```

Listing 17-1: An `AveragedCollection` struct that average of the items in the collection

The struct is marked `pub` so that other code can remain private. This is important in this case because whenever a value is added or removed from the collection, we do this by implementing `add`, `remove`, and `average` shown in Listing 17-2:

Filename: `src/lib.rs`

```
impl AveragedCollection {
    pub fn add(&mut self, value: i32) {
        self.list.push(value);
        self.update_average();
    }

    pub fn remove(&mut self) -> Option<i32> {
        let result = self.list.pop();
        match result {
            Some(value) => {
                self.update_average();
                Some(value)
            },
            None => None,
        }
    }

    pub fn average(&self) -> f64 {
        self.average
    }

    fn update_average(&mut self) {
        let total: i32 = self.list.iter().sum();
        self.average = total as f64 / self.list.len() as f64;
    }
}
```

Listing 17-2: Implementations of the public methods of `AveragedCollection`

The public methods `add`, `remove`, and `average` are implemented on an instance of `AveragedCollection`. When an item is added to the collection using the `add` method or removed using the `remove` method, the private `update_average` method that handles updating the average is called.

We leave the `list` and `average` fields private so that only the `add` or `remove` items to the `list` field directly; the `average` field is updated by the `update_average` method.

become out of sync when the `list` changes. Then, by adding `average` to the `AveragedCollection` field, allowing external code to read the average.

Because we've encapsulated the implementation of `AveragedCollection`, we can easily change aspects of its behavior in the future. For instance, we could use a `HashSet` instead of a `Vec` as long as the signatures of the `add`, `remove`, and `average` methods remain the same. Code using `AveragedCollection` wouldn't need to be changed. If `average` were public instead, this wouldn't necessarily be the case. If `add` and `remove` were public methods for adding and removing items, so the change if it were modifying `list` directly.

If encapsulation is a required aspect for a language, then Rust meets that requirement. The option to encapsulate code enables encapsulation of implementation details.

## Inheritance as a Type System and as Code

*Inheritance* is a mechanism whereby an object class definition can inherit from a parent object's definition, thus gaining the parent object's data and methods without having to define them again.

If a language must have inheritance to be an object-oriented language, then Rust is not one. There is no way to define a struct that inherits from another struct and has its own method implementations. However, if you're using Rust as part of a larger programming toolbox, you can use other solutions to achieve the benefits of inheritance in the first place.

You choose inheritance for two main reasons. One is to reuse code to implement particular behavior for one type, and the other is to implement that behavior for a different type. You can also use trait implementations instead, which you saw in the previous chapter. A default implementation of the `summarize` method for a `Summary` trait implementing the `Summary` trait would have the same behavior as a parent class without any further code. This is similar to a parent class having a method and an inheriting child class also having that method. We can also override the default implementation of the `summarize` method. We can also implement the `Summary` trait, which is similar to a parent class implementing a method inherited from a parent class.

The other reason to use inheritance relates to the fact that a child class can be used in the same places as the parent type, which means that you can substitute multiple objects of a child class for a single object of a parent class.

share certain characteristics.

---

## Polymorphism

To many people, polymorphism is synonymous with inheritance, a more general concept that refers to code that can be used for many types. For inheritance, those types are generally

Rust instead uses generics to abstract over different types. Bounded bounds to impose constraints on what those types can be. Sometimes called *bounded parametric polymorphism*.

---

Inheritance has recently fallen out of favor as a design pattern in programming languages because it's often at risk of being over-engineered. Subclasses shouldn't always share all code, but will do so with inheritance. This can make a program more complex and introduces the possibility of calling methods on objects that cause errors because the methods don't apply. Modern programming languages will only allow a subclass to inherit from a base class if the flexibility of a program's design.

For these reasons, Rust takes a different approach to inheritance. Let's look at how trait objects enable

## Using Trait Objects that Allow Any Static Types

In Chapter 8, we mentioned that one limitation of enums is that they can only hold elements of only one type. We created a workaround by using a `SpreadsheetCell` enum that had variants to hold different types of data, which meant we could store different types of data in a single enum. This represented a row of cells. This is a perfectly good design if the items are a fixed set of types that we know when

However, sometimes we want our library user to be able to use our library with types that are valid in a particular situation. To show how to do this, we'll create an example graphical user interface (GUI) tool that takes a list of trait objects and calling a `draw` method on each one to draw it to the screen. We'll create a library crate called `gui`

GUI library. This crate might include some types like `TextField`. In addition, `gui` users will want to call `draw` on these types. For instance, one programmer might add a `SelectBox`.

We won't implement a fully fledged GUI library for now, but we can see how the pieces would fit together. At the time of writing this, we have a `Component` trait that all the types other programmers might want to call `draw` on. `Component` needs to keep track of many values of different types, and it has a `draw` method on each of these differently typed values. We need to define what will happen when we call the `draw` method on a `Component`, and a `draw` method is available for us to call.

To do this in a language with inheritance, we might have a `Component` trait that has a method named `draw` on it. The other types, like `SelectBox`, would inherit from `Component` and could each override the `draw` method to define their own behavior. The framework could treat all of the types as if they implemented `draw` on them. But because Rust doesn't have inheritance, we need to structure the `gui` library to allow users to extend it.

## Defining a Trait for Common Behavior

To implement the behavior we want `gui` to have, we need a way to call `draw` on an `object` that will have one method named `draw`. Then we can call `draw` on the `object`. A trait object points to an instance of a type that implements the trait. We specify a trait object by specifying some type, a reference or a `Box<T>` smart pointer, and then the `dyn` keyword. (We'll talk about the reason trait objects exist in section 19 in the section "Dynamically Sized Types & Sizes of a generic or concrete type. Wherever we use a trait object, we ensure at compile time that any value used in the trait object's trait. Consequently, we don't need to know the type of the trait object at runtime.

We've mentioned that in Rust, we refrain from calling trait objects "objects" to distinguish them from other languages' objects. In Rust, the data fields and the behavior in `impl` blocks are separate. In other languages, the data and behavior combined into a single object. However, trait objects *are* more like objects in other languages that they combine data and behavior. But trait objects are not objects in that we can't add data to a trait object. Trait objects are objects in that they combine data and behavior.



objects in other languages: their specific purposes and common behavior.

Listing 17-3 shows how to define a trait named

Filename: src/lib.rs

```
pub trait Draw {  
    fn draw(&self);  
}
```

Listing 17-3: Definition of the `Draw` trait

This syntax should look familiar from our discussion in Chapter 10. Next comes some new syntax: Listing 17-4 defines a `Screen` struct that holds a vector named `components`, which is a trait object; it's a stand-in for any type that implements the `Draw` trait.

Filename: src/lib.rs

```
pub struct Screen {  
    pub components: Vec<Box<dyn Draw>>,  
}
```

Listing 17-4: Definition of the `Screen` struct with a vector of trait objects that implement the `Draw` trait

On the `Screen` struct, we'll define a method named `run` that calls `draw` on each of its `components`, as shown in Listing 17-5.

Filename: src/lib.rs

```
impl Screen {  
    pub fn run(&self) {  
        for component in self.components.iter() {  
            component.draw();  
        }  
    }  
}
```

Listing 17-5: A `run` method on `Screen` that calls `draw` on each component

This works differently than defining a struct that

trait bounds. A generic type parameter can only type at a time, whereas trait objects allow for multiple trait objects at runtime. For example, we could have a generic type and a trait bound as in Listing 17-6:

Filename: src/lib.rs

```
pub struct Screen<T: Draw> {
    pub components: Vec<T>,
}

impl<T> Screen<T>
    where T: Draw {
    pub fn run(&self) {
        for component in self.components.iter() {
            component.draw();
        }
    }
}
```

Listing 17-6: An alternate implementation of the `Screen` struct using generics and trait bounds

This restricts us to a `Screen` instance that has a `Vec` of components, where each component is of type `Draw` or all of type `TextField`. If you'll only ever have one type of component, using generics and trait bounds is preferable because the code is monomorphized at compile time to use the concrete type.

On the other hand, with the method using trait objects, the `Screen` struct can hold a `Vec` that contains a `Box<Button>` as well as a `Box<TextField>`. We'll see how this works, and then we'll talk about the runtime performance implications.

## Implementing the Trait

Now we'll add some types that implement the `Draw` trait. Again, actually implementing a GUI library is a bit more complex, but the `draw` method won't have any useful implementation. For example, the implementation might look like, a `Button` struct with `width`, `height`, and `label`, as shown in Listing 17-7:

Filename: src/lib.rs

```
pub struct Button {
    pub width: u32,
    pub height: u32,
    pub label: String,
}

impl Draw for Button {
    fn draw(&self) {
        // code to actually draw a button
    }
}
```

Listing 17-7: A `Button` struct that implements the

The `width`, `height`, and `label` fields on `Button` components, such as a `TextField` type, that might have a `placeholder` field instead. Each of the types we implement the `Draw` trait but will use different code to draw that particular type, as `Button` has which is beyond the scope of this chapter). The `Button` has an additional `impl` block containing methods related to clicking the button. These kinds of methods won't

If someone using our library decides to implement a `SelectBox` type with `width`, `height`, and `options` fields, they implement the `Draw` trait for `SelectBox` as well, as shown in Listing 17-8

Filename: `src/main.rs`

```
extern crate gui;
use gui::Draw;

struct SelectBox {
    width: u32,
    height: u32,
    options: Vec<String>,
}

impl Draw for SelectBox {
    fn draw(&self) {
        // code to actually draw a select box
    }
}
```

Listing 17-8: Another crate using `gui` and implementing the `Draw` trait for `SelectBox` struct

Our library's user can now write their `main` function. When they create the `Screen` instance, they can add a `SelectBox` to become a trait object. They can then call `draw` on each of the components, which will call `draw` on each of the component's implementation:

Filename: `src/main.rs`

```
use gui::{Screen, Button};

fn main() {
    let screen = Screen {
        components: vec![
            Box::new(SelectBox {
                width: 75,
                height: 10,
                options: vec![
                    String::from("Yes"),
                    String::from("Maybe"),
                    String::from("No")
                ],
            }),
            Box::new(Button {
                width: 50,
                height: 10,
                label: String::from("OK"),
            }),
        ],
    };

    screen.run();
}
```

Listing 17-9: Using trait objects to store values of different types that implement the same trait

When we wrote the library, we didn't know that `Screen` was a `dyn Draw` type, but our `Screen` implementation was able to call `draw` on it because `SelectBox` implements the `Draw` trait and has a `draw` method.

This concept—of being concerned only with the trait that the value implements rather than the value's concrete type—is similar to the concept in dynamically typed languages: if it walks like a duck and quacks like a duck, it's a duck. In the implementation of `run` on `Screen` in Listing 17-9, we don't know what the concrete type of each component is. It could be an instance of a `Button` or a `SelectBox`, but it just needs to implement the `Draw` trait. By specifying `Box<dyn Draw>` as the type of each component, we can store any component that implements the `Draw` trait.

components vector, we've defined `Screen` to ne  
method on.

The advantage of using trait objects and Rust's `dyn` code using duck typing is that we never have to implement a particular method at runtime or worry about getting it implemented but we call it anyway. Rust doesn't implement the traits that the trait objects implement.

For example, Listing 17-10 shows what happens `String` as a component:

Filename: src/main.rs

```
extern crate gui;
use gui::Screen;

fn main() {
    let screen = Screen {
        components: vec![
            Box::new(String::from("Hi")),
        ],
    };

    screen.run();
}
```

Listing 17-10: Attempting to use a type that does

We'll get this error because `String` doesn't imp

```
error[E0277]: the trait bound `std::string`  
satisfied  
    --> src/main.rs:7:13  
      |  
7     |             Box::new(String::from("Hi"  
      |             ^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^  
implemented for `std::string::String`  
      |  
= note: required for the cast to the optional
```

This error lets us know that either we're passing to pass and we should pass a different type or we should pass a `String` so that `Screen` is able to call `draw` on it.

## Trait Objects Perform Dynamic Dispatch

Recall in the “Performance of Code Using Generics” discussion on the monomorphization process that we use trait bounds on generics: the compiler generates functions and methods for each concrete type that implements the trait for each parameter. The code that results from monomorphization is when the compiler knows what method to call, which is opposed to *dynamic dispatch*, which is when the compiler doesn't know which method to call, which method you're calling. In dynamic dispatch, the compiler at runtime will figure out which method to call.

When we use trait objects, Rust must use dynamic dispatch. The compiler doesn't know all the types that might be used with the trait object, so it doesn't know which method implemented on which type to call. Rust uses the pointers inside the trait object to look up the method at runtime. This has a runtime cost when this lookup happens that doesn't exist in static dispatch. Dynamic dispatch also prevents the compiler from generating specialized code, which in turn prevents some optimizations. This is the trade-off in the code that we wrote in Listing 17-5 and we'll discuss it more when it's a trade-off to consider.

## Object Safety Is Required for Trait Objects

You can only make *object-safe* traits into trait objects. A trait is object safe if it has the properties that make a trait object safe, but not all traits are object safe. A trait is object safe if all the methods and associated constants have the following properties:

- The return type isn't `Self`.
- There are no generic type parameters.

The `Self` keyword is an alias for the type we're working with. Trait objects must be object safe because the compiler no longer knows the concrete type that's implementing the trait. A trait object returns the concrete `Self` type, but a trait object doesn't know the concrete type. There is no way the method can use the original generic type parameters that are filled in with concrete types when the trait is used: the concrete types become part of the trait object. When the type is forgotten through the use of a trait object, the compiler doesn't know what types to fill in the generic type parameters with.

An example of a trait whose methods are not object safe is the `Clone` trait. The signature for the `clone` method is

```
pub trait Clone {
    fn clone(&self) -> Self;
}
```

The `String` type implements the `Clone` trait, and on an instance of `String` we get back an instance of `String`. On an instance of `Vec`, we get back an instance of `Vec`. The `clone` method needs to know what type will stand in for the original type.

The compiler will indicate when you're trying to use a trait object of object safety in regard to trait objects. For example, if you try to use the `Screen` struct in Listing 17-4 to hold types that are not object safe, like the `Draw` trait, like this:

```
pub struct Screen {
    pub components: Vec<Box<dyn Clone>>,
}
```

We would get this error:

```
error[E0038]: the trait `std::clone::Clone`
--> src/lib.rs:2:5
|
|
2 |         pub components: Vec<Box<dyn Clone>
|         ^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^ th
cannot be
made into an object
|
= note: the trait cannot require that `S
```

This error means you can't use this trait as a trait interested in more details on object safety, see [this](#)

## Implementing an Object-Oriented

The *state pattern* is an object-oriented design pattern where an object's value has some internal state, which is represented by a field. The object's behavior changes based on the internal state, thus changing its functionality: in Rust, of course, we use structs and enums instead of classes and inheritance. Each state object is responsible for its own behavior and when it should change into another state. The value of the object is not about nothing about the different behavior of the state objects.

states.

Using the state pattern means when the business rules change, we won't need to change the code of the objects that use the value. We'll only need to update the objects to change its rules or perhaps add more state objects. In the next section, we'll look at the state design pattern and how to use it in Rust.

We'll implement a blog post workflow in an incremental way. The final functionality will look like this:

1. A blog post starts as an empty draft.
2. When the draft is done, a review of the post is requested.
3. When the post is approved, it gets published.
4. Only published blog posts return content to the user. Drafts should not accidentally be published.

Any other changes attempted on a post should be rejected. For example, to approve a draft blog post before we've requested a review, or to publish an unpublished draft.

Listing 17-11 shows this workflow in code form: we'll implement it in a library crate named `blog`. We haven't implemented the `blog` crate yet.

Filename: `src/main.rs`

```
extern crate blog;
use blog::Post;

fn main() {
    let mut post = Post::new();

    post.add_text("I ate a salad for lunch today");
    assert_eq!("", post.content());

    post.request_review();
    assert_eq!("", post.content());

    post.approve();
    assert_eq!("I ate a salad for lunch today", post.content());
}
```

Listing 17-11: Code that demonstrates the desired behavior of the blog post workflow

We want to allow the user to create a new draft blog post. We also want to allow text to be added to the blog post via the `add_text` method.



get the post's content immediately, before approving the post is still a draft. We've added `assert_eq!` for testing purposes. An excellent unit test for this would be one that returns an empty string from the `content` method for this example.

Next, we want to enable a request for a review call to return an empty string while waiting for the review to be published, meaning the text of the review is called.

Notice that the only type we're interacting with for the `Post` type will use the state pattern and will hold a value of `State` objects representing the various states a post can be in. Changing from one state to another is done by the `Post` type. The states change in response to the actions on the `Post` instance, but they don't have to make any changes. Users can't make a mistake with the states, like p

## Defining `Post` and Creating a New Instance

Let's get started on the implementation of the little `Post` struct that holds some content, so we'll start with a `new` function to create a new `Post`. Listing 17-12. We'll also make a private `State` trait of `Box<dyn State>` inside an `Option<T>` in a private module. Why the `Option<T>` is necessary in a bit.

Filename: `src/lib.rs`

```

pub struct Post {
    state: Option<Box<dyn State>>,
    content: String,
}

impl Post {
    pub fn new() -> Post {
        Post {
            state: Some(Box::new(Draft {})),
            content: String::new(),
        }
    }
}

trait State {}

struct Draft {}

impl State for Draft {}

```

Listing 17-12: Definition of a `Post` struct and an instance, a `State` trait, and a `Draft` struct

The `State` trait defines the behavior shared by `PendingReview`, and `Published` states will all inherit from the trait. The trait doesn't have any methods, and we'll start with an empty implementation because that is the state we want a post to start in.

When we create a new `Post`, we set its `state` field to `Some(Box::new(Draft {}))`. This `Box` points to a new instance of the `Draft` struct. When we create a new instance of `Post`, it will start out as a `Draft`. Because `Post` is private, there is no way to create a `Post` directly. In the `new` function, we set the `content` field to a new, empty `String`.

## Storing the Text of the Post Content

Listing 17-11 showed that we want to be able to pass it a `&str` that is then added to the text content. We'll implement this as a method rather than exposing the `content` field directly. We'll implement a method later that will control how the text is added. The `add_text` method is pretty straightforward, so I'll leave it to Listing 17-13 to the `impl Post` block:

Filename: `src/lib.rs`

```
impl Post {
    // --snip--
    pub fn add_text(&mut self, text: &str) {
        self.content.push_str(text);
    }
}
```

Listing 17-13: Implementing the `add_text` method

The `add_text` method takes a mutable reference to the `Post` instance that we're calling `add_text` on. It takes a `String` in `content` and pass the `text` argument. The behavior doesn't depend on the state the post is in. The `add_text` method doesn't interact with any other part of the behavior we want to support.

## Ensuring the Content of a Draft Post Is Empty

Even after we've called `add_text` and added some content to the `content` method to return an empty string slice, we need to ensure the draft state, as shown on line 8 of Listing 17-11. For now, we'll implement the `content` method with the simplest thing that will fulfill the requirement: an empty string slice. We'll change this later once we've added the `publish` method to the `Post` struct. So far, posts in the `draft` state should always be empty. Listing 17-14 shows the implementation:

Filename: `src/lib.rs`

```
impl Post {
    // --snip--
    pub fn content(&self) -> &str {
        ""
    }
}
```

Listing 17-14: Adding a placeholder implementation for the `Post` that always returns an empty string slice

With this added `content` method, everything is working as intended.

## Requesting a Review of the Post Change

Next, we need to add functionality to request a review of the post and change its state from `Draft` to `PendingReview`. Listing 17-15 shows the code.

Filename: `src/lib.rs`

```
impl Post {
    // --snip--
    pub fn request_review(&mut self) {
        if let Some(s) = self.state.take() {
            self.state = Some(s.request_review())
        }
    }
}

trait State {
    fn request_review(self: Box<Self>) -> Box<Self>;
}

struct Draft {}

impl State for Draft {
    fn request_review(self: Box<Self>) -> Box<Self> {
        Box::new(PendingReview {})
    }
}

struct PendingReview {}

impl State for PendingReview {
    fn request_review(self: Box<Self>) -> Box<Self> {
        self
    }
}
```

Listing 17-15: Implementing `request_review` method

We give `Post` a public method named `request_review`. This method takes a mutable reference to `self`. Then we call an internal `request_review` method on the state of `Post`, and this second `request_review` method returns a new state.

We've added the `request_review` method to the `Post` struct. The trait that the trait will now need to implement is `State`. Instead of having `self`, `&self`, or `&mut self` as the first argument to the method, we have `self: Box<Self>`. This syntax is called `Box::new` and is called on a `Box` holding the type. This syntax takes a `Box` and returns a `Box` of the same type.

invalidating the old state so the state value of the state.

To consume the old state, the `request_review` the state value. This is where the `Option` in the `take` method to take the `Some` value out of its place, because Rust doesn't let us have `unpin` move the `state` value out of `Post` rather than `state` value to the result of this operation.

We need to set `state` to `None` temporarily rather like `self.state = self.state.request_review` value. This ensures `Post` can't use the old `state` into a new state.

The `request_review` method on `Draft` needs to new `PendingReview` struct, which represents the review. The `PendingReview` struct also implements doesn't do any transformations. Rather, it returns review on a post already in the `PendingReview` : `PendingReview` state.

Now we can start seeing the advantages of the `state` method on `Post` is the same no matter its `state` its own rules.

We'll leave the `content` method on `Post` as is, can now have a `Post` in the `PendingReview` state we want the same behavior in the `PendingReview` to line 11!

## Adding the `approve` Method that Chan

The `approve` method will be similar to the `request_review` to the value that the current state says it should shown in Listing 17-16:

Filename: `src/lib.rs`

```

impl Post {
    // --snip--
    pub fn approve(&mut self) {
        if let Some(s) = self.state.take()
            self.state = Some(s.approve())
        }
    }
}

trait State {
    fn request_review(self: Box<Self>) ->
    fn approve(self: Box<Self>) -> Box<dyr
}

struct Draft {}

impl State for Draft {
    // --snip--
    fn approve(self: Box<Self>) -> Box<dyr
        self
    }
}

struct PendingReview {}

impl State for PendingReview {
    // --snip--
    fn approve(self: Box<Self>) -> Box<dyr
        Box::new(Published {})
    }
}

struct Published {}

impl State for Published {
    fn request_review(self: Box<Self>) ->
        self
    }

    fn approve(self: Box<Self>) -> Box<dyr
        self
    }
}

```

Listing 17-16: Implementing the `approve` methc

We add the `approve` method to the `State` trait implements `State`, the `Published` state.

Similar to `request_review`, if we call the `approv`

effect because it will return `self`. When we call `publish`, it returns a new, boxed instance of the `Published` struct. `Post` implements the `State` trait, and for both the `publish` and `approve` methods, it returns itself, because the `state` is `self` in those cases.

Now we need to update the `content` method on `Post`. We want to return the value in the post's `content` field. If it's an empty string slice, as shown in Listing 17-17:

Filename: `src/lib.rs`

```
impl Post {  
    // --snip--  
    pub fn content(&self) -> &str {  
        self.state.as_ref().unwrap().content  
    }  
    // --snip--  
}
```

Listing 17-17: Updating the `content` method on `Post` to use the `content` method on `State`

Because the goal is to keep all these rules inside `Post`, we call a `content` method on the value in `state` (`self.state`) as an argument. Then we return the value returned by the `content` method on the `state` value.

We call the `as_ref` method on the `Option` because it's inside the `Option` rather than ownership of the `Option<Box<dyn State>>`, when we call `as_ref` it returns a `&Option`. If we didn't call `as_ref`, we would get the `state` out of the borrowed `&self` of the function.

We then call the `unwrap` method, which we know is safe because the methods on `Post` ensure that `state` will always be `Some` when those methods are done. This is one of the cases in “You Have More Information Than the Compiler” where we know that a `None` value is never possible, even though the compiler doesn't understand that.

At this point, when we call `content` on the `&Box`, it has the same effect on the `&` and the `Box` so the `content` method is called on a type that implements the `State` trait. That means

`State` trait definition, and that is where we'll put the `content` method depending on which state we have, as shown in Listing 17-18.

Filename: `src/lib.rs`

```
trait State {
    // --snip--
    fn content<'a>(&self, post: &'a Post)
        -> String;
}

// --snip--
struct Published {}

impl State for Published {
    // --snip--
    fn content<'a>(&self, post: &'a Post)
        -> String {
        post.content
    }
}
```

Listing 17-18: Adding the `content` method to the `State` trait

We add a default implementation for the `content` method on the `Published` struct. That means we don't need to implement the `content` method on the `PendingReview` struct. The `Published` struct will return the value in `post.content`.

Note that we need lifetime annotations on this `content` method. We're taking a reference to a `post` as an argument, so the lifetime of the returned `String` is tied to the lifetime of the `post` argument.

And we're done—all of Listing 17-11 now works! We've centralized the logic for the blog post workflow. The logic for creating state objects rather than being scattered throughout the code.

## Trade-offs of the State Pattern

We've shown that Rust is capable of implementing the State pattern. To encapsulate the different kinds of behavior as methods on `Post`, the methods on `Post` know nothing about the various states. In the code, we have to look in only one place to know how a `Post` can behave: the implementation of the `State` trait.



If we were to create an alternative implementation we might instead use `match` expressions in the `main` code that checks the state of the post and That would mean we would have to look in several implications of a post being in the published states we added: each of those `match` expressions

With the state pattern, the `Post` methods and the `match` expressions, and to add a new state, we can and implement the trait methods on that one state

The implementation using the state pattern is easy functionality. To see the simplicity of maintaining a few of these suggestions:

- Add a `reject` method that changes the post to `Draft`.
- Require two calls to `approve` before the state
- Allow users to add text content only when have the state object responsible for what not responsible for modifying the `Post`.

One downside of the state pattern is that, because transitions between states, some of the states are another state between `PendingReview` and `Published` would have to change the code in `PendingReview`. It would be less work if `PendingReview` didn't need a new state, but that would mean switching to another

Another downside is that we've duplicated some duplication, we might try to make default implementations and `approve` methods on the `State` trait that violate object safety, because the trait doesn't know exactly. We want to be able to use `State` as a trait to be object safe.

Other duplication includes the similar implementations of `approve` methods on `Post`. Both methods delegate the same method on the value in the `state` field of `Post` to the result. If we had a lot of methods in the state pattern, we might consider defining a macro to reduce (D for more on macros).

By implementing the state pattern exactly as it's

languages, we're not taking as full advantage of look at some changes we can make to the `blog` and transitions into compile time errors.

## Encoding States and Behavior as Types

We'll show you how to rethink the state pattern. Rather than encapsulating the states and transitions, with no knowledge of them, we'll encode the states in Rust's type checking system will prevent attempts to publish posts are allowed by issuing a compile

Let's consider the first part of `main` in Listing 17

Filename: `src/main.rs`

```
fn main() {  
    let mut post = Post::new();  
  
    post.add_text("I ate a salad for lunch");  
    assert_eq!("", post.content());  
}
```

We still enable the creation of new posts in the `new` method, but the ability to add text to the post's content. But instead of a draft post that returns an empty string, we'll make the `content` method at all. That way, if we try to get the content, the compiler error telling us the method doesn't exist will prevent us to accidentally display draft post content in production. Listing 17-19 shows the definition of the `Post` struct, as well as methods on each:

Filename: `src/lib.rs`

```

pub struct Post {
    content: String,
}

pub struct DraftPost {
    content: String,
}

impl Post {
    pub fn new() -> DraftPost {
        DraftPost {
            content: String::new(),
        }
    }

    pub fn content(&self) -> &str {
        &self.content
    }
}

impl DraftPost {
    pub fn add_text(&mut self, text: &str) {
        self.content.push_str(text);
    }
}

```

Listing 17-19: A `Post` with a `content` method a method

Both the `Post` and `DraftPost` structs have a private field for the blog post text. The structs no longer have the `serialize` method for encoding the state to the types of the structs. `Post` is a published post, and it has a `content` method that returns the content of the post.

We still have a `Post::new` function, but instead of returning a `Post`, it returns an instance of `DraftPost`. Because `content` is a method on `Post`, functions that return `Post`, it's not possible to call `content` on a `DraftPost`.

The `DraftPost` struct has an `add_text` method for adding text to the draft post. The program ensures all posts start as draft posts. Any attempt to get the content of a draft post results in a compiler error.

## Implementing Transitions as Transformation

So how do we get a published post? We want to

to be reviewed and approved before it can be published. The `PendingReviewPost` state should still not display any content. Let's introduce another struct, `PendingReviewPost`, defining the `DraftPost` to return a `PendingReviewPost`, and `PendingReviewPost` to return a `Post`, as shown in Listing 17-20.

Filename: `src/lib.rs`

```
impl DraftPost {
    // --snip--

    pub fn request_review(self) -> PendingReviewPost {
        PendingReviewPost {
            content: self.content,
        }
    }
}

pub struct PendingReviewPost {
    content: String,
}

impl PendingReviewPost {
    pub fn approve(self) -> Post {
        Post {
            content: self.content,
        }
    }
}
```

Listing 17-20: A `PendingReviewPost` that gets created from a `DraftPost` and an `approve` method that turns a `PendingReviewPost` into a published `Post`

The `request_review` and `approve` methods take a `DraftPost` and turn it into a `PendingReviewPost` by consuming the `DraftPost` and so forth. The `PendingReviewPost` struct does not have any lingering `DraftPost` instances after we call `request_review` on it, so attempting to read its content results in a runtime error. Because the only way to get a published `Post` from a `PendingReviewPost` is to call the `approve` method, we've now encoded the blog post workflow.

But we also have to make some small changes to the `Post` struct to support the new workflow.

`approve` methods return new instances rather than being called on, so we need to add more `let post =` returned instances. We also can't have the assertion review post's contents be empty strings, nor do we want `main` that tries to use the content of posts in those states. `main` is shown in Listing 17-21:

Filename: src/main.rs

```
extern crate blog;
use blog::Post;

fn main() {
    let mut post = Post::new();

    post.add_text("I ate a salad for lunch today");

    let post = post.request_review();

    let post = post.approve();

    assert_eq!("I ate a salad for lunch today", post.content());
}
```

Listing 17-21: Modifications to `main` to use the review workflow

The changes we needed to make to `main` to realize this design are not too different from the original implementation. The review workflow doesn't quite follow the object-oriented pattern, but the transformations between the states are no longer as complex as in the original `Post` implementation. However, our gain is that because of the type system and the type checker, we can ensure that certain bugs, such as display of the content of a post that has not been approved, can be discovered before they make it to production.

Try the tasks suggested for additional requirements in this section on the `blog` crate as it is after Listing 17-21. This is the design of this version of the code. Note that the review workflow was already completed in this design.

We've seen that even though Rust is capable of implementing object-oriented patterns, other patterns, such as encoding state in a struct, are also available in Rust. These patterns have different trade-offs than object-oriented patterns. Rethinking Rust's features can provide benefits, such as preventing bugs that object-oriented patterns won't always be able to catch. Features, like ownership, that object-oriented languages

## Summary

No matter whether or not you think Rust is an OOP language, this chapter, you now know that you can use traits and dynamic dispatch in Rust. Dynamic dispatch can give you a bit of runtime performance. You can use this for patterns that can help your code's maintainability. Ownership, that object-oriented languages don't have, won't always be the best way to take advantage of Rust's option.

Next, we'll look at patterns, which are another tool for flexibility. We've looked at them briefly through the lens of capability yet. Let's go!

## Patterns and Matching

Patterns are a special syntax in Rust for matching complex and simple values. Using patterns in conjunction with `match` constructs gives you more control over a program's flow, some combination of the following:

- Literals
- Deconstructed arrays, enums, structs, or tuples
- Variables
- Wildcards
- Placeholders

These components describe the shape of the data you want to match against values to determine whether our program should continue running a particular piece of code.

To use a pattern, we compare it to some value. If it matches, we use the value parts in our code. Recall the `match` statement and patterns, such as the coin-sorting machine example. In the pattern, we can use the named pieces. If it doesn't match, the pattern won't run.

This chapter is a reference on all things related to patterns. It shows places to use patterns, the difference between `match` and `if`, the different kinds of pattern syntax that you might encounter, and you'll know how to use patterns to express many

# All the Places Patterns Can Be

Patterns pop up in a number of places in Rust, a without realizing it! This section discusses all the

## `match` Arms

As discussed in Chapter 6, we use patterns in the `match` expression. Formally, `match` expressions are defined as the and one or more match arms that consist of a pattern and a value that matches that arm's pattern, like this:

```
match VALUE {  
    PATTERN => EXPRESSION,  
    PATTERN => EXPRESSION,  
    PATTERN => EXPRESSION,  
}
```

One requirement for `match` expressions is that the pattern must cover all possibilities for the value in the `match` expression. One way to ensure you've covered every possibility is to use the `_` pattern for the last arm: for example, a variable name `m` thus covers every remaining case.

A particular pattern `_` will match anything, but it is often used in the last match arm. The `_` pattern ignores any value not specified, for example. We'll see this in the "Ignoring Values in a Pattern" section later.

## Conditional `if let` Expressions

In Chapter 6 we discussed how to use `if let` expressions to write the equivalent of a `match` that only matches a value and has a corresponding `else` containing code to execute if the value doesn't match.

Listing 18-1 shows that it's also possible to mix `if let` and `else if let` expressions. Doing so gives us multiple conditions in which we can express only one value to compare. We can express a series of `if let`, `else if`, `else` conditions that relate to each other.

The code in Listing 18-1 shows a series of checks for what the background color should be. For this example, we use hardcoded values that a real program might receive from the user.

Filename: src/main.rs

```
fn main() {
    let favorite_color: Option<&str> = None;
    let is_tuesday = false;
    let age: Result<u8, _> = "34".parse();

    if let Some(color) = favorite_color {
        println!("Using your favorite color as the background color");
    } else if is_tuesday {
        println!("Tuesday is green day!");
    } else if let Ok(age) = age {
        if age > 30 {
            println!("Using purple as the background color");
        } else {
            println!("Using orange as the background color");
        }
    } else {
        println!("Using blue as the background color");
    }
}
```

Listing 18-1: Mixing `if let`, `else if`, and `else`

If the user specifies a favorite color, that color is used as the background color. If it's Tuesday, the background color is green. If the user doesn't specify a color, we can parse the age as a number successfully, the background color is determined depending on the value of the number. If none of these conditions are met, the background color is blue.

This conditional structure lets us support complex conditions. For the values we have here, this example will print `Using blue as the background color`.

You can see that `if let` can also introduce shadowed variables: the line `if let Ok(age) = age` introduces a new variable `age` that contains the value inside the `Ok` variant. Then, the `if age > 30` condition within that block: we're actually referring to the `age` variable introduced in `if let Ok(age) = age` & `age > 30`. The shadowed `age` to 30 isn't valid until the new scope starts with the `if` block.

The downside of using `if let` expressions is their exhaustiveness, whereas with `match` expressions you can be more specific about the values you want to handle.



block and therefore missed handling some case the possible logic bug.

## `while let` Conditional Loops

Similar in construction to `if let`, the `while let` loop to run for as long as a pattern continues to shows a `while let` loop that uses a vector as a vector in the opposite order in which they were

```
let mut stack = Vec::new();

stack.push(1);
stack.push(2);
stack.push(3);

while let Some(top) = stack.pop() {
    println!("{}", top);
}
```

Listing 18-2: Using a `while let` loop to print val returns `Some`

This example prints 3, 2, and then 1. The `pop` m the vector and returns `Some(value)` . If the vecto `while` loop continues running the code in its bl When `pop` returns `None` , the loop stops. We can element off our stack.

## `for` Loops

In Chapter 3, we mentioned that the `for` loop is in Rust code, but we haven't yet discussed the p the pattern is the value that directly follows the `|` `x` is the pattern.

Listing 18-3 demonstrates how to use a pattern apart, a tuple as part of the `for` loop.

```
let v = vec!['a', 'b', 'c'];

for (index, value) in v.iter().enumerate()
    println!("{}", value, index)
}
```

Listing 18-3: Using a pattern in a `for` loop to de

The code in Listing 18-3 will print the following:

```
a is at index 0
b is at index 1
c is at index 2
```

We use the `enumerate` method to adapt an iterator to include each value's index in the iterator, placed into a tuple. In this case, the tuple `(0, 'a')`. When this value is matched, the `index` will be `0` and `value` will be `'a'`, printing

## let Statements

Prior to this chapter, we had only explicitly discussed the `if let` statement, but in fact, we've used patterns in other statements. For example, consider this straightforward example:

```
let x = 5;
```

Throughout this book, we've used `let` like this. If you might not have realized it, you were using pattern matching. It looks like this:

```
let PATTERN = EXPRESSION;
```

In statements like `let x = 5;` with a variable name, the name is just a particularly simple form of a pattern. The pattern against the pattern and assigns any names it finds. In this case, it is a pattern that means "bind what matches here to the name `x`". Since the whole pattern is `x`, this pattern effectively binds the variable `x` to whatever the value is.

To see the pattern matching aspect of `let` more

uses a pattern with `let` to destructure a tuple.

```
let (x, y, z) = (1, 2, 3);
```

Listing 18-4: Using a pattern to destructure a tuple

Here, we match a tuple against a pattern. Rust compiler sees that the value matches `x`, `2` to `y`, and `3` to `z`. You can think of this tuple as having individual variable patterns inside it.

If the number of elements in the pattern doesn't match the tuple, the overall type won't match and we'll get an error. Listing 18-5 shows an attempt to destructure a tuple into two variables, which won't work.

```
let (x, y) = (1, 2, 3);
```

Listing 18-5: Incorrectly constructing a pattern with a different number of elements in the tuple

Attempting to compile this code results in this type error:

```
error[E0308]: mismatched types
--> src/main.rs:2:9
   |
2 |     let (x, y) = (1, 2, 3);
   |               ^^^^^^ expected a tuple with 3 elements
   |
   = note: expected type `{integer}, {integer}`
           found type `{integer}, {integer}, {integer}`
```

If we wanted to ignore one or more of the values in the tuple, as you'll see in the "Ignoring Values in a Pattern" section, we can use the underscore. If we have too many variables in the pattern, the solution is to use the underscore to ignore some of the values, removing variables so the number of variables matches the number of elements in the tuple.

## Function Parameters

Function parameters can also be patterns. The code below shows a function named `foo` that takes one parameter. This code should look familiar.

```
fn foo(x: i32) {
    // code goes here
}
```

Listing 18-6: A function signature uses patterns in

The `x` part is a pattern! As we did with `let`, we arguments to the pattern. Listing 18-7 splits the function.

Filename: src/main.rs

```
fn print_coordinates(&(x, y): &(i32, i32))
    println!("Current location: ({}, {})",
}

fn main() {
    let point = (3, 5);
    print_coordinates(&point);
}
```

Listing 18-7: A function with parameters that de

This code prints `Current location: (3, 5)`. Th `&(x, y)`, so `x` is the value `3` and `y` is the valu

We can also use patterns in closure parameter li parameter lists, because closures are similar to

At this point, you've seen several ways of using p same in every place we can use them. In some p irrefutable; in other circumstances, they can be concepts next.

## Refutability: Whether a Pattern

Patterns come in two forms: refutable and irrefutable. Patterns that can match any possible value passed are *irrefutable*. An example is `let x = _`; because `x` matches anything and therefore always matches. Patterns that can fail to match for some possible values would be `Some(x)` in the expression `if let Some(x) = a_value` in the `a_value` variable is `None` rather than `Some(x)` match.

Function parameters, `let` statements, and `for` patterns, because the program cannot do anything to match. The `if let` and `while let` expressions are allowed because by definition they're intended to handle the case where the conditional is in its ability to perform differently.

In general, you shouldn't have to worry about the irrefutable patterns; however, you do need to be aware of refutability so you can respond when you see it. If you'll need to change either the pattern or the code depending on the intended behavior of the code, you'll need to use a refutable pattern.

Let's look at an example of what happens when you use a refutable pattern where Rust requires an irrefutable pattern and a `let` statement, but for the pattern we've specified `Some(x)` might expect, this code will not compile.

```
let Some(x) = some_option_value;
```

Listing 18-8: Attempting to use a refutable pattern where an irrefutable one is required

If `some_option_value` was a `None` value, it would mean the pattern is refutable. However, the `let` statement requires an irrefutable pattern because there is nothing valid to do if the pattern doesn't match. At compile time, Rust will complain that we've tried to use a refutable pattern where an irrefutable one is required:

```
error[E0005]: refutable pattern in local binding
-->
  |
3 | let Some(x) = some_option_value;
  |          ^^^^^^^ pattern `None` not covered
```

Because we didn't cover (and couldn't cover!) even the `Some(x)` case, Rust rightfully produces a compiler error.

To fix the problem where we have a refutable pattern where an irrefutable one is needed, we can change the code that uses the pattern to use `if let`. Then if the pattern doesn't match, we can use curly brackets, giving it a way to continue validly. Here's the corrected code in Listing 18-8.

```
if let Some(x) = some_option_value {
    println!("{}", x);
}
```

Listing 18-9: Using `if let` and a block with `refutable`

We've given the code an out! This code is perfect for use an irrefutable pattern without receiving an error. It will always match, such as `x`, as shown in Listing 18-10.

```
if let x = 5 {
    println!("{}", x);
};
```

Listing 18-10: Attempting to use an irrefutable pattern

Rust complains that it doesn't make sense to use an irrefutable pattern in a block.

```
error[E0162]: irrefutable if-let pattern
--> <anon>:2:8
   |
 2 | if let x = 5 {
   |           ^ irrefutable pattern
```

For this reason, match arms must use `refutable` patterns, which should match any remaining values with a `catch-all` pattern. To use an irrefutable pattern in a `match` with one arm is particularly useful and could be replaced with a `catch-all` pattern.

Now that you know where to use patterns and the difference between irrefutable patterns, let's cover all the syntax we can use in patterns.

## Pattern Syntax

Throughout the book, you've seen examples of patterns. In this chapter, we gather all the syntax valid in patterns and discuss each one.

### Matching Literals

As you saw in Chapter 6, you can match patterns with literals. For example,

following code gives some examples:

```
let x = 1;

match x {
  1 => println!("one"),
  2 => println!("two"),
  3 => println!("three"),
  _ => println!("anything"),
}
```

This code prints `one` because the value in `x` is `1`. You can use `match` to tell your code to take an action if it gets a particular value.

## Matching Named Variables

Named variables are irrefutable patterns that match a single value. We use them many times in the book. However, there is a common pattern for using named variables in `match` expressions. Because `match` is declared as part of a pattern inside the `match` expression, it can't have the same name outside the `match` construct, as is the case with `let`. In Listing 18-11, we declare a variable named `x` with the value `10`. We then create a `match` expression with two patterns in the match arms and `println!` at the end of each arm. The code will print before running this code or reading it.

Filename: src/main.rs

```
fn main() {
  let x = Some(5);
  let y = 10;

  match x {
    Some(50) => println!("Got 50"),
    Some(y) => println!("Matched, y = {y}"),
    _ => println!("Default case, x = {x}"),
  }

  println!("at the end: x = {x}, y = {y}");
}
```

Listing 18-11: A `match` expression with an arm that matches the value of `y`

Let's walk through what happens when the `match` expression is evaluated.

first match arm doesn't match the defined value

The pattern in the second match arm introduces match any value inside a `Some` value. Because `y` is a new variable, this is a new `y` variable, not the `y` variable that had the value 10. This new `y` binding will match any value that has in `x`. Therefore, this new `y` binds to the inner value, so the expression for that arm executes and prints `at the end`.

If `x` had been a `None` value instead of `Some(5)`, the first match arm wouldn't have matched, so the value would have been `None`. The `match` expression didn't introduce the `x` variable in the pattern of the second arm, so the `x` in the expression is still the outer `x` that hasn't been shadowed. Therefore, the `match` would print `Default case, x = None`.

When the `match` expression is done, its scope ends, and the inner `y` is no longer in scope. The last `println!` produces `at the end`.

To create a `match` expression that compares the value of `x` without introducing a shadowed variable, we would use a `match guard` instead. We'll talk about match guards in the "Match Guards" section.

## Multiple Patterns

In `match` expressions, you can match multiple patterns, which means *or*. For example, the following code matches `x` against three patterns, the first of which has an *or* option, meaning that if either of the values in that arm, that arm's code will run:

```
let x = 1;

match x {
  1 | 2 => println!("one or two"),
  3 => println!("three"),
  _ => println!("anything"),
}
```

This code prints `one or two`.

## Matching Ranges of Values with ...



The `...` syntax allows us to match to an inclusive range of values, when a pattern matches any of the values in the range, the code will execute:

```
let x = 5;

match x {
    1 ... 5 => println!("one through five"),
    _ => println!("something else"),
}
```

If `x` is 1, 2, 3, 4, or 5, the first arm will match. This is equivalent to using the `|` operator to express the same idea; for example, to specify `1 | 2 | 3 | 4 | 5` if we used `|`. Specifically, this is useful especially if we want to match, say, any number in a range.

Ranges are only allowed with numeric values or character values. Checks that the range isn't empty at compile time. We can tell if a range is empty or not are `char` and `num`.

Here is an example using ranges of `char` values:

```
let x = 'c';

match x {
    'a' ... 'j' => println!("early ASCII letters"),
    'k' ... 'z' => println!("late ASCII letters"),
    _ => println!("something else"),
}
```

Rust can tell that `c` is within the first pattern's range.

## Destructuring to Break Apart Values

We can also use patterns to destructure structs, tuples, and other composite values. Let's walk through an example.

### Destructuring Structs

Listing 18-12 shows a `Point` struct with two fields. We can destructure it using a pattern with a `let` statement.

Filename: src/main.rs

```

struct Point {
    x: i32,
    y: i32,
}

fn main() {
    let p = Point { x: 0, y: 7 };

    let Point { x: a, y: b } = p;
    assert_eq!(0, a);
    assert_eq!(7, b);
}

```

Listing 18-12: Destructuring a struct's fields into

This code creates the variables `a` and `b` that match the fields of the `p` variable. This example shows that the variable names don't have to match the field names of the struct variable names to match the field names to make the variables come from which fields.

Because having variable names match the fields of a struct, `let Point { x: x, y: y } = p;` contains a lot of boilerplate for patterns that match struct fields: you only need to write the struct name and the variables created from the pattern will have the same names. Listing 18-13 shows code that behaves in the same way as the code in Listing 18-12, but the variables created in the `let` pattern are `x` and `y`.

Filename: src/main.rs

```

struct Point {
    x: i32,
    y: i32,
}

fn main() {
    let p = Point { x: 0, y: 7 };

    let Point { x, y } = p;
    assert_eq!(0, x);
    assert_eq!(7, y);
}

```

Listing 18-13: Destructuring struct fields using `struct`

This code creates the variables `x` and `y` that match the fields of the `p` variable. The outcome is that the variables `x` and `y` are the same as the struct fields.

We can also destructure with literal values as pa creating variables for all the fields. Doing so allo particular values while creating variables to dest

Listing 18-14 shows a `match` expression that se cases: points that lie directly on the `x` axis (whic (`x = 0`)), or neither.

Filename: src/main.rs

```
fn main() {  
    let p = Point { x: 0, y: 7 };  
  
    match p {  
        Point { x, y: 0 } => println!("On  
        Point { x: 0, y } => println!("On  
        Point { x, y } => println!("On nei  
    }  
}
```

Listing 18-14: Destructuring and matching literal

The first arm will match any point that lies on th matches if its value matches the literal `0`. The p we can use in the code for this arm.

Similarly, the second arm matches any point on field matches if its value is `0` and creates a varia The third arm doesn't specify any literals, so it m variables for both the `x` and `y` fields.

In this example, the value `p` matches the second so this code will print `On the y axis at 7`.

## Destructuring Enums

We've deconstructed enums earlier in this book, f `Option<i32>` in Listing 6-5 in Chapter 6. One de that the pattern to destructure an enum should stored within the enum is defined. As an exampl `Message` enum from Listing 6-2 and write a `mat` each inner value.

Filename: src/main.rs

```

enum Message {
    Quit,
    Move { x: i32, y: i32 },
    Write(String),
    ChangeColor(i32, i32, i32),
}

fn main() {
    let msg = Message::ChangeColor(0, 160,

    match msg {
        Message::Quit => {
            println!("The Quit variant has
        },
        Message::Move { x, y } => {
            println!(
                "Move in the x direction {
                x,
                y
            );
        }
        Message::Write(text) => println!(
        Message::ChangeColor(r, g, b) => {
            println!(
                "Change the color to red {
                r,
                g,
                b
            )
        }
    }
}

```

Listing 18-15: Deconstructing enum variants that

This code will print `Change the color to red 0` changing the value of `msg` to see the code from

For enum variants without any data, like `Message` value any further. We can only match on the lite variables are in that pattern.

For struct-like enum variants, such as `Message::` the pattern we specify to match structs. After the brackets and then list the fields with variables so the code for this arm. Here we use the shorthand

For tuple-like enum variants, like `Message::Write` element and `Message::ChangeColor` that holds

pattern is similar to the pattern we specify to match. The pattern in the pattern must match the number of elements in the value.

## Destructuring Nested Structs & Enums

Up until now, all of our examples have been matching on simple values. Matching can work on nested structures too.

We can refactor the example above to support both RGB and HSV color representations.

```
enum Color {
    Rgb(i32, i32, i32),
    Hsv(i32, i32, i32)
}

enum Message {
    Quit,
    Move { x: i32, y: i32 },
    Write(String),
    ChangeColor(Color),
}

fn main() {
    let msg = Message::ChangeColor(Color::Rgb(0, 0, 0));

    match msg {
        Message::ChangeColor(Color::Rgb(r, g, b)) => {
            println!(
                "Change the color to red {r}, green {g}, blue {b}"
            );
        },
        Message::ChangeColor(Color::Hsv(h, s, v)) => {
            println!(
                "Change the color to hue {h}, saturation {s}, value {v}"
            );
        },
        _ => {}
    }
}
```

## Destructuring References

When the value we're matching to our pattern contains a reference, we can use the `&` keyword to match the reference.

destructure the reference from the value, which pattern. Doing so lets us get a variable holding the value rather than getting a variable that holds the reference. This is useful in closures where we have iterators that iterate over references to use the values in the closure rather than the references.

The example in Listing 18-16 iterates over references and destructures the reference and the struct so we can access the `x` and `y` values easily.

```
let points = vec![
    Point { x: 0, y: 0 },
    Point { x: 1, y: 5 },
    Point { x: 10, y: -3 },
];

let sum_of_squares: i32 = points
    .iter()
    .map(|&Point { x, y }| x * x + y * y)
    .sum();
```

Listing 18-16: Destructuring a reference to a struct

This code gives us the variable `sum_of_squares` with the result of squaring the `x` value and the `y` value, adding the result for each `Point` in the `points` vector.

If we had not included the `&` in `&Point { x, y }`, we would get an error because `iter` would then iterate over references to `Point` rather than the actual values. The error would look like

```
error[E0308]: mismatched types
-->
   |
14 |         .map(|Point { x, y }| x * x + y * y)
   |               ^^^^^^^^^^^^^ expected `&Point`, found `Point`
   |
   = note: expected type `&Point`
           found type `Point`
```

This error indicates that Rust was expecting our closure to match directly to a `Point` value, not a reference to a `Point`.

## Destructuring Structs and Tuples

We can mix, match, and nest destructuring patterns to access fields of structs and tuples.

following example shows a complicated destruct inside a tuple and destructure all the primitive v

```
let ((feet, inches), Point {x, y}) = ((3,
```

This code lets us break complex types into their values we're interested in separately.

Destructuring with patterns is a convenient way value from each field in a struct, separately from

## Ignoring Values in a Pattern

You've seen that it's sometimes useful to ignore arm of a `match`, to get a catchall that doesn't ac for all remaining possible values. There are a few of values in a pattern: using the `_` pattern (whic within another pattern, using a name that starts ignore remaining parts of a value. Let's explore l patterns.

### Ignoring an Entire Value with `_`

We've used the underscore (`_`) as a wildcard pa not bind to the value. Although the underscore last arm in a `match` expression, we can use it in parameters, as shown in Listing 18-17.

Filename: src/main.rs

```
fn foo(_: i32, y: i32) {  
    println!("This code only uses the y pa  
}  
  
fn main() {  
    foo(3, 4);  
}
```

Listing 18-17: Using `_` in a function signature

This code will completely ignore the value passe  
print `This code only uses the y parameter:`

In most cases when you no longer need a particular parameter, you can change the signature so it doesn't include the unused parameter. This can be especially useful in some cases, such as when you need a certain type signature but the implementation doesn't need one of the parameters. This is about unused function parameters, as it would be in the case of a trait.

## Ignoring Parts of a Value with a Nested `_`

We can also use `_` inside another pattern to ignore parts of a value when we want to test for only part of a value but still run the corresponding code we want to run. Listing 18-18 shows managing a setting's value. The business requirement is that a user should be allowed to overwrite an existing customization setting and can give the setting a value if it is currently `None`.

```
let mut setting_value = Some(5);
let new_setting_value = Some(10);

match (setting_value, new_setting_value) {
    (Some(_), Some(_)) => {
        println!("Can't overwrite an existing setting")
    }
    _ => {
        setting_value = new_setting_value;
    }
}

println!("setting is {:?}", setting_value)
```

Listing 18-18: Using an underscore within a pattern to ignore parts of a value that don't need to be used inside the `Some` variant.

This code will print `Can't overwrite an existing setting is Some(5)`. In the first match arm, we check if both `setting_value` and `new_setting_value` are `Some` variants, but we do not need the values inside either `Some` variant, so we use `_` to ignore them. This is why we're not changing `setting_value`, and it's why we're not changing `new_setting_value`.

In all other cases (if either `setting_value` or `new_setting_value` is `None`), the code is expressed by the `_` pattern in the second arm, which assigns the value to become `setting_value`.

We can also use underscores in multiple places to ignore particular values. Listing 18-19 shows an example of this.



values in a tuple of five items.

```
let numbers = (2, 4, 8, 16, 32);

match numbers {
    (first, _, third, _, fifth) => {
        println!("Some numbers: {}, {}, {}",
    },
}
```

Listing 18-19: Ignoring multiple parts of a tuple

This code will print `Some numbers: 2, 8, 32`, all

## Ignoring an Unused Variable by Starting Its Name with an Underscore

If you create a variable but don't use it anywhere, the compiler will warn you about it because that could be a bug. But sometimes it's useful to ignore a variable, such as when you're prototyping or just want to ignore a value. You can tell Rust not to warn you about the unused variable with an underscore. In Listing 18-20, we use a variable that we don't use, and when we run this code, we should only get a warning about the unused variable.

Filename: src/main.rs

```
fn main() {
    let _x = 5;
    let y = 10;
}
```

Listing 18-20: Starting a variable name with an underscore to ignore unused variable warnings

Here we get a warning about not using the variable `_x`. The warning message says "variable preceded by the unused variable `_x`".

Note that there is a subtle difference between `unused` and `unbound`. `unused` starts with an underscore. The syntax `_x` still binds the variable `x`, but `_` doesn't bind at all. To show a case where this difference matters, we provide us with an error.

```
let s = Some(String::from("Hello!"));

if let Some(_s) = s {
    println!("found a string");
}

println!("{:?}", s);
```

Listing 18-21: An unused variable starting with a `_` which might take ownership of the value

We'll receive an error because the `s` value will shield us from using `s` again. However, using the underscore `_` shields the value. Listing 18-22 will compile without any errors into `_`.

```
let s = Some(String::from("Hello!"));

if let Some(_) = s {
    println!("found a string");
}

println!("{:?}", s);
```

Listing 18-22: Using an underscore does not bind the variable

This code works just fine because we never bind the variable `s`.

## Ignoring Remaining Parts of a Value with `..`

With values that have many parts, we can use the `..` pattern and ignore the rest, avoiding the need to list unused parts of the pattern. In Listing 18-23, we have a `Point` in three-dimensional space. In the `match` expression, we use the `x` coordinate and ignore the values in the `y` and `z` coordinates.

```

struct Point {
    x: i32,
    y: i32,
    z: i32,
}

let origin = Point { x: 0, y: 0, z: 0 };

match origin {
    Point { x, .. } => println!("x is {}",
}

```

Listing 18-23: Ignoring all fields of a `Point` except

We list the `x` value and then just include the `..` to list `y: _` and `z: _`, particularly when we're viewing fields in situations where only one or two fields

The syntax `..` will expand to as many values as how to use `..` with a tuple.

Filename: src/main.rs

```

fn main() {
    let numbers = (2, 4, 8, 16, 32);

    match numbers {
        (first, .., last) => {
            println!("Some numbers: {}, {}",
        },
    }
}

```

Listing 18-24: Matching only the first and last values

In this code, the first and last value are matched and ignore everything in the middle.

However, using `..` must be unambiguous. If it is for matching and which should be ignored, Rust shows an example of using `..` ambiguously, so

Filename: src/main.rs

```
fn main() {
    let numbers = (2, 4, 8, 16, 32);

    match numbers {
        (.., second, ..) => {
            println!("Some numbers: {}", s
                },
        }
    }
}
```

Listing 18-25: An attempt to use `..` in an ambiguous match

When we compile this example, we get this error:

```
error: `..` can only be used once per tuple pattern
--> src/main.rs:5:22
   |
5 |         (.., second, ..) => {
   |                      ^^
```

It's impossible for Rust to determine how many values to match with `second` and then how many to ignore thereafter. This code could mean that we want to match `2` and then ignore `8`, `16`, and `32`; or that we want to match `2` and `4` and then ignore `16` and `32`; and so forth. The compiler has nothing special to Rust, so we get a compiler error like this is ambiguous.

## Extra Conditionals with Match Guards

A *match guard* is an additional `if` condition specified for a match arm that must also match, along with the pattern chosen. Match guards are useful for expressing conditions that alone allows.

The condition can use variables created in the pattern where the first arm has the pattern `Some(x)` and an `if x < 5`.

```
let num = Some(4);

match num {
    Some(x) if x < 5 => println!("less than five: {}", x),
    Some(x) => println!("{}", x),
    None => (),
}
```

Listing 18-26: Adding a match guard to a pattern

This example will print `less than five: 4`. When the first arm, it matches, because `Some(4)` matches. It checks whether the value in `x` is less than `5`, and if so, it is selected.

If `num` had been `Some(10)` instead, the match guard would be false because 10 is not less than 5. Rust would then try the second arm, which would match because the second arm doesn't have a match guard. It matches any `Some` variant.

There is no way to express the `if x < 5` condition without a match guard. A match guard gives us the ability to express this logic.

In Listing 18-11, we mentioned that we could use variable shadowing to solve the shadowing problem. Recall that a new variable `y` shadows the previous `y`. Using a `match` expression instead of using the variable `y` meant we couldn't test against the value of the original `y`. Now we can use a match guard to fix this problem.

Filename: src/main.rs

```
fn main() {
    let x = Some(5);
    let y = 10;

    match x {
        Some(50) => println!("Got 50"),
        Some(n) if n == y => println!("Matched y: {}", n),
        _ => println!("Default case, x = {:?}", x),
    }

    println!("at the end: x = {:?}", x, y = y);
}
```

Listing 18-27: Using a match guard to test for equality

This code will now print `Default case, x = Some(5), y = 10`.

match arm doesn't introduce a new variable `y` that means we can use the outer `y` in the match guard pattern as `Some(y)`, which would have shadowed the outer `y`. This creates a new variable `n` that doesn't shadow the variable outside the `match`.

The match guard `if n == y` is not a pattern and doesn't introduce new variables. This `y` is the outer `y` rather than a new variable. It's a value that has the same value as the outer `y`.

You can also use the `or` operator `|` in a match guard. The match guard condition will apply to all the patterns. The precedence of combining a match guard with a pattern is that the `if y` match guard applies to all the patterns. part of this example is that the `if y` match guard applies to all the patterns though it might look like `if y` only applies to `6`.

```
let x = 4;
let y = false;

match x {
  4 | 5 | 6 if y => println!("yes"),
  _ => println!("no"),
}
```

Listing 18-18: Combining multiple patterns with a match guard

The match condition states that the arm only matches `5`, or `6` and if `y` is `true`. When this code runs, because `x` is `4`, but the match guard `if y` is `false`, the code moves on to the second arm, which does `println!("no")`. The reason is that the `if` condition applies only to the last value `6`. In other words, the precedence of combining a match guard with a pattern behaves like this:

```
(4 | 5 | 6) if y => ...
```

rather than this:

```
4 | 5 | (6 if y) => ...
```

After running the code, the precedence behavior is applied only to the final value in the list of values. If `y` were `true`, the first arm would have matched and the program would have printed `yes`.

## @ Bindings

The *at* operator ( `@` ) lets us create a variable that testing that value to see whether it matches a pattern. For example where we want to test that a `Message::Hello` has an `id` in the range `3...7`. But we also want to bind the value to the variable `id_variable` in the code associated with the arm. We could use the `id` field, but for this example we'll use a different variable.

```
enum Message {
    Hello { id: i32 },
}

let msg = Message::Hello { id: 5 };

match msg {
    Message::Hello { id: id_variable @ 3..7 } => {
        println!("Found an id in range: {}", id_variable),
    },
    Message::Hello { id: 10..12 } => {
        println!("Found an id in another range: {}", id),
    },
    Message::Hello { id } => {
        println!("Found some other id: {}", id),
    },
}
```

Listing 18-19: Using `@` to bind to a value in a pattern

This example will print `Found an id in range: 3...7`, before the range `3...7`, we're capturing whatever value is in the `id` field, also testing that the value matched the range pattern.

In the second arm, where we only have a range pattern associated with the arm doesn't have a variable `id` field. The `id` field's value could have been 10 or 11, but the pattern doesn't know which it is. The pattern captures the value from the `id` field, because we haven't saved the value in a variable.

In the last arm, where we've specified a variable `id` field, we have a value available to use in the arm's code in a variable `id`. We've used the struct field shorthand syntax. But we haven't saved the value in the `id` field in this arm, as we did with the first arm, so we can't match this pattern.

Using `@` lets us test a value and save it in a variable.

## Legacy patterns: `ref` and `ref mut`

In older versions of Rust, `match` would assume `name` was `Some`. But sometimes, that's not what you want.

```
let robot_name = &Some(String::from("Borsalino"));

match robot_name {
    Some(name) => println!("Found a name: {}", name),
    None => (),
}

println!("robot_name is: {:?}", robot_name);
```

Here, `robot_name` is a `&Option<String>`. Rust v1.0 doesn't match up with `&Option<T>`, so you'd have to write:

```
let robot_name = &Some(String::from("Borsalino"));

match robot_name {
    &Some(name) => println!("Found a name: {}", name),
    None => (),
}

println!("robot_name is: {:?}", robot_name);
```

Next, Rust would complain that `name` is trying to move out of `robot_name` but because it's a reference to an option, it's borrowed. This is where the `ref` keyword comes into play:

```
let robot_name = &Some(String::from("Borsalino"));

match robot_name {
    &Some(ref name) => println!("Found a name: {}", name),
    None => (),
}

println!("robot_name is: {:?}", robot_name);
```

The `ref` keyword is like the opposite of `&` in that it can't be a `&String`, don't try to move it out. In other words, it's not `ref` against a reference, but `ref` creates a reference to mutable references.

Anyway, today's Rust doesn't work like this. If you borrow, then all of the bindings you create will be borrowed.



means that the original code works as you'd expect.

Because Rust is backwards compatible, we could use these patterns in older Rust code, so knowing what they do is still useful. Sometimes they're sometimes useful in obscure situations, but they're part of a struct as mutable and another part as immutable.

## Summary

Rust's patterns are very useful in that they help you work with data. When used in `match` expressions, Rust ensures that every possible value, or your program won't compile. In function parameters make those constructs more complex. Patterns of values into smaller parts at the same time as simple or complex patterns to suit our needs.

Next, for the penultimate chapter of the book, we'll cover a variety of Rust's features.

## Advanced Features

By now, you've learned the most commonly used features of the language. Before we do one more project in Chapter 16, the language you might run into every once in a while. The reference for when you encounter any unknown feature. The features you learn to use in this chapter are useful in very specific situations. If you don't reach for them often, we want to make sure you know Rust has to offer.

In this chapter, we'll cover:

- Unsafe Rust: how to opt out of some of Rust's guarantees for manually upholding those guarantees
- Advanced lifetimes: syntax for complex lifetime relationships
- Advanced traits: associated types, default trait methods, syntax, supertraits, and the newtype pattern
- Advanced types: more about the newtype pattern, dynamically sized types
- Advanced functions and closures: function pointers

It's a panoply of Rust features with something for everyone.

# Unsafe Rust

All the code we've discussed so far has had Rust enforced at compile time. However, Rust has a `unsafe` mode that doesn't enforce these memory safety guarantees like regular Rust, but gives us extra superpowers:

Unsafe Rust exists because, by nature, static analysis and the compiler tries to determine whether or not code is safe. Sometimes, for it to reject some valid programs rather than accept them. Although the code might be okay, as far as Rust is concerned, you can use unsafe code to tell the compiler, "Trust me, this is safe." The downside is that you use it at your own risk: if you have memory problems due to memory unsafety, such as null pointer dereference, it's your fault.

Another reason Rust has an unsafe alter ego is that some hardware is inherently unsafe. If Rust didn't let you do certain tasks, Rust needs to allow you to do low-level tasks as directly interacting with the operating system or hardware. Working with low-level systems programming language. Let's explore what we can do with unsafe Rust.

## Unsafe Superpowers

To switch to unsafe Rust, use the `unsafe` keyword to mark blocks of code that hold the unsafe code. You can take four actions, or *superpowers*, that you can't in safe Rust. Those are:

- Dereference a raw pointer
- Call an unsafe function or method
- Access or modify a mutable static variable
- Implement an unsafe trait

It's important to understand that `unsafe` doesn't disable any other of Rust's safety checks: if you use `unsafe`, you still be checked. The `unsafe` keyword only gives you a degree of safety inside of an unsafe block.

In addition, `unsafe` does not mean the code inside is unsafe or that it will definitely have memory safety problems. As a programmer, you'll ensure the code inside an `unsafe` block is valid way.

People are fallible, and mistakes will happen, but operations to be inside blocks annotated with `unsafe` related to memory safety must be within an `unsafe` block; you'll be thankful later when you investigate.

To isolate unsafe code as much as possible, it's better to provide a safe abstraction and provide a safe API, which we examine unsafe functions and methods. Part of the Rust standard library is implemented as safe abstractions over unsafe code. Using unsafe code in a safe abstraction prevents uses of unsafe code in the places that you or your users might want to use with `unsafe` code, because using a safe abstraction is safer.

Let's look at each of the four unsafe superpower abstractions that provide a safe interface to unsafe code.

## Dereferencing a Raw Pointer

In Chapter 4, in the “Dangling References” section, we saw how Rust ensures references are always valid. Unsafe Rust has *raw pointers* that are similar to references. As with references, raw pointers can be `immutable` or `mutable` and are written as `*const T` or `*mut T`. The asterisk isn't the dereference operator; it's part of the pointer type. *Raw pointers*, *immutable* means that the pointer being dereferenced is not mutable.

Different from references and smart pointers, raw pointers:

- Are allowed to ignore the borrowing rules (e.g., multiple mutable pointers or multiple mutable pointers to the same memory)
- Aren't guaranteed to point to valid memory
- Are allowed to be null
- Don't implement any automatic cleanup

By opting out of having Rust enforce these guarantees, you trade safety in exchange for greater performance or to interface with legacy code or hardware where Rust's guarantees don't apply.

Listing 19-1 shows how to create an immutable raw pointer to a reference.

```
let mut num = 5;

let r1 = &num as *const i32;
let r2 = &mut num as *mut i32;
```

Listing 19-1: Creating raw pointers from references

Notice that we don't include the `unsafe` keyword for creating raw pointers in safe code; we just can't dereference them as you'll see in a bit.

We've created raw pointers by using `as` to cast a reference into their corresponding raw pointer type directly from references guaranteed to be valid, so the raw pointers are valid, but we can't make that assumption.

Next, we'll create a raw pointer whose validity we can't guarantee. This shows how to create a raw pointer to an arbitrary memory address. Arbitrary memory is undefined: there might be data there or not, the compiler might optimize the code so that the data is not there, the program might error with a segmentation fault. We can write code like this, but it is possible.

```
let address = 0x012345usize;
let r = address as *const i32;
```

Listing 19-2: Creating a raw pointer to an arbitrary memory address

Recall that we can create raw pointers in safe code and dereference them to read the data being pointed to. In Listing 19-3, we use the dereference operator `*` on a raw pointer that requires an `unsafe` block.

```
let mut num = 5;

let r1 = &num as *const i32;
let r2 = &mut num as *mut i32;

unsafe {
    println!("r1 is: {}", *r1);
    println!("r2 is: {}", *r2);
}
```

Listing 19-3: Dereferencing raw pointers within an `unsafe` block

Creating a pointer does no harm; it's only when it points at that we might end up dealing with an issue.

Note also that in Listing 19-1 and 19-3, we create pointers that both pointed to the same memory instead of trying to create an immutable and a mutable pointer. This code would not have compiled because Rust's ownership rules require that there be at the same time as any immutable references. With one mutable pointer and an immutable pointer to the same memory, through the mutable pointer, potentially creating a dangling pointer.

With all of these dangers, why would you ever use `unsafe`? The answer is when interfacing with C code, as you'll see in the next chapter. Another case is when building a library that the Rust borrow checker doesn't understand. We'll introduce an example of a safe abstraction that uses `unsafe` in the next chapter.

## Calling an Unsafe Function or Method

The second type of operation that requires an `unsafe` block is calling unsafe functions and methods. Unsafe functions and methods look like regular functions and methods, but they have an extra `unsafe` before the function name. The `unsafe` keyword in this context indicates the function or method does not uphold Rust's safety requirements. By calling an unsafe function with `unsafe`, we're saying we've read this function's documentation and take responsibility for its contracts.

Here is an unsafe function named `dangerous` that takes a `String` and returns a `String`.

```
unsafe fn dangerous() {}

unsafe {
    dangerous();
}
```

We must call the `dangerous` function within a `unsafe` block. If we call `dangerous` without the `unsafe` block, we'll get a compiler error.

```

error[E0133]: call to unsafe function requires unsafe
  -->
   |
4  |     dangerous();
   |     ^^^^^^^^^^^^^ call to unsafe function

```

By inserting the `unsafe` block around our call to `dangerous()`, that we've read the function's documentation, we can mark it as `unsafe` and we've verified that we're fulfilling the contract.

Bodies of unsafe functions are effectively `unsafe` operations. Within an unsafe function, we don't need to mark individual operations as `unsafe`.

## Creating a Safe Abstraction over Unsafe Code

Just because a function contains unsafe code doesn't mean the entire function is unsafe. In fact, wrapping unsafe code in a `unsafe` block is a common abstraction. As an example, let's study `split_at_mut`, that requires some unsafe code to implement it. This safe method is defined on `Vec` and makes it two by splitting the slice at the index given. Listing 19-4 shows how to use `split_at_mut`.

```

let mut v = vec![1, 2, 3, 4, 5, 6];

let r = &mut v[..];

let (a, b) = r.split_at_mut(3);

assert_eq!(a, &mut [1, 2, 3]);
assert_eq!(b, &mut [4, 5, 6]);

```

Listing 19-4: Using the safe `split_at_mut` function

We can't implement this function using only safe code, something like Listing 19-5, which won't compile because `split_at_mut` is implemented as a function rather than a method. Listing 19-6 shows how to implement `split_at_mut` for a generic type `T`.

```
fn split_at_mut(slice: &mut [i32], mid: usize)
{
    let len = slice.len();

    assert!(mid <= len);

    (&mut slice[..mid],
     &mut slice[mid..])
}
```

Listing 19-5: An attempted implementation of `split_at_mut`

This function first gets the total length of the slice as a parameter is within the slice by checking whether the `mid` length. The assertion means that if we pass an invalid `mid` to split the slice at, the function will panic before it returns.

Then we return two mutable slices in a tuple: one from the start to the `mid` index and another from `mid` to the end of the slice.

When we try to compile the code in Listing 19-5,

```
error[E0499]: cannot borrow `*slice` as mutable more than once at a time
-->
  |
6 |         (&mut slice[..mid],
  |         ----- first mutable borrow occurs here
7 |         &mut slice[mid..])
  |         ^^^^^ second mutable borrow occurs here
8 |     }
  |     - first borrow ends here
```

Rust's borrow checker can't understand that we're borrowing from different parts of a slice; it only knows that we're borrowing from the same slice. Borrowing from different parts of a slice is fundamentally okay because the slices don't overlap, but Rust isn't smart enough to know that. But Rust doesn't, it's time to reach for unsafe code.

Listing 19-6 shows how to use an `unsafe` block, `unsafe` functions to make the implementation of `split_at_mut` work.

```

use std::slice;

fn split_at_mut(slice: &mut [i32], mid: usize)
{
    let len = slice.len();
    let ptr = slice.as_mut_ptr();

    assert!(mid <= len);

    unsafe {
        (slice::from_raw_parts_mut(ptr, mid),
         slice::from_raw_parts_mut(ptr.offset(mid as isize), len - mid))
    }
}

```

Listing 19-6: Using unsafe code in the implementation

Recall from “The Slice Type” section in Chapter 4 and the length of the slice. We use the `len` method to get the length of the slice. We use the `as_mut_ptr` method to access the raw pointer to the slice. We have a mutable slice to `i32` values, `as_mut_ptr` returns a `*mut i32`, which we’ve stored in the variable `ptr`.

We keep the assertion that the `mid` index is within the slice. In the unsafe code: the `slice::from_raw_parts_mut` function takes a raw pointer, a length, and it creates a slice. We use this function to create a slice from `ptr` and is `mid` items long. Then we call the `offset` method with `mid` as an argument to get a raw pointer that starts at `mid` items from the start pointer and the remaining number of items after `mid`.

The function `slice::from_raw_parts_mut` is unsafe and must trust that this pointer is valid. The `offset` method is also unsafe, because it must trust that the offset location is valid. We had to put an `unsafe` block around our calls to `from_raw_parts_mut` and `offset` so we could call them. By looking at the code, we can see that `mid` must be less than or equal to `len`, we can be sure that the pointers within the `unsafe` block will be valid pointers to the slice. This is an acceptable and appropriate use of `unsafe`.

Note that we don’t need to mark the resulting slices as `unsafe`. We can call this function from safe Rust. We’ve called the `split_at_mut` function with an implementation of the function, because it creates only valid pointers from the slice.

In contrast, the use of `slice::from_raw_parts_mut` is unsafe.



when the slice is used. This code takes an arbitrary slice 10,000 items long.

```
use std::slice;

let address = 0x01234usize;
let r = address as *mut i32;

let slice : &[i32] = unsafe {
    slice::from_raw_parts_mut(r, 10000)
};
```

Listing 19-7: Creating a slice from an arbitrary memory

We don't own the memory at this arbitrary location, so the slice this code creates contains valid `i32` values, but though it's a valid slice results in undefined behavior. To align `address` to 4 (the alignment of `i32`), the `slice::from_raw_parts_mut` would already be required to always be aligned, even if they are not used (and

## Using `extern` Functions to Call External Code

Sometimes, your Rust code might need to interact with a foreign language. For this, Rust has a keyword, `extern`, to declare the signature of a *Foreign Function Interface (FFI)*. An FFI is a way to define functions and enable a different (foreign) language to call them. Rust can't check them, so responsible programmers ensure safety.

Listing 19-8 demonstrates how to set up an interface to the C standard library. Functions declared within the `extern "C"` block can be called from Rust code. The reason is that other languages have different calling conventions, and Rust can't check them, so responsible programmers ensure safety.

Filename: `src/main.rs`

```
extern "C" {
    fn abs(input: i32) -> i32;
}

fn main() {
    unsafe {
        println!("Absolute value of -3 acc")
    }
}
```

Listing 19-8: Declaring and calling an `extern` function

Within the `extern "C"` block, we list the names from another language we want to call. The `"C"` *interface (ABI)* the external function uses: the ABI at the assembly level. The `"C"` ABI is the most common language's ABI.

---

## Calling Rust Functions from Other Languages

We can also use `extern` to create an interface to call Rust functions. Instead of an `extern` block, we specify the ABI to use just before the `fn` definition with the `#[no_mangle]` annotation to tell the Rust compiler not to mangle this function. *Mangling* is when a compiler changes a function to a different name that contains more information about the compilation process to consume but is less readable. Most programming language compilers mangle names to make a Rust function to be nameable by other language compilers' name mangling.

In the following example, we make the `call_from_c` code, after it's compiled to a shared library and

```
#[no_mangle]
pub extern "C" fn call_from_c() {
    println!("Just called a Rust function")
}
```

This usage of `extern` does not require `unsafe` code.

---

## Accessing or Modifying a Mutable Static Variable

Until now, we've not talked about *global variable* problematic with Rust's ownership rules. If two threads access a mutable global variable, it can cause a data race.

In Rust, global variables are called *static* variable. The declaration and use of a static variable with a string literal is as follows:

Filename: src/main.rs

```
static HELLO_WORLD: &str = "Hello, world!"

fn main() {
    println!("name is: {}", HELLO_WORLD);
}
```

Listing 19-9: Defining and using an immutable static variable

Static variables are similar to constants, which we discussed in the "Between Variables and Constants" section in Chapter 19. Static variables, which are in `SCREAMING_SNAKE_CASE` by convention, are of type `&'static str`, which is `&'static str` in this example. Static variables live for the entire lifetime of the program, which means they are references with the `'static` lifetime, which means they live for the entire lifetime; we don't need to annotate it explicitly because the variable is safe.

Constants and immutable static variables might seem similar, but one difference is that values in a static variable have a fixed address in memory and always access the same data. Constants, on the other hand, are not stored in memory; their data is calculated whenever they're used.

Another difference between constants and static variables is that static variables can be mutable. Accessing and modifying mutable static variables is discussed in the next section, which shows how to declare, access, and modify a mutable static variable.

Filename: src/main.rs

```

static mut COUNTER: u32 = 0;

fn add_to_count(inc: u32) {
    unsafe {
        COUNTER += inc;
    }
}

fn main() {
    add_to_count(3);

    unsafe {
        println!("COUNTER: {}", COUNTER);
    }
}

```

Listing 19-10: Reading from or writing to a mutable static

As with regular variables, we specify mutability with `mut`. Reads or writes from `COUNTER` must be within an `unsafe` block, and prints `COUNTER: 3` as we would expect because multiple threads accessing `COUNTER` would likely race.

With mutable data that is globally accessible, it's easy to have races, which is why Rust considers mutable statics `unsafe`. If possible, it's preferable to use the concurrency tools or pointers we discussed in Chapter 16 so the communication between different threads is done safely.

## Implementing an Unsafe Trait

The final action that works only with `unsafe` is implementing a trait. A trait is `unsafe` when at least one of its methods has some `unsafe` code. We can declare that a trait is `unsafe` by a `unsafe trait` and marking the implementation of the trait with `unsafe impl`. Listing 19-11.

```

unsafe trait Foo {
    // methods go here
}

unsafe impl Foo for i32 {
    // method implementations go here
}

```

Listing 19-11: Defining and implementing an unsafe

By using `unsafe impl`, we're promising that we're doing what the compiler can't verify.

As an example, recall the `Sync` and `Send` markers from “Extensible Concurrency with the `Sync` and `Send` Traits”. The compiler implements these traits automatically for `Send` and `Sync` types. If we implement a type that is `Send` or `Sync`, such as raw pointers, and we want to make it unsafe, we must use `unsafe`. Rust can't verify that our type can be safely sent across threads or accessed from multiple threads, so we have to do those checks manually and indicate as such.

## When to Use Unsafe Code

Using `unsafe` to take one of the four actions (such as `ptr::read`) is not even frowned upon. But it is trickier to get the compiler to help uphold memory safety. When you write unsafe code, you can do so, and having the explicit `unsafe` keyword track down the source of problems if they occur.

## Advanced Lifetimes

In Chapter 10 in the “Validating References with Lifetimes” section, we saw how to annotate references with lifetime parameters to show how references relate. You saw how every reference has a lifetime. Rust will let you elide lifetimes. Now we'll look at some advanced lifetime rules that we haven't covered yet:

- Lifetime subtyping: ensures that one lifetime is shorter than another
- Lifetime bounds: specifies a lifetime for a reference
- Inference of trait object lifetimes: allows the compiler to infer lifetimes and when they need to be specified
- The anonymous lifetime: making elision more explicit

## Ensuring One Lifetime Outlives Another

*Lifetime subtyping* specifies that one lifetime should outlive another.

explore lifetime subtyping, imagine we want to write a function called `Context` that holds a reference to the string that will parse this string and return success or failure. The `Context` to do the parsing. Listing 19-12 implements this code, but it doesn't have the required lifetime annotations.

Filename: `src/lib.rs`

```
struct Context(&str);

struct Parser {
    context: &Context,
}

impl Parser {
    fn parse(&self) -> Result<(), &str> {
        Err(&self.context.0[1..])
    }
}
```

Listing 19-12: Defining a parser without lifetime annotations

Compiling the code results in errors because Rust requires a lifetime annotation on the string slice in `Context` and the reference to a `Context` in `Parser`.

For simplicity's sake, the `parse` function returns a `Result`. The `parse` function will do nothing on success and, on failure, it will return an error slice that didn't parse correctly. A real implementation would return a structured data type containing error information and would return a structured data type. We won't be discussing those details because they are beyond the scope of this example.

To keep this code simple, we won't write any parsing logic. Somewhere in the parsing logic we would have a function that returns an error that references the part of the input that is invalid. The code example interesting in regard to lifetime subtyping is that the input is invalid after the first byte. The first byte is not on a valid character boundary. We're using this example to focus on the lifetimes involved.

To get this code to compile, we need to fill in the lifetime annotations. The straightforward way to do this is to use the same lifetime parameter in Listing 19-13. Recall from the "Lifetime Annotations" section in Chapter 10 that each of `struct Context<'a>`, `struct Parser<'a>`, and `impl Parser<'a>` is declaring a new lifetime parameter. While their lifetimes are independent, they are all related to the lifetime of the input string.

three lifetime parameters declared in this exam

Filename: src/lib.rs

```
struct Context<'a>(&'a str);

struct Parser<'a> {
    context: &'a Context<'a>,
}

impl<'a> Parser<'a> {
    fn parse(&self) -> Result<(), &str> {
        Err(&self.context.0[1..])
    }
}
```

Listing 19-13: Annotating all references in `Context` parameters

This code compiles just fine. It tells Rust that a `Parser` holds a `Context` with lifetime `'a` and that `Context` holds a reference to the `Context` in `Parser`. Rust requires lifetime parameters for these references.

Next, in Listing 19-14, we'll add a function that takes a `Parser` to parse that context, and returns what it finds. It's a bit of work.

Filename: src/lib.rs

```
fn parse_context(context: Context) -> Result<(), &str> {
    Parser { context: &context }.parse()
}
```

Listing 19-14: An attempt to add a `parse_context` function that uses a `Parser`

We get two verbose errors when we try to compile the `parse_context` function:

```

error[E0597]: borrowed value does not live long enough
--> src/lib.rs:14:5
   |
14 |         Parser { context: &context }.parse
   |         ^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^ does not live long enough
15 |     }
   |     - temporary value only lives until here
   |
note: borrowed value must be valid for the duration of the function
on the function body at 13:1...
--> src/lib.rs:13:1
   |
13 | / fn parse_context(context: Context)
14 | |     Parser { context: &context }.parse
15 | | }
   | |_^

```

```

error[E0597]: `context` does not live long enough
--> src/lib.rs:14:24
   |
14 |         Parser { context: &context }.parse
   |                        ^^^^^^^^^^ does not live long enough
15 |     }
   |     - borrowed value only lives until here
   |
note: borrowed value must be valid for the duration of the function
on the function body at 13:1...
--> src/lib.rs:13:1
   |
13 | / fn parse_context(context: Context)
14 | |     Parser { context: &context }.parse
15 | | }
   | |_^

```

These errors state that the `Parser` instance that is created as a function parameter live only until the end of the `parse_context` function. We need `Parser` to live for the entire lifetime of the function.

In other words, `Parser` and `context` need to be valid before the function starts as well as after it ends. The `Parser` we're creating and `context` are both in scope at the end of the function, because `parse_context` returns `Result`.

To figure out why these errors occur, let's look at the signature of the `parse` function, specifically the references in the signature of the `Parser` struct.

```
fn parse(&self) -> Result<(), &str> {
```



Remember the elision rules? If we annotate the eliding, the signature would be as follows:

```
fn parse<'a>(&'a self) -> Result<(), &
```

That is, the error part of the return value of `parse` is the lifetime of the `Parser` instance (that of `&self` in the original signature) makes sense: the returned string slice reference is an instance held by the `Parser`, and the definition of the lifetime of the reference to `Context` and the `Context` holds should be the same.

The problem is that the `parse_context` function is annotated with `parse`, so the lifetime of the return value of `parse` is tied to the `Parser` as well. But the `Parser` instance created in `main` won't live past the end of the function (it's temporary) and its scope at the end of the function (`parse_context`).

Rust thinks we're trying to return a reference to a temporary at the end of the function, because we annotated all the parameters with `parse`. The annotations told Rust the lifetime of the return value holds is the same as that of the lifetime of the parameters.

The `parse_context` function can't see that with the `Parser` returned will outlive `Context` and `Parser` and that the `Context` returned refers to the string slice, not to `Context`.

By knowing what the implementation of `parse` is, we can tell Rust that the return value of `parse` is tied to the `Parser` instance's `Context`, which is referencing the lifetime of the string slice that `parse_context` returns. We can tell Rust that the string slice in `Context` and the `Context` in `Parser` have different lifetimes and that the return value of `parse` has the lifetime of the string slice in `Context`.

First, we'll try giving `Parser` and `Context` different lifetimes. Listing 19-15. We'll use `'s` and `'c` as lifetime parameters. The `'s` lifetime goes with the string slice in `Context` and the `'c` lifetime goes with the `Context` in `Parser`. Note that this solution won't work. We'll look at why this fix isn't sufficient with Listing 19-16.

Filename: src/lib.rs

```

struct Context<'s>(&'s str);

struct Parser<'c, 's> {
    context: &'c Context<'s>,
}

impl<'c, 's> Parser<'c, 's> {
    fn parse(&self) -> Result<(), &'s str>
        Err(&self.context.0[1..])
}

fn parse_context(context: Context) -> Result<'c>
    Parser { context: &context }.parse()
}

```

Listing 19-15: Specifying different lifetime parameters and to `Context`

We've annotated the lifetimes of the references annotated them in Listing 19-13. But this time we're depending on whether the reference goes with the lifetime of the struct or the lifetime of the string slice. We've also added an annotation to the string slice to indicate that it goes with the lifetime of the struct.

When we try to compile now, we get the following error:

```

error[E0491]: in type `&'c Context<'s>`, the lifetime `'c` outlives the data it references
--> src/lib.rs:4:5
   |
4 |         context: &'c Context<'s>,
   |                   ^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^
note: the pointer is valid for the lifetime `'c` because it is borrowed from
3:1
--> src/lib.rs:3:1
   |
3 | / struct Parser<'c, 's> {
4 | |     context: &'c Context<'s>,
5 | | }
   | |_^
note: but the referenced data is only valid for the lifetime `'s`
on the struct at 3:1
--> src/lib.rs:3:1
   |
3 | / struct Parser<'c, 's> {
4 | |     context: &'c Context<'s>,
5 | | }
   | |_^

```

Rust doesn't know of any relationship between referenced data in `Context` with lifetime `'s` ne that it lives longer than the reference with lifetin the reference to `Context` might not be valid.

Now we get to the point of this section: the Rust that one lifetime parameter lives at least as long where we declare lifetime parameters, we can d declare a lifetime `'b` that lives at least as long a syntax `'b: 'a` .

In our definition of `Parser` , to say that `'s` (the guaranteed to live at least as long as `'c` (the life we change the lifetime declarations to look like t

Filename: src/lib.rs

```
struct Parser<'c, 's: 'c> {  
    context: &'c Context<'s>,  
}
```

Now the reference to `Context` in the `Parser` ar the `Context` have different lifetimes; we've ensi slice is longer than the reference to the `Context`

That was a very long-winded example, but as we chapter, Rust's advanced features are very speci we described in this example, but in such situati something and give it the necessary lifetime.

## Lifetime Bounds on References to Gen

In the "Trait Bounds" section in Chapter 10, we d generic types. We can also add lifetime paramet these are called *lifetime bounds*. Lifetime bounds generic types won't outlive the data they're refer

As an example, consider a type that is a wrapper `RefCell<T>` type from the "`RefCell<T>`" and the in Chapter 15: its `borrow` and `borrow_mut` meth `RefMut` , respectively. These types are wrappers borrowing rules at runtime. The definition of the without lifetime bounds for now.

Filename: src/lib.rs

```
struct Ref<'a, T>(&'a T);
```

Listing 19-16: Defining a struct to wrap a reference with lifetime bounds

Without explicitly constraining the lifetime `'a` in `Ref`, Rust will error because it doesn't know how long

```
error[E0309]: the parameter type `T` may not live long enough
--> src/lib.rs:1:19
   |
1 | struct Ref<'a, T>(&'a T);
   |                   ^^^^^^^
   |
   = help: consider adding an explicit lifetime bound
note: ...so that the reference type `&'a T` points at the data it points at
--> src/lib.rs:1:19
   |
1 | struct Ref<'a, T>(&'a T);
   |                   ^^^^^^^
```

Because `T` can be any type, `T` could be a reference type. If `T` is a reference type, each of which could have their own lifetime, as long as `'a`.

Fortunately, the error provides helpful advice on how to fix this case:

```
consider adding an explicit lifetime bound
type
`&'a T` does not outlive the data it points to
```

Listing 19-17 shows how to apply this advice by adding a lifetime bound to `T`: we declare the generic type `T` to be `'a`.

```
struct Ref<'a, T: 'a>(&'a T);
```

Listing 19-17: Adding lifetime bounds on `T` to ensure it lives at least as long as `'a`

This code now compiles because the `T: 'a` syntax ensures that `T` lives at least as long as `'a`, but if it contains any references, the references will be valid for at least as long as `'a`.

We could solve this problem in a different way, a `StaticRef` struct in Listing 19-18, by adding the means if `T` contains any references, they must live for the entire program.

```
struct StaticRef<T: 'static>(&'static T);
```

Listing 19-18: Adding a `'static` lifetime bound means only `'static` references or no references

Because `'static` means the reference must live for the entire program (because there are no references concerned about references living long enough, type that has no references and a type that has the same for determining whether or not a reference it refers to.

## Inference of Trait Object Lifetimes

In Chapter 17 in the “Using Trait Objects that Allow for Dynamic Dispatch” section, we discussed trait objects, consisting of a trait object that we use to use dynamic dispatch. We haven’t yet discussed implementing the trait in the trait object has a lifetime parameter. Listing 19-19 where we have a trait `Red` and a struct `Ball` that has a lifetime parameter (and thus has a lifetime parameter) and we want to use an instance of `Ball` as the trait object.

Filename: src/main.rs

```
trait Red { }

struct Ball<'a> {
    diameter: &'a i32,
}

impl<'a> Red for Ball<'a> { }

fn main() {
    let num = 5;

    let obj = Box::new(Ball { diameter: &num })
}
```

Listing 19-19: Using a type that has a lifetime parameter

This code compiles without any errors, even though the lifetimes involved in `obj`. This code works both with lifetimes and trait objects:

- The default lifetime of a trait object is `'static`
- With `&'a Trait` or `&'a mut Trait`, the default lifetime is `'a`
- With a single `T: 'a` clause, the default lifetime is `'a`
- With multiple clauses like `T: 'a`, there is no explicit lifetime.

When we must be explicit, we can add a lifetime bound `Box<dyn Red>` using the syntax `Box<dyn Red + 'a>` depending on whether the reference lives for the whole duration of the function or is bound by other bounds, the syntax adding a lifetime bound `Red` trait that has references inside the type must be the same as those references.

## The anonymous lifetime

Let's say that we have a struct that's a wrapper around a string:

```
struct StrWrap<'a>(&'a str);
```

We can write a function that returns one of these:

```
fn foo<'a>(string: &'a str) -> StrWrap<'a> {
    StrWrap(string)
}
```

But that's a lot of `'a`s! To cut down on some of the boilerplate, we can use the anonymous lifetime, `'_`, like this:

```
fn foo(string: &str) -> StrWrap<'_> {
    StrWrap(string)
}
```

The `'_` says "use the elided lifetime here." This `StrWrap` contains a reference, but we don't need to specify a lifetime for it.

It works in `impl` headers too; for example:

```
// verbose
impl<'a> fmt::Debug for StrWrap<'a> {

// elided
impl fmt::Debug for StrWrap<'_> {
```

Next, let's look at some other advanced features

## Advanced Traits

We first covered traits in the “Traits: Defining Shared Behavior” chapter, but as with lifetimes, we didn't discuss the more advanced features. Now that we're more comfortable with Rust, we can get into the nitty-gritty details.

### Specifying Placeholder Types in Trait Definitions

*Associated types* connect a type placeholder with a concrete type. Trait definitions can use these placeholder types in their signatures, and trait implementations will specify the concrete type to be used in their implementation. That way, we can define a trait without knowing exactly what those types are up front.

We've described most of the advanced features needed. Associated types are somewhere in the middle of the book, between the “Advanced Features” chapter and the “Error Handling” chapter. Other features explained in the rest of the book include `const` and `unsafe` traits, and other features discussed in this chapter.

One example of a trait with an associated type is the `Iterator` trait in the `std::iter` module. The associated type is named `Item`, and it represents the type implementing the `Iterator` trait. The `next` method of the `Iterator` trait and the `next` method of the `Iterator` trait is as shown in Listing 19-2.

```
pub trait Iterator {
    type Item;

    fn next(&mut self) -> Option<Self::Item>
}
```

Listing 19-20: The definition of the `Iterator` tra

The type `Item` is a placeholder type, and the `next` will return values of type `Option<Self::Item>`. `Item` will specify the concrete type for `Item`, and the `next` containing a value of that concrete type.

Associated types might seem like a similar concept to us to define a function without specifying what types associated types?

Let's examine the difference between the two. In Chapter 13 that implements the `Iterator` trait 13-21, we specified that the `Item` type was `u32`

Filename: src/lib.rs

```
impl Iterator for Counter {  
    type Item = u32;  
  
    fn next(&mut self) -> Option<Self::Item> {  
        // --snip--  
    }  
}
```

This syntax seems comparable to that of generic `Iterator` trait with generics, as shown in Listing 19-21.

```
pub trait Iterator<T> {  
    fn next(&mut self) -> Option<T>;  
}
```

Listing 19-21: A hypothetical definition of the `Iterator` trait

The difference is that when using generics, as in Listing 19-21, we have to specify the type `T` in each implementation; because we can't have a single implementation of `Iterator<String> for Counter` or any other type. In Listing 19-20, we have a single implementation of `Iterator` for `Counter`. In Listing 19-21, it can be implemented for a type multiple times. For each type of the generic type parameters each time. For `Counter`, we would have to provide type annotations for each implementation of `Iterator` we want to use.

With associated types, we don't need to annotate the trait on a type multiple times. In Listing 19-20 with associated types, we can only choose what the type of `Item` will be once. We can have one `impl Iterator for Counter`. We don't have to



iterator of `u32` values everywhere that we call `iter_u32`

## Default Generic Type Parameters and Default Values

When we use generic type parameters, we can specify a default generic type. This eliminates the need for implementing a concrete type if the default type works. The syntax for specifying a default generic type is `<PlaceholderType=ConcreteType>`.

A great example of a situation where this technique is useful is operator overloading. *Operator overloading* is customizing the behavior of the `+` operator in particular situations.

Rust doesn't allow you to create your own operators. But you can overload the operations and corresponding traits by implementing the traits associated with the operations. For example, to overload the `+` operator to add two `Point` instances, you can implement the `Add` trait on a `Point` struct:

Filename: src/main.rs

```
use std::ops::Add;

#[derive(Debug, PartialEq)]
struct Point {
    x: i32,
    y: i32,
}

impl Add for Point {
    type Output = Point;

    fn add(self, other: Point) -> Point {
        Point {
            x: self.x + other.x,
            y: self.y + other.y,
        }
    }
}

fn main() {
    assert_eq!(Point { x: 1, y: 0 } + Point { x: 3, y: 3 },
               Point { x: 4, y: 3 });
}
```

Listing 19-22: Implementing the `Add` trait to overload the `+` operator for `Point` instances

The `add` method adds the `x` values of two `Point` instances to create a new `Point`. The `AcOutput` that determines the type returned from

The default generic type in this code is within the

```
trait Add<RHS=Self> {  
    type Output;  
  
    fn add(self, rhs: RHS) -> Self::Output  
}
```

This code should look generally familiar: a trait with a type. The new part is `RHS=Self`: this syntax is called a generic type parameter (short for “right hand side parameter” in the `add` method. If we don’t specify a type parameter in the `add` method, the type of `RHS` will default to `Self`, which is the type we’re implementing `Add` on.

When we implemented `Add` for `Point`, we used `Point` as the `RHS`. We wanted to add two `Point` instances. Let’s look at an example of the `Add` trait where we want to customize the `RHS`.

We have two structs, `Millimeters` and `Meters`, and we want to add values in millimeters to values in meters. We want the `Add` trait to do the conversion correctly. We can implement `Add` for `Meters` as the `RHS`, as shown in Listing 19-23.

Filename: `src/lib.rs`

```
use std::ops::Add;  
  
struct Millimeters(u32);  
struct Meters(u32);  
  
impl Add<Meters> for Millimeters {  
    type Output = Millimeters;  
  
    fn add(self, other: Meters) -> Millimeters {  
        Millimeters(self.0 + (other.0 * 1000))  
    }  
}
```

Listing 19-23: Implementing the `Add` trait on `Millimeters` and `Meters`

To add `Millimeters` and `Meters`, we specify in the `RHS` type parameter instead of using the default

You'll use default type parameters in two main ways:

- To extend a type without breaking existing code
- To allow customization in specific cases

The standard library's `Add` trait is an example of how to add two like types, but the `Add` trait provides the flexibility to use a default type parameter. Using a default type parameter in the `Add` trait allows you to specify the extra parameter most of the time. In this way, boilerplate isn't needed, making it easier to use.

The first purpose is similar to the second but in a different way. When you add a parameter to an existing trait, you can give it a default value. This way, the functionality of the trait without breaking the existing code.

## Fully Qualified Syntax for Disambiguating Same Name

Nothing in Rust prevents a trait from having a method with the same name as another trait's method, nor does Rust prevent you from having two methods with the same name on one type. It's also possible to implement a method with the same name as methods from traits.

When calling methods with the same name, you need to specify which one you want to use. Consider the code in Listing 19-24, which defines `Wizard` and `Human`, that both have a method called `fly`. The `Human` type already has a method named `fly`, but the `Human` method does something different.

Filename: `src/main.rs`

```

trait Pilot {
    fn fly(&self);
}

trait Wizard {
    fn fly(&self);
}

struct Human;

impl Pilot for Human {
    fn fly(&self) {
        println!("This is your captain speaking");
    }
}

impl Wizard for Human {
    fn fly(&self) {
        println!("Up!");
    }
}

impl Human {
    fn fly(&self) {
        println!("*waving arms furiously*");
    }
}

```

Listing 19-24: Two traits are defined to have a `fly` method, and the `Human` type, and a `fly` method is implemented for the `Human` type.

When we call `fly` on an instance of `Human`, the method that is directly implemented on the type is called.

Filename: src/main.rs

```

fn main() {
    let person = Human;
    person.fly();
}

```

Listing 19-25: Calling `fly` on an instance of `Human`.

Running this code will print `*waving arms furiously*`, which is the `fly` method implemented on `Human` directly.

To call the `fly` methods from either the `Pilot` or `Wizard` traits, we use more explicit syntax to specify which `fly` method we want to call.

demonstrates this syntax.

Filename: src/main.rs

```
fn main() {  
    let person = Human;  
    Pilot::fly(&person);  
    Wizard::fly(&person);  
    person.fly();  
}
```

Listing 19-26: Specifying which trait's `fly` method

Specifying the trait name before the method name in the implementation of `fly` we want to call. We could also write `person.fly()` which is equivalent to the `person.fly()` that we no longer need to write if we don't need to disambiguate.

Running this code prints the following:

```
This is your captain speaking.  
Up!  
*waving arms furiously*
```

Because the `fly` method takes a `self` parameter, Rust could figure out which trait to implement one *trait*, based on the type of `self`.

However, associated functions that are part of a trait are not associated with a type. When two types in the same scope implement the same trait, you need to specify the type you mean unless you use *fully qualified syntax*. Listing 19-27 has the associated function `baby_rattle` for the struct `Dog`, and the associated function

Filename: src/main.rs

```

trait Animal {
    fn baby_name() -> String;
}

struct Dog;

impl Dog {
    fn baby_name() -> String {
        String::from("Spot")
    }
}

impl Animal for Dog {
    fn baby_name() -> String {
        String::from("puppy")
    }
}

fn main() {
    println!("A baby dog is called a {}",
}

```

Listing 19-27: A trait with an associated function and a function of the same name that also implements the trait.

This code is for an animal shelter that wants to have a function implemented in the `baby_name` associated function. The `Dog` type also implements the trait `Animal`, which defines a `baby_name` function that all animals have. Baby dogs are called puppies, and the code implements the `Animal` trait on `Dog` in the `impl` block with the `Animal` trait.

In `main`, we call the `Dog::baby_name` function, which is defined on `Dog` directly. This code prints the following output:

```
A baby dog is called a Spot
```

This output isn't what we wanted. We want to call the `baby_name` function of the `Animal` trait that we implemented on `Dog`. The output is `A baby dog is called a puppy`. The technique used in Listing 19-26 doesn't help here; if we change Listing 19-28, we'll get a compilation error.

Filename: src/main.rs

```
fn main() {
    println!("A baby dog is called a {}"),
}
```

Listing 19-28: Attempting to call the `baby_name` Rust doesn't know which implementation to use

Because `Animal::baby_name` is an associated function thus doesn't have a `self` parameter, Rust can't `Animal::baby_name` we want. We'll get this com

```
error[E0283]: type annotations required: c
--> src/main.rs:20:43
|
20 |     println!("A baby dog is called a
|
|
= note: required by `Animal::baby_name`
```

To disambiguate and tell Rust that we want to use `Dog`, we need to use fully qualified syntax. Listing 19-29 shows the fully qualified syntax.

Filename: src/main.rs

```
fn main() {
    println!("A baby dog is called a {}"),
}
```

Listing 19-29: Using fully qualified syntax to specify the `baby_name` function from the `Animal` trait as implemented by `Dog`

We're providing Rust with a type annotation with `Dog` so it knows we want to call the `baby_name` method from the `Dog` type. This code will now print what we want:

```
A baby dog is called a puppy
```

In general, fully qualified syntax is defined as follows:

```
<Type as Trait>::function(receiver_if_method)
```

For associated functions, there would not be a receiver. For methods, there would be a receiver. You could use fully qualified syntax for functions or methods. However, you're allowed to use the `Self` keyword for methods.

Rust can figure out from other information in the more verbose syntax in cases where there are the same name and Rust needs help to identify which

## Using Supertraits to Require One Trait Another Trait

Sometimes, you might need one trait to use and you need to rely on the dependent traits also being on is a *supertrait* of the trait you're implementing

For example, let's say we want to make an `OutlinePrint` method that will print a value for `Point` struct that implements `Display` to result `outline_print` on a `Point` instance that has `1` the following:

```
*****  
*           *  
* (1, 3) *  
*           *  
*****
```

In the implementation of `outline_print`, we will use `Display` functionality. Therefore, we need to specify that for types that also implement `Display` and provide `OutlinePrint` needs. We can do that in the trait `OutlinePrint: Display`. This technique is similar to Listing 19-30 shows an implementation of the `OutlinePrint`

Filename: src/main.rs



```

use std::fmt;

trait OutlinePrint: fmt::Display {
    fn outline_print(&self) {
        let output = self.to_string();
        let len = output.len();
        println!("{}", "*".repeat(len + 4));
        println!("*{}*", " ".repeat(len + 1));
        println!("* {} *", output);
        println!("*{}*", " ".repeat(len + 1));
        println!("{}", "*".repeat(len + 4));
    }
}

```

Listing 19-30: Implementing the `OutlinePrint` trait from `Display`

Because we've specified that `OutlinePrint` requires the `to_string` function that is automatically implemented by `Display`. If we tried to use `to_string` without specifying the `Display` trait after the trait name, the compiler would find a method named `to_string` was found for the type.

Let's see what happens when we try to implement `OutlinePrint` for a type that doesn't implement `Display`, such as the `Point` struct.

Filename: src/main.rs

```

struct Point {
    x: i32,
    y: i32,
}

impl OutlinePrint for Point {}

```

We get an error saying that `Display` is required for `OutlinePrint`:

```

error[E0277]: the trait bound `Point: std::fmt::Display` is not satisfied
  --> src/main.rs:20:6
   |
20 | impl OutlinePrint for Point {}
   |          ^^^^^^^^^^^^^^^^^ `Point` cannot be implemented as `std::fmt::Display`;
   |          try using `:?` instead if you are using a variable
   |
   = help: the trait `std::fmt::Display` is not implemented for `Point`

```

To fix this, we implement `Display` on `Point` and `OutlinePrint` requires, like so:

Filename: src/main.rs

```
use std::fmt;

impl fmt::Display for Point {
    fn fmt(&self, f: &mut fmt::Formatter)
        write!(f, "({}, {})", self.x, self.y)
    }
}
```

Then implementing the `OutlinePrint` trait on `Point`, we can call `outline_print` on a `Point` instance and see the output with asterisks.

## Using the Newtype Pattern to Implement Traits on Types

In Chapter 10 in the “Implementing a Trait on a Type” section, we discussed the orphan rule that states we’re allowed to implement a trait if either the trait or the type are local to our crate. It’s possible to work around this rule by using the *newtype pattern*, which involves creating a new type that wraps the original type (covered tuple structs in the “Using Tuple Structs to Create Different Types” section of Chapter 5.) The tuple struct wraps the type we want to implement the trait on, and the trait implementation is local to our crate, and we can implement the trait without incurring the performance penalty for using this pattern, and without violating the orphan rule.

As an example, let’s say we want to implement `Display` on a `Vec` type. The orphan rule prevents us from doing directly because `Vec` and `Display` are defined outside our crate. We can create a new type `VecWrapper` that is an instance of `Vec`; then we can implement `Display` on `VecWrapper` value, as shown in Listing 19-31.

Filename: src/main.rs

```

use std::fmt;

struct Wrapper(Vec<String>);

impl fmt::Display for Wrapper {
    fn fmt(&self, f: &mut fmt::Formatter)
        write!(f, "[{}]", self.0.join(", '
    }
}

fn main() {
    let w = Wrapper(vec![String::from("hel
println!("w = {}", w);
}

```

Listing 19-31: Creating a `Wrapper` type around `Vec`

The implementation of `Display` uses `self.0` to access the inner `Vec`. Because `Wrapper` is a tuple struct and `Vec<T>` is the inner type, we can use the functionality of the `Display` type on `Wrapper`.

The downside of using this technique is that `Wrapper` doesn't have the methods of the value it's holding. We would have to implement `Vec<T>` directly on `Wrapper` such that the methods of `Vec` allow us to treat `Wrapper` exactly like a `Vec<T>`. Instead of implementing every method the inner type has, implementing `Deref` (see 15 in the “Treating Smart Pointers like Regular References” section) on the `Wrapper` to return the inner type would allow the `Wrapper` type to have all the methods of the inner type. If we wanted the `Wrapper` type's behavior—we would have to implement it manually.

Now you know how the newtype pattern is used in Rust. We'll see the pattern even when traits are not involved. Let's look at some advanced ways to interact with Rust's type system.

## Advanced Types

The Rust type system has some features that we haven't yet discussed. We'll start by discussing newtypes and why newtypes are useful as types. Then we'll mention generics, which are similar to newtypes but with slightly different semantics. Finally, we'll discuss dynamically sized types.

---

Note: The next section assumes you've read the [Newtype Pattern to Implement External Traits on External Types](#) article.

---

## Using the Newtype Pattern for Type Safety

The newtype pattern is useful for tasks beyond statically enforcing that values are never confused. You saw an example of using newtypes to indicate the `Millimeters` and `Meters` structs wrapped a function with a parameter of type `Millimeter` that accidentally tried to call that function with a

Another use of the newtype pattern is in abstracting details of a type: the new type can expose a public interface while hiding the private inner type if we used the new type directly. For example,

Newtypes can also hide internal implementation details. For example, the `People` type to wrap a `HashMap<i32, String>` to represent people with their name. Code using `People` would only need to know about the public interface, such as a method to add a name string. The inner implementation wouldn't need to know that we assign an `i32` ID. The newtype pattern is a lightweight way to achieve encapsulation, which we discussed in the "Encapsulation that Hides Implementation" section of Chapter 17.

## Creating Type Synonyms with Type Aliases

Along with the newtype pattern, Rust provides the `type` keyword to create an alias for an existing type with another name. For this we use the `type` keyword. We can create the alias `Kilometers` to `i32` like so:

```
type Kilometers = i32;
```

Now, the alias `Kilometers` is a *synonym* for `i32`. The `Meters` types we created in Listing 19-23, `Kilometers` values that have the type `Kilometers` will be treated as `i32`.

```

type Kilometers = i32;

let x: i32 = 5;
let y: Kilometers = 5;

println!("x + y = {}", x + y);

```

Because `Kilometers` and `i32` are the same type and we can pass `Kilometers` values to functions. However, using this method, we don't get the type of the new type pattern discussed earlier.

The main use case for type synonyms is to reduce the length of a type. For example, we can have a lengthy type like this:

```
Box<dyn Fn() + Send + 'static>
```

Writing this lengthy type in function signatures and code can be tiresome and error prone. Imagine the code in Listing 19-32.

```

let f: Box<dyn Fn() + Send + 'static> = Box::new(|| {});

fn takes_long_type(f: Box<dyn Fn() + Send + 'static>) {
    // --snip--
}

fn returns_long_type() -> Box<dyn Fn() + Send + 'static> {
    // --snip--
}

```

Listing 19-32: Using a long type in many places

A type alias makes this code more manageable. In Listing 19-33, we've introduced an alias named `Thunk` for all uses of the type with the shorter alias `Thunk`.

```

type Thunk = Box<dyn Fn() + Send + 'static>

let f: Thunk = Box::new(|| println!("hi"))

fn takes_long_type(f: Thunk) {
    // --snip--
}

fn returns_long_type() -> Thunk {
    // --snip--
}

```

Listing 19-33: Introducing a type alias `Thunk` to

This code is much easier to read and write! Choosing an alias can help communicate your intent as well (as discussed in the next chapter), so it's an appropriate name.

Type aliases are also commonly used with the `Result` type to avoid repetition. Consider the `std::io` module in the standard library, which returns a `Result<T, Error>` to handle situations where a function can fail. The `std::io::Error` struct represents all possible errors, so functions in `std::io` will be returning `Result<T, std::io::Error>`, such as these functions in the

```

use std::io::Error;
use std::fmt;

pub trait Write {
    fn write(&mut self, buf: &[u8]) -> Result<(), Error>;
    fn flush(&mut self) -> Result<(), Error>;

    fn write_all(&mut self, buf: &[u8]) -> Result<(), Error>;
    fn write_fmt(&mut self, fmt: fmt::Arguments) -> Result<(), Error>;
}

```

The `Result<..., Error>` is repeated a lot. As such, the `std::io` module has a type alias declaration:

```

type Result<T> = Result<T, std::io::Error>

```

Because this declaration is in the `std::io` module, you can use `std::io::Result<T>`—that is, a `Result<T, Error>`—instead of `Result<T, std::io::Error>`. The `Write` trait function signature becomes

```
pub trait Write {
    fn write(&mut self, buf: &[u8]) -> Res
    fn flush(&mut self) -> Result<()>;

    fn write_all(&mut self, buf: &[u8]) ->
    fn write_fmt(&mut self, fmt: Arguments
}

```

The type alias helps in two ways: it makes code consistent across all of `std::io`. Because `Result<T, E>`, which means we can use any method with it, as well as special syntax like the `?` operator.

## The Never Type that Never Returns

Rust has a special type named `!` that's known as the never type because it has no values. We prefer to call it the never type in place of the return type when a function will never return.

```
fn bar() -> ! {
    // --snip--
}

```

This code is read as “the function `bar` returns a never type”. Functions that never return are called *diverging functions*. We can't create values of this type, so they're not possibly return.

But what use is a type you can never create values of? In Listing 19-25; we've reproduced part of it here in Listing 19-34.

```
let guess: u32 = match guess.trim().parse() {
    Ok(num) => num,
    Err(_) => continue,
};

```

Listing 19-34: A `match` with an arm that ends in `continue`

At the time, we skipped over some details in this “Control Flow Operator” section, we discussed the `continue` keyword. It has the same type. So, for example, the following code could be written as:

```
let guess = match guess.trim().parse() {
    Ok(_) => 5,
    Err(_) => "hello",
}
```

The type of `guess` in this code would have to be `u32` because it requires that `guess` have only one type. So what if we allowed to return a `u32` from one arm and a `String` from `continue` in Listing 19-34?

As you might have guessed, `continue` has a type of `!T`. When the type of `guess`, it looks at both match arms, the latter with a `!` value. Because `!` can never be a value, the type of `guess` is `u32`.

The formal way of describing this behavior is that `continue` coerces its value into any other type. We're allowed to use `continue` because `continue` doesn't return a value; instead, it loops back to the start of the loop, so in the `Err` case, we never assign a value to `guess`.

The never type is useful with the `panic!` macro, a function that we call on `Option<T>` values to provide a definition:

```
impl<T> Option<T> {
    pub fn unwrap(self) -> T {
        match self {
            Some(val) => val,
            None => panic!("called `Option::unwrap` on a `None` value"),
        }
    }
}
```

In this code, the same thing happens as in the `match` expression. `val` has the type `T` and `panic!` has the type `!T`. The expression is `T`. This code works because `panic!` never returns, so the program never reaches the `None` case, we won't be returning a value. The code is valid.

One final expression that has the type `!T` is a `loop` expression:

```
loop {
    print!("forever ");
    loop {
        print!("and ever ");
    }
}
```



Here, the loop never ends, so `!` is the value of `t` be true if we included a `break`, because the loop `break`.

## Dynamically Sized Types and the `SizeOf`

Due to Rust's need to know certain details, such as the value of a particular type, there is a corner of its type system the concept of *dynamically sized types*. Sometimes these types let us write code using values whose

Let's dig into the details of a dynamically sized type using `String` throughout the book. That's right, not `&str` can't know how long the string is until runtime, nor type `str`, nor can we take an argument of type `str` which does not work:

```
let s1: str = "Hello there!";  
let s2: str = "How's it going?";
```

Rust needs to know how much memory to allocate and all values of a type must use the same amount of space. If we write this code, these two `str` values would need different space. But they have different lengths: `s1` needs 15. This is why it's not possible to create a variable of type `str`.

So what do we do? In this case, you already know `s1` and `s2` are `&str` rather than a `str`. Recall that in Chapter 4, we said the slice data structure stores the address of the slice.

So although a `&T` is a single value that stores the address of a `T` located in memory, a `&str` is *two* values: the address of the slice and the size of the slice. We know the size of a `&str` value at compile time: it's always 16 bytes. In other words, we always know the size of a `&str`, no matter what. In general, this is the way in which dynamically sized types work: they have an extra bit of metadata that stores the size of the value. The rule of dynamically sized types is that we must always store them behind a pointer of some kind.

We can combine `str` with all kinds of pointers: `String`, `Vec<str>`, etc. In fact, you've seen this before but with a different trait. A trait is a dynamically sized type we can refer to by `dyn`.

Chapter 17 in the “Using Trait Objects that Allow” section, we mentioned that to use traits as trait pointers, such as `&dyn Trait` or `Box<dyn Trait>`.

To work with DSTs, Rust has a particular trait called `Sized` that checks whether or not a type’s size is known at compile time. This trait is implemented for everything whose size is known at compile time, and it implicitly adds a bound on `Sized` to every generic function definition like this:

```
fn generic<T>(t: T) {  
    // --snip--  
}
```

is actually treated as though we had written this:

```
fn generic<T: Sized>(t: T) {  
    // --snip--  
}
```

By default, generic functions will work only on types that are known at compile time. However, you can use the following restriction:

```
fn generic<T: ?Sized>(t: &T) {  
    // --snip--  
}
```

A trait bound on `?Sized` is the opposite of a trait bound on `Sized`; this as “`T` may or may not be `Sized`.” This syntax is used to restrict other traits.

Also note that we switched the type of the `t` parameter from `T` to `&T`; a type might not be `Sized`, so we need to use a reference to it when we’ve chosen a reference.

Next, we’ll talk about functions and closures!

## Advanced Functions and Closures

Finally, we’ll explore some advanced features related to functions, which include function pointers and returning closures.

## Function Pointers

We've talked about how to pass closures to functions to functions! This technique is useful if you've already defined rather than defining a new function. Function pointers will allow you to use functions as arguments. You can coerce to the type `fn` (with a lowercase f), not to a trait. The `fn` type is called a *function pointer*. The parameter is a function pointer is similar to that in Listing 19-35.

Filename: src/main.rs

```
fn add_one(x: i32) -> i32 {
    x + 1
}

fn do_twice(f: fn(i32) -> i32, arg: i32) -> i32 {
    f(arg) + f(arg)
}

fn main() {
    let answer = do_twice(add_one, 5);

    println!("The answer is: {}", answer);
}
```

Listing 19-35: Using the `fn` type to accept a function pointer

This code prints `The answer is: 12`. We specify `fn` is an `fn` that takes one parameter of type `i32` and returns a `i32`. `f` in the body of `do_twice`. In `main`, we can pass `add_one` as the first argument to `do_twice`.

Unlike closures, `fn` is a type rather than a trait, so you can use the type directly rather than declaring a generic type and using it as a trait bound.

Function pointers implement all three of the closure traits, so you can always pass a function pointer as an argument to a function that takes a closure. It's best to write functions using a generic type so your functions can accept either functions or closures.

An example of where you would want to only accept function pointers is when interfacing with external code that doesn't have closures. C doesn't have closures, so you would want to use function pointers.

As an example of where you could use either a `collect` function, let's look at a use of `map`. To use the `map` numbers into a vector of strings, we could use a

```
let list_of_numbers = vec![1, 2, 3];
let list_of_strings: Vec<String> = list_of_numbers
    .iter()
    .map(|i| i.to_string())
    .collect();
```

Or we could name a function as the argument to

```
let list_of_numbers = vec![1, 2, 3];
let list_of_strings: Vec<String> = list_of_numbers
    .iter()
    .map(ToString::to_string)
    .collect();
```

Note that we must use the fully qualified syntax `ToString::to_string` from the “Advanced Traits” section because there are multiple `to_string` functions in scope. Here, we're using the `to_string` function from `ToString` which the standard library has implemented for

Some people prefer this style, and some people prefer the other. Both are compiling to the same code, so use whichever style you prefer.

## Returning Closures

Closures are represented by traits, which means that in most cases where you might want to return a trait object type that implements the trait as the return value, you can't use closures because they don't have a concrete type. However, you are not allowed to use the function pointer `fn` as a return type either.

The following code tries to return a closure directly:

```
fn returns_closure() -> Fn(i32) -> i32 {
    |x| x + 1
}
```

The compiler error is as follows:

```

error[E0277]: the trait bound `std::ops::Fn
std::marker::Sized` is not satisfied
-->
|
1 | fn returns_closure() -> Fn(i32) -> i32
|                                     ^^^^^^^^^^^^^^^^^^^
'static`
does not have a constant size known at c
|
= help: the trait `std::marker::Sized` i
`std::ops::Fn(i32) -> i32 + 'static`
= note: the return type of a function mu

```

The error references the **Sized** trait again! Rust need to store the closure. We saw a solution to t trait object:

```

fn returns_closure() -> Box<dyn Fn(i32) ->
    Box::new(|x| x + 1)
}

```

This code will compile just fine. For more about 'Objects That Allow for Values of Different Types'

## Summary

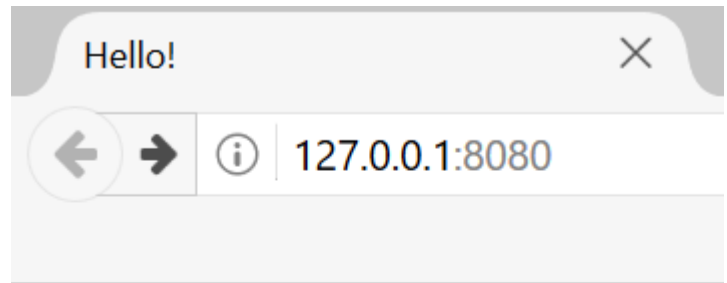
Whew! Now you have some features of Rust in y but you'll know they're available in very particul several complex topics so that when you encour suggestions or in other peoples' code, you'll be a syntax. Use this chapter as a reference to guide

Next, we'll put everything we've discussed throu one more project!

## Final Project: Building Web Server

It's been a long journey, but we've reached the e build one more project together to demonstrate the final chapters, as well as recap some earlier

For our final project, we'll make a web server that prints "Hello, World" in a web browser.



# Hello!

## Hi from Rust

Figure 20-1: Our final shared project

Here is the plan to build the web server:

1. Learn a bit about TCP and HTTP.
2. Listen for TCP connections on a socket.
3. Parse a small number of HTTP requests.
4. Create a proper HTTP response.
5. Improve the throughput of our server with

But before we get started, we should mention one thing: there are many crates available on <https://crates.io/> that provide more powerful pool implementations than we'll build.

However, our intention in this chapter is to help you learn the basics of Rust. Because Rust is a systems programming language, we want to work with low-level abstraction and can go to the practical in other languages. We'll write the basic code manually so you can learn the general ideas and how they might be used in the future.

## Building a Single-Threaded Web Server

We'll start by getting a single-threaded web server.

at a quick overview of the protocols involved in these protocols are beyond the scope of this book, but you have the information you need.

The two main protocols involved in web servers are *HTTP* and the *Transmission Control Protocol* (*TCP*), meaning a *client* initiates requests and the server provides a response to the client. The contents of the requests and responses are defined by the protocols.

TCP is the lower-level protocol that describes how to connect one server to another but doesn't specify what to do with the data of TCP by defining the contents of the requests and responses to use HTTP with other protocols, but in the vast majority of cases, data is sent over TCP. We'll work with the raw bytes of TCP responses.

## Listening to the TCP Connection

Our web server needs to listen to a TCP connection. The standard library offers a `std::net` module to create a new project in the usual fashion:

```
$ cargo new hello
   Created binary (application) `hello`
$ cd hello
```

Now enter the code in Listing 20-1 in `src/main.rs` to listen on address `127.0.0.1:7878` for incoming TCP streams. When a connection is established, it will print `Connection established!`.

Filename: `src/main.rs`

```
use std::net::TcpListener;

fn main() {
    let listener = TcpListener::bind("127.0.0.1:7878").unwrap();

    for stream in listener.incoming() {
        let stream = stream.unwrap();

        println!("Connection established!");
    }
}
```

## Listing 20-1: Listening for incoming streams and a stream

Using `TcpListener`, we can listen for TCP connections on `127.0.0.1:7878`. In the address, the section before the colon represents your computer (this is the same or represents the authors' computer specifically), and the section after the colon represents the port for two reasons: HTTP is normally accepted on a telephone.

The `bind` function in this scenario works like the `new TcpListener` instance. The reason the function is called `bind` is because it's like binding a telephone line to a port to listen to is known as binding.

The `bind` function returns a `Result<T, E>`, where `T` is the type of the stream and `E` is the error type. For example, connecting to port 80 requires administrative privileges (nonadministrators can listen only on ports higher than 1024). If you try to connect to port 80 without being an administrator, you'll get an error. For example, binding wouldn't work if we ran two instances of the program listening to the same port. Because of security learning purposes, we won't worry about handling errors; we'll just use `unwrap` to stop the program if errors happen.

The `incoming` method on `TcpListener` returns an iterator of streams (more specifically, streams of type `TcpStream`). Each stream represents an open connection between the client and the server. The `incoming` method returns the full request and response process in which a client connects to the server, the server generates a response, and the server closes the connection. The `TcpStream` will read from itself to see what the client sent and produce our response to the stream. Overall, this `for` loop iterates over the incoming streams and produces a series of streams for us to handle.

For now, our handling of the stream consists of printing the stream to the console. The program will stop if the stream has any errors; if there are no errors, it will print the message. We'll add more functionality for the server in the next chapter. The reason we might receive errors from the `incoming` method is that we're not actually iterating over *connection attempts*. The connection might fail for many reasons, many of them operating system specific. For example, some systems have a limit to the number of simultaneous open connections; new connection attempts beyond that limit will fail. Some of the open connections are closed.

Let's try running this code! Invoke `cargo run` in the directory where the code is located.



127.0.0.1:7878 in a web browser. The browser shows “Connection reset,” because the server isn’t running when you look at your terminal, you should see the message when the browser connected to the server!

```
Running `target/debug/hello`  
Connection established!  
Connection established!  
Connection established!
```

Sometimes, you’ll see multiple messages printed. This might be that the browser is making a request for other resources, like the *favicon.ico* icon that appears on the page.

It could also be that the browser is trying to connect because the server isn’t responding with any data. When the connection is dropped at the end of the loop, the connection is closed. Browsers sometimes deal with this because the problem might be temporary. The browser successfully gotten a handle to a TCP connection.

Remember to stop the program by pressing ctrl-c. Then restart `cargo run` to make sure you’re running the new code changes to make sure you’re running the new code.

## Reading the Request

Let’s implement the functionality to read the request. The concerns of first getting a connection and then handling the connection, we’ll start a new function for processing the connection. In the `handle_connection` function, we’ll read data from the connection and can see the data being sent from the browser. C20-2.

Filename: src/main.rs

```

use std::io::prelude::*;
use std::net::TcpStream;
use std::net::TcpListener;

fn main() {
    let listener = TcpListener::bind("127.0.0.1:8080");

    for stream in listener.incoming() {
        let stream = stream.unwrap();

        handle_connection(stream);
    }
}

fn handle_connection(mut stream: TcpStream) {
    let mut buffer = [0; 512];

    stream.read(&mut buffer).unwrap();

    println!("Request: {}", String::from_utf8_lossy(&buffer[0..stream.read().unwrap().len()]));
}

```

Listing 20-2: Reading from the `TcpStream` and printing the request

We bring `std::io::prelude` into scope to get `read` and `write` from `io::prelude`. In the `for` loop in `main`, we are printing a message that says we made a connection. We call the `handle_connection` function and pass the `stream` as an argument.

In the `handle_connection` function, we’ve made `mut` for `stream`. The reason is that the `TcpStream` instance keeps track of the data read internally. It might read more data than we asked for. It might read less data than we asked for. It therefore needs to be mutable; usually, we think of “reading” as not needing the `mut` keyword.

Next, we need to actually read from the stream. We declare a `buffer` on the stack to hold the data that we read. It is 512 bytes in size, which is big enough to hold the data that we need. It is sufficient for our purposes in this chapter. If we needed an arbitrary size, buffer management would need to be more complex. We pass the buffer to `stream.read`. `TcpStream` and put them in the buffer.

Second, we convert the bytes in the buffer to a `String`. The `String::from_utf8_lossy` function takes a `&[u8]` and returns a `String`. The “lossy” part of the name indicates the behavior of the function: it will convert the bytes to a `String` even if the bytes are not valid UTF-8. This is useful for displaying data that might not be valid UTF-8.

## U+FFFD REPLACEMENT CHARACTER

Note that we'll still get an error message

[illegible]

printing the request data, we can see why we get

program

```
headers CRLF
message-body
```

requesting. The first part of the request line indi

(URI) the client is requesting: a URI is almost, but *Resource Locator (URL)*. The difference between the two for the purposes in this chapter, but the HTTP spec uses URI. I will mentally substitute URL for URI here.

The last part is the HTTP version the client uses, *CRLF sequence*. (CRLF stands for *carriage return and line feed*!) The CRLF sequence can also be represented as `\r\n`. The CRLF sequence separates the request line from the rest of the request data. Note that when a line starts rather than `\r\n`.

Looking at the request line data we received from the browser, that `GET` is the method, `/` is the request URI, and

After the request line, the remaining lines starting with `GET` requests have no body.

Try making a request from a different browser client, such as `127.0.0.1:7878/test`, to see how the request data changes.

Now that we know what the browser is asking for,

## Writing a Response

Now we'll implement sending data in response to the following format:

```
HTTP-Version Status-Code Reason-Phrase CRLF
headers CRLF
message-body
```

The first line is a *status line* that contains the HTTP version, a numeric status code that summarizes the result of the request, and a text description of the status code. This is followed by headers, another CRLF sequence, and the body.

Here is an example response that uses HTTP version 1.1, a 200 OK reason phrase, no headers, and no body:

```
HTTP/1.1 200 OK\r\n\r\n
```

The status code 200 is the standard success response for an HTTP response. Let's write this to the stream as

From the `handle_connection` function, remove request data and replace it with the code in Listing 20-3.

Filename: `src/main.rs`

```
fn handle_connection(mut stream: TcpStream) {
    let mut buffer = [0; 512];

    stream.read(&mut buffer).unwrap();

    let response = "HTTP/1.1 200 OK\r\n\r\n";

    stream.write(response.as_bytes()).unwrap();
    stream.flush().unwrap();
}
```

Listing 20-3: Writing a tiny successful HTTP response

The first new line defines the `response` variable with the response data. Then we call `as_bytes` on our `response` to get a `&[u8]`. The `write` method on `stream` takes a `&[u8]` and sends the data to the connection.

Because the `write` operation could fail, we use `unwrap` to handle the error. Again, in a real application you would add error handling to prevent the program from continuing until the connection is closed. `TcpStream` contains an internal buffer and uses the operating system's `write` system call.

With these changes, let's run our code and make a request. We won't see any output data to the terminal, so we won't see any output from Cargo. When you load `127.0.0.1:7878` in a web browser, you'll see the response instead of an error. You've just hand-coded an HTTP response.

## Returning Real HTML

Let's implement the functionality for returning a response with a file, `hello.html`, in the root of your project directory. We'll use the `File` struct to read the file and return its contents as the response input any HTML you want; Listing 20-4 shows the code.

Filename: `hello.html`

```

<!DOCTYPE html>
<html lang="en">
  <head>
    <meta charset="utf-8">
    <title>Hello!</title>
  </head>
  <body>
    <h1>Hello!</h1>
    <p>Hi from Rust</p>
  </body>
</html>

```

Listing 20-4: A sample HTML file to return in a re

This is a minimal HTML5 document with a headi  
the server when a request is received, we'll mod  
Listing 20-5 to read the HTML file, add it to the r

Filename: src/main.rs

```

use std::fs;
// --snip--

fn handle_connection(mut stream: TcpStream) {
    let mut buffer = [0; 512];
    stream.read(&mut buffer).unwrap();

    let contents = fs::read_to_string("hel

    let response = format!("HTTP/1.1 200 C

    stream.write(response.as_bytes()).unwr
    stream.flush().unwrap();
}

```

Listing 20-5: Sending the contents of *hello.html* a

We've added a line at the top to bring the stand  
code for opening a file and reading the contents  
Chapter 12 when we read the contents of a file f

Next, we use `format!` to add the file's contents

Run this code with `cargo run` and load `127.0.0.`  
see your HTML rendered!

Currently, we're ignoring the request data in `bu`  
contents of the HTML file unconditionally. That r

127.0.0.1:7878/something-else in your browser, you get a 404 response. Our server is very limited and is not able to customize our responses depending on the request. We need to return a well-formed request to /.

## Validating the Request and Selectively

Right now, our web server will return the HTML file requested. Let's add functionality to check that the request is for the HTML file and return an error if the request is for something else. This we need to modify `handle_connection`, as shown in Listing 20-6. We'll check the content of the request received again like we did in Listing 20-5 and add `if` and `else` blocks to treat requests differently.

Filename: src/main.rs

```
// --snip--

fn handle_connection(mut stream: TcpStream) {
    let mut buffer = [0; 512];
    stream.read(&mut buffer).unwrap();

    let get = b"GET / HTTP/1.1\r\n";

    if buffer.starts_with(get) {
        let contents = fs::read_to_string("src/main.html").unwrap();

        let response = format!("HTTP/1.1 200 OK\r\nContent-Length: {}\r\n\r\n{}",
                                contents.len(), contents);

        stream.write(response.as_bytes()).unwrap();
        stream.flush().unwrap();
    } else {
        // some other request
    }
}
```

Listing 20-6: Matching the request and handling requests

First, we hardcode the data corresponding to the request for the HTML file. Because we're reading raw bytes into the buffer, we use the `b""` byte string syntax at the start of the `get` variable. We then check whether `buffer` starts with the bytes in `get`. If the request is a well-formed request to /, which is the success case, we return the contents of our HTML file.

If `buffer` does *not* start with the bytes in `get`, it's not a `GET` request. We'll add code to the `else` block in `main.rs` to handle other requests.

Run this code now and request `127.0.0.1:7878/y`. If you make any other request, such as `127.0.0.1:7878/`, you'll see a connection error like those you saw when running `20-2`.

Now let's add the code in Listing 20-7 to the `else` block. Status code 404, which signals that the content file was not found, also returns some HTML for a page to render in the browser to the end user.

Filename: `src/main.rs`

```
// --snip--

} else {
    let status_line = "HTTP/1.1 404 NOT FOUND";
    let contents = fs::read_to_string("404.html")
        .unwrap_or_else(|_| "404 Not Found".into());

    let response = format!("{}", status_line).as_bytes().concat(contents.as_bytes());

    stream.write(response.as_bytes()).unwrap();
    stream.flush().unwrap();
}
```

Listing 20-7: Responding with status code 404 and some HTML when a file was requested

Here, our response has a status line with status code 404 and the text `NOT FOUND`. We're still not returning headers, and we're returning the HTML in the file `404.html`. You'll need to create an `error.html` file; again feel free to use any HTML you want. See Listing 20-8.

Filename: `404.html`



```

<!DOCTYPE html>
<html lang="en">
  <head>
    <meta charset="utf-8">
    <title>Hello!</title>
  </head>
  <body>
    <h1>Oops!</h1>
    <p>Sorry, I don't know what you're ask</p>
  </body>
</html>

```

Listing 20-8: Sample content for the page to send

With these changes, run your server again. Request the contents of *hello.html*, and any other request will return the error HTML from *404.html*.

## A Touch of Refactoring

At the moment the `if` and `else` blocks have a lot of code for reading the files and writing the contents of the files to the status line and the filename. Let's make the code more concise by moving the differences into separate `if` and `else` lines that take the status line and the filename to variables; we can then use the code to read the file and write the response. We'll do this after replacing the large `if` and `else` blocks.

Filename: `src/main.rs`

```
// --snip--

fn handle_connection(mut stream: TcpStream) {
    // --snip--

    let (status_line, filename) = if buffer
        ("HTTP/1.1 200 OK\r\n\r\n", "hello")
    } else {
        ("HTTP/1.1 404 NOT FOUND\r\n\r\n",
        );

    let contents = fs::read_to_string(filename)
        .unwrap();

    let response = format!("{}", status_line, contents);

    stream.write(response.as_bytes()).unwrap();
    stream.flush().unwrap();
}
```

Listing 20-9: Refactoring the `if` and `else` block between the two cases

Now the `if` and `else` blocks only return the `status_line` and `filename` in a tuple; we then use destructuring to get `status_line` and `filename` using a pattern in the `let` statement in Listing 20-9. This is covered in Chapter 18.

The previously duplicated code is now outside the `if` and `else` blocks. This makes it easier to maintain between the two cases, and it means we have one place to change how the file reading and response code in Listing 20-9 will be the same as that in Listing 20-8.

Awesome! We now have a simple web server in Listing 20-10 that responds to one request with a page of content and to all other requests with a 404 response.

Currently, our server runs in a single thread, meaning it can only handle one request at a time. Let's examine how that can be a problem and how to fix it. Then we'll fix it so our server can handle multiple requests at once.

## Turning Our Single-Threaded Server into a Multithreaded Server

Right now, the server will process each request in a second connection until the first is finished processing and more requests, this serial execution would block. If it receives a request that takes a long time to process, it will wait until the long request is finished, even if there are other requests waiting to be processed quickly. We'll need to fix this, but first, we'll look at how to simulate a slow request.

## Simulating a Slow Request in the Current Server Implementation

We'll look at how a slow-processing request can be simulated in the current server implementation. Listing 20-10 implements a simulated slow response that will cause the server to wait before responding.

Filename: src/main.rs

```
use std::thread;
use std::time::Duration;
// --snip--

fn handle_connection(mut stream: TcpStream) {
    // --snip--

    let get = b"GET / HTTP/1.1\r\n";
    let sleep = b"GET /sleep HTTP/1.1\r\n";

    let (status_line, filename) = if buffer.starts_with(get) {
        ("HTTP/1.1 200 OK\r\n\r\n", "hello.html")
    } else if buffer.starts_with(sleep) {
        thread::sleep(Duration::from_secs(5));
        ("HTTP/1.1 200 OK\r\n\r\n", "hello.html")
    } else {
        ("HTTP/1.1 404 NOT FOUND\r\n\r\n", "")
    };

    // --snip--
}
```

Listing 20-10: Simulating a slow request by recognizing a request to /sleep and sleeping for 5 seconds

This code is a bit messy, but it's good enough for our purposes. The second request `sleep`, whose data our server receives after the `if` block to check for the request to `/sleep`, causes the server to sleep for 5 seconds before rendering a response.

You can see how primitive our server is: real libr multiple requests in a much less verbose way!

Start the server using `cargo run`. Then open two `http://127.0.0.1:7878/` and the other for `http://127.0.0.1:7878/sleep` URI a few times, as before, you'll see it respond ( then load `/`, you'll see that `/` waits until `sleep` has loading.

There are multiple ways we could change how o more requests back up behind a slow request; t pool.

## Improving Throughput with a Thread P

A *thread pool* is a group of spawned threads that task. When the program receives a new task, it a to the task, and that thread will process the task are available to handle any other tasks that com processing. When the first thread is done proces of idle threads, ready to handle a new task. A thi connections concurrently, increasing the throug

We'll limit the number of threads in the pool to a Denial of Service (DoS) attacks; if we had our pro request as it came in, someone making 10 millio havoc by using up all our server's resources and to a halt.

Rather than spawning unlimited threads, we'll h in the pool. As requests come in, they'll be sent t will maintain a queue of incoming requests. Each off a request from this queue, handle the request another request. With this design, we can proces is the number of threads. If each thread is respo subsequent requests can still back up in the que of long-running requests we can handle before i

This technique is just one of many ways to impro Other options you might explore are the fork/joi async I/O model. If you're interested in this topic solutions and try to implement them in Rust; wit these options are possible.

Before we begin implementing a thread pool, let's see what the code should look like. When you're trying to design code, the compiler can help guide your design. Write the API of the interface you want to call it; then implement the functionality. This is called compiler-driven development. We'll write the code to do what we want, and then we'll look at errors from the compiler to change the code next to get the code to work.

Similar to how we used test-driven development, we'll use compiler-driven development here. We'll write the code to do what we want, and then we'll look at errors from the compiler to change the code next to get the code to work.

## Code Structure If We Could Spawn a Thread for Each Stream

First, let's explore how our code might look if it could spawn a thread for each connection. As mentioned earlier, this isn't our final solution because it potentially spawning an unlimited number of threads. Listing 20-11 shows the changes to make to `main` to spawn a thread for each stream within the `for` loop.

Filename: `src/main.rs`

```
fn main() {
    let listener = TcpListener::bind("127.0.0.1:8080").unwrap();

    for stream in listener.incoming() {
        let stream = stream.unwrap();

        thread::spawn(|| {
            handle_connection(stream);
        });
    }
}
```

Listing 20-11: Spawning a new thread for each stream

As you learned in Chapter 16, `thread::spawn` works by taking a closure and running the code in the closure in the new thread. If you open a browser, then / in two more browser tabs, you'll see that the browser doesn't have to wait for `/sleep` to finish. But as we add more tabs, we overwhelm the system because you'd be making too many threads.

## Creating a Similar Interface for a Finite Number of Threads

We want our thread pool to work in a similar, fair way. A thread pool doesn't require large changes to the code, but it does require a change to the interface.

20-12 shows the hypothetical interface for a `ThreadPool` instead of `thread::spawn`.

Filename: `src/main.rs`

```
fn main() {
    let listener = TcpListener::bind("127.0.0.1:8080");
    let pool = ThreadPool::new(4);

    for stream in listener.incoming() {
        let stream = stream.unwrap();

        pool.execute(|| {
            handle_connection(stream);
        });
    }
}
```

Listing 20-12: Our ideal `ThreadPool` interface

We use `ThreadPool::new` to create a new thread pool with four threads, in this case four. Then, in the `for` loop, we use `pool.execute` instead of `thread::spawn` in that it takes a closure the pool needs to implement `pool.execute` so it takes the closure to the pool to run. This code won't yet compile, but we'll see how to fix it.

## Building the `ThreadPool` Struct Using Compile-Time Configuration

Make the changes in Listing 20-12 to `src/main.rs`, then run `cargo check` to drive our development.

```
$ cargo check
   Compiling hello v0.1.0 (file:///project)
error[E0433]: failed to resolve. Use of undeclared type `ThreadPool`
  --> src/main.rs:10:16
   |
10 |         let pool = ThreadPool::new(4);
   |                        ^^^^^^^^^^^^^^^^^ Use of undeclared type `ThreadPool`

error: aborting due to previous error
```

Great! This error tells us we need a `ThreadPool` struct now. Our `ThreadPool` implementation will be in `src/lib.rs` and the web server is doing. So, let's switch the `hello` crate to use `ThreadPool`.

crate to hold our `ThreadPool` implementation. We could also use the separate thread pool library `futures` for thread pool, not just for serving web requests.

Create a `src/lib.rs` that contains the following, with a `ThreadPool` struct that we can have for now:

Filename: src/lib.rs

```
pub struct ThreadPool;
```

Then create a new directory, `src/bin`, and move it into `src/bin/main.rs`. Doing so will make the library directory; we can still run the binary in `src/bin/main.rs` file, edit it to bring the library crate in by adding the following code to the top of `src/bin/main.rs`:

Filename: src/bin/main.rs

```
extern crate hello;
use hello::ThreadPool;
```

This code still won't work, but let's check it again  
address:

```
$ cargo check
    Compiling hello v0.1.0 (file:///project)
error[E0599]: no function or associated item found for struct `hello::ThreadPool` in the current scope
--> src/bin/main.rs:13:16
   |
13 |         let pool = ThreadPool::new(4);
   |                        ^^^^^^^^^^^^^^^^^^^^^^^^^^ function or associated item not found in `hello::ThreadPool`
in `hello::ThreadPool`
```

This error indicates that next we need to create for `ThreadPool`. We also know that `new` needs to accept `4` as an argument and should return a `Task`. The simplest `new` function that will have those c

Filename: src/lib.rs

```
pub struct ThreadPool;

impl ThreadPool {
    pub fn new(size: usize) -> ThreadPool
    {
        ThreadPool
    }
}
```

We chose `usize` as the type of the `size` parameter because a negative number of threads doesn't make any sense, and the number of elements in a collection of threads is a natural number, as discussed in the "Integer Types" section of the Rust book.

Let's check the code again:

```
$ cargo check
   Compiling hello v0.1.0 (file:///project)
warning: unused variable: `size`
  --> src/lib.rs:4:16
  |
4 |     pub fn new(size: usize) -> ThreadPool {
  |                   ^^^^^
  |
  = note: #[warn(unused_variables)] on by default
  = note: to avoid this warning, consider `size` as a mutable state
error[E0599]: no method named `execute` found for struct `ThreadPool` in the current scope
  --> src/bin/main.rs:18:14
  |
18 |         pool.execute(|| {
  |                   ^^^^^^^
  |
```

Now we get a warning and an error. Ignoring the warning for now, the error occurs because we don't have an `execute` method on `ThreadPool`. "Creating a Similar Interface for a Finite Number of Threads" says our thread pool should have an interface similar to `std::thread::JoinHandle`. We'll implement the `execute` function so it takes the closure to run on the idle thread in the pool to run.

We'll define the `execute` method on `ThreadPool`. Recall from the "Storing Closures Using Generic Parameters" section in Chapter 13 that we can take closures as arguments. We also know the traits: `Fn`, `FnMut`, and `FnOnce`. We need to decide which one to use. We know we'll end up doing something similar to `std::thread::spawn` in our implementation, so we can look at what bounds



on its parameter. The documentation shows us

```
pub fn spawn<F, T>(f: F) -> JoinHandle<T>
    where
        F: FnOnce() -> T + Send + 'static,
        T: Send + 'static
```

The `F` type parameter is the one we're concerned with, related to the return value, and we're not concerned with `T`. `spawn` uses `FnOnce` as the trait bound on `F`. That's because we'll eventually pass the argument we get to `spawn` further confident that `FnOnce` is the trait we want. When running a request will only execute that request once, the `Once` in `FnOnce`.

The `F` type parameter also has the trait bounds `'static`, which are useful in our situation: we need to move from one thread to another and `'static` because the thread will take to execute. Let's create an `execute` function that takes a generic parameter of type `F` with these bounds.

Filename: src/lib.rs

```
impl ThreadPool {
    // --snip--

    pub fn execute<F>(&self, f: F)
        where
            F: FnOnce() + Send + 'static
    {
        // ...
    }
}
```

We still use the `()` after `FnOnce` because this `F` has no parameters and doesn't return a value. Just like `FnOnce` type can be omitted from the signature, but even without it we need the parentheses.

Again, this is the simplest implementation of the `ThreadPool` but we're trying only to make our code compile.

```

$ cargo check
   Compiling hello v0.1.0 (file:///project)
warning: unused variable: `size`
--> src/lib.rs:4:16
4 |         pub fn new(size: usize) -> ThreadP
   |                                ^^^^^
   |
   = note: #[warn(unused_variables)] on by default
   = note: to avoid this warning, consider

warning: unused variable: `f`
--> src/lib.rs:8:30
8 |         pub fn execute<F>(&self, f: F)
   |                                ^
   |
   = note: to avoid this warning, consider

```

We're receiving only warnings now, which means we can `cargo run` and make a request in the browser, that we saw at the beginning of the chapter. Our closure passed to `execute` yet!

---

Note: A saying you might hear about language Haskell and Rust, is "if the code compiles, it works". This is not universally true. Our project compiles, but it's not building a real, complete project, this would be a good idea to add tests to check that the code compiles *and* has

---

## Validating the Number of Threads in `new`

We'll continue to get warnings because we aren't adding a `size` parameter to `new` and `execute`. Let's implement the bodies we want. To start, let's think about `new`. Earlier we had a `size` parameter, because a pool with a negative number of threads doesn't make sense. However, a pool with zero threads also makes no sense. We'll add code to check that `size` is greater than zero. We'll add a `ThreadPool` instance and have the program panic with `assert!` macro, as shown in Listing 20-13.

Filename: src/lib.rs

```

impl ThreadPool {
    /// Create a new ThreadPool.
    ///
    /// The size is the number of threads
    ///
    /// # Panics
    ///
    /// The `new` function will panic if 1
    pub fn new(size: usize) -> ThreadPool {
        assert!(size > 0);

        ThreadPool
    }

    // --snip--
}

```

Listing 20-13: Implementing `ThreadPool::new` to

We’ve added some documentation for our `ThreadPool` that we followed good documentation practices in situations in which our function can panic, as discussed in chapter 19. The docs for `ThreadPool` docs for `new` look like!

Instead of adding the `assert!` macro as we’ve done with `Config::new` in the `Config` module, we decided in this case that trying to create a thread should be an unrecoverable error. If you’re feeling ambitious, you can try the following signature to compare both versions:

```

pub fn new(size: usize) -> Result<ThreadPool, Error>

```

## Creating Space to Store the Threads

Now that we have a way to know we have a valid `ThreadPool`, we can create those threads and store them, returning it. But how do we “store” a thread? Let’s look at the `thread::spawn` signature:

```

pub fn spawn<F, T>(f: F) -> JoinHandle<T>
    where
        F: FnOnce() -> T + Send + 'static,
        T: Send + 'static

```

The `spawn` function returns a `JoinHandle<T>`, `v` returns. Let's try using `JoinHandle` too and see closures we're passing to the thread pool will have anything, so `T` will be the unit type `()`.

The code in Listing 20-14 will compile but doesn't change the definition of `ThreadPool` to hold a vector of instances, initialize the vector with a capacity of `size`, run some code to create the threads, and return them.

Filename: `src/lib.rs`

```
use std::thread;

pub struct ThreadPool {
    threads: Vec<thread::JoinHandle<()>>,
}

impl ThreadPool {
    // --snip--
    pub fn new(size: usize) -> ThreadPool {
        assert!(size > 0);

        let mut threads = Vec::with_capacity(size);

        for _ in 0..size {
            // create some threads and store them in the vector
        }

        ThreadPool {
            threads
        }
    }

    // --snip--
}
```

Listing 20-14: Creating a vector for `ThreadPool`

We've brought `std::thread` into scope in the lib and `thread::JoinHandle` as the type of the items in the vector.

Once a valid size is received, our `ThreadPool` creates `size` items. We haven't used the `with_capacity` method, but it performs the same task as `Vec::new` but with a capacity of `size` in the vector. Because we know we need to create `size` items, doing this allocation up front is slightly more efficient.

resizes itself as elements are inserted.

When you run `cargo check` again, you'll get a few more successes.

## A `Worker` Struct Responsible for Sending Code to Threads

We left a comment in the `for` loop in Listing 20-14. Here, we'll look at how we actually create threads using `thread::spawn` as a way to create threads, and how we code the thread should run as soon as the thread is created. We want to create the threads and have them *wait* for the results. The standard library's implementation of threads does this, but we have to implement it manually.

We'll implement this behavior by introducing a new struct, `ThreadPool`, and the threads that will manage the threads. We'll use the structure `Worker`, which is a common term in programming to refer to people working in the kitchen at a restaurant: they take orders from customers, and then they're responsible for preparing the food.

Instead of storing a vector of `JoinHandle<T>` in `ThreadPool`, we'll store instances of the `Worker` struct. Each `Worker` will have a `thread::spawn` instance. Then we'll implement a method on `Worker` to run and send it to the already running thread pool. We'll give each worker an `id` so we can distinguish between them when logging or debugging.

Let's make the following changes to what happens in Listing 20-14. We'll implement the code that sends the closure to the threads set up in this way:

1. Define a `Worker` struct that holds an `id` and a `thread::spawn` instance.
2. Change `ThreadPool` to hold a vector of `Worker` instead of `JoinHandle`.
3. Define a `Worker::new` function that takes a closure and returns an instance that holds the `id` and a thread spawn handle.
4. In `ThreadPool::new`, use the `for` loop to create `n` `Worker` with that `id`, and store the workers in the `ThreadPool`.

If you're up for a challenge, try implementing the `Worker` struct by looking at the code in Listing 20-15.

Ready? Here is Listing 20-15 with one way to make the `Worker` struct.

Filename: src/lib.rs

```
use std::thread;

pub struct ThreadPool {
    workers: Vec<Worker>,
}

impl ThreadPool {
    // --snip--
    pub fn new(size: usize) -> ThreadPool {
        assert!(size > 0);

        let mut workers = Vec::with_capacity(size);

        for id in 0..size {
            workers.push(Worker::new(id));
        }

        ThreadPool {
            workers
        }
    }
    // --snip--
}

struct Worker {
    id: usize,
    thread: thread::JoinHandle<()>,
}

impl Worker {
    fn new(id: usize) -> Worker {
        let thread = thread::spawn(|| {});

        Worker {
            id,
            thread,
        }
    }
}
```

Listing 20-15: Modifying `ThreadPool` to hold `Worker` threads directly

We've changed the name of the field on `ThreadPool` because it's now holding `Worker` instances instead of `Thread`. We also use the counter in the `for` loop as an argument to `new Worker` in the vector named `workers`.

External code (like our server in `src/bin/main.rs`) implementation details regarding using a `Worker` make the `Worker` struct and its `new` function provide the `id` we give it and stores a `JoinHandle<()>` a new thread using an empty closure.

This code will compile and will store the number as an argument to `ThreadPool::new`. But we're still getting in `execute`. Let's look at how to do that next.

## Sending Requests to Threads via Channels

Now we'll tackle the problem that the closures get nothing. Currently, we get the closure we want to execute. But we need to give `thread::spawn` a closure to execute during the creation of the `ThreadPool`.

We want the `Worker` structs that we just created to be held in the `ThreadPool` and send that code to it.

In Chapter 16, you learned about *channels*—a way to communicate between two threads—that would be perfect for this use case. The `execute` method will send a job to the `Worker` instances, which will send the job to its

1. The `ThreadPool` will create a channel and hold it.
2. Each `Worker` will hold on to the receiving end of the channel.
3. We'll create a new `Job` struct that will hold the closure and the channel.
4. The `execute` method will send the job to the sending side of the channel.
5. In its thread, the `Worker` will loop over its channel and execute the closures of any jobs it receives.

Let's start by creating a channel in `ThreadPool::new`. The `ThreadPool` instance, as shown in Listing 20, will do nothing for now but will be the type of item we

Filename: `src/lib.rs`

```

// --snip--
use std::sync::mpsc;

pub struct ThreadPool {
    workers: Vec<Worker>,
    sender: mpsc::Sender<Job>,
}

struct Job;

impl ThreadPool {
    // --snip--
    pub fn new(size: usize) -> ThreadPool {
        assert!(size > 0);

        let (sender, receiver) = mpsc::channel();

        let mut workers = Vec::with_capacity(size);

        for id in 0..size {
            workers.push(Worker::new(id));
        }

        ThreadPool {
            workers,
            sender,
        }
    }
    // --snip--
}

```

Listing 20-16: Modifying `ThreadPool` to store the `Job` instances

In `ThreadPool::new`, we create our new channel. This will successfully compile, still with warnings.

Let's try passing a receiving end of the channel in `new`. We know we want to use the `receiver` to create the channel. We know we want to use the `workers` to spawn, so we'll reference the `receiver` in Listing 20-17 won't quite compile yet.

Filename: `src/lib.rs`



```

impl ThreadPool {
    // --snip--
    pub fn new(size: usize) -> ThreadPool {
        assert!(size > 0);

        let (sender, receiver) = mpsc::channel();

        let mut workers = Vec::with_capacity(size);

        for id in 0..size {
            workers.push(Worker::new(id, receiver));
        }

        ThreadPool {
            workers,
            sender,
        }
    }
    // --snip--
}

// --snip--

impl Worker {
    fn new(id: usize, receiver: mpsc::Receiver<Job>) {
        let thread = thread::spawn(|| {
            for job in receiver {
                // ...
            }
        });

        Worker {
            id,
            thread,
        }
    }
}

```

Listing 20-17: Passing the receiving end of the channel to workers

We've made some small and straightforward changes to the code from Listing 20-16. We pass the channel into `Worker::new`, and then we use `Receiver` to receive jobs.

When we try to check this code, we get this error:

```

$ cargo check
   Compiling hello v0.1.0 (file:///project)
error[E0382]: use of moved value: `receiver`
  --> src/lib.rs:27:42
   |
27 |         workers.push(Worker::new(
   |         ^^^^^^^^^^^^^^^^^^^^^^^
   |         |
   |         previous iteration of loop
   |
   = note: move occurs because `receiver`
         has type `std::sync::mpsc::Receiver<Job>`, which
         implements the Drop trait

```

The code is trying to pass `receiver` to multiple workers, but you'll recall from Chapter 16: the channel implements the *multiple producer, single consumer*. This means we need to change the channel to fix this code. Even if we could, we wouldn't want to use it; instead, we want to distribute the jobs among all the workers.

Additionally, taking a job off the channel queue is a shared operation; the threads need a safe way to share and modify state under race conditions (as covered in Chapter 16).

Recall the thread-safe smart pointers discussed in Chapter 16: `Arc<Mutex<T>>`. The `Arc` type will let multiple threads share a value, and `Mutex` will ensure that only one worker gets a job from the channel at a time. This shows the changes we need to make.

Filename: src/lib.rs

```

use std::sync::Arc;
use std::sync::Mutex;
// --snip--

impl ThreadPool {
    // --snip--
    pub fn new(size: usize) -> ThreadPool {
        assert!(size > 0);

        let (sender, receiver) = mpsc::channel();

        let receiver = Arc::new(Mutex::new(receiver));

        let mut workers = Vec::with_capacity(size);

        for id in 0..size {
            workers.push(Worker::new(id, receiver.clone()));
        }

        ThreadPool {
            workers,
            sender,
        }
    }
    // --snip--
}

impl Worker {
    fn new(id: usize, receiver: Arc<Mutex<Receiver>>) -> Worker {
        {
            // --snip--
        }
    }
}

```

Listing 20-18: Sharing the receiving end of the channel and `Mutex`

In `ThreadPool::new`, we put the receiving end of the channel in a `Mutex`. For each new worker, we clone the `Arc` to bump the reference count so it can share ownership of the receiving end.

With these changes, the code compiles! We're good to go!

## Implementing the `execute` Method

Let's finally implement the `execute` method on `ThreadPool`, changing it from a struct to a type alias for a trait object that implements `Executor`.

`execute` receives. As discussed in the “Creating section of Chapter 19, type aliases allow us to m 20-19.

Filename: src/lib.rs

```
// --snip--

type Job = Box<dyn FnOnce() + Send + 'static>

impl ThreadPool {
    // --snip--

    pub fn execute<F>(&self, f: F)
        where
            F: FnOnce() + Send + 'static
    {
        let job = Box::new(f);

        self.sender.send(job).unwrap();
    }
}

// --snip--
```

Listing 20-19: Creating a `Job` type alias for a `Box` sending the job down the channel

After creating a new `Job` instance using the closure, we send the job down the sending end of the channel. We’re not sure that sending fails. This might happen if, for example, the receiver is not executing, meaning the receiving end has stopped. At that moment, we can’t stop our threads from executing. They’ll keep running as long as the pool exists. The reason we use `unwrap` is that it won’t happen, but the compiler doesn’t know that.

But we’re not quite done yet! In the worker, our `thread::spawn` still only *references* the receiving end of the channel to loop forever, asking the receiver to send a job when it gets one. Let’s make the `Worker::new` function.

Filename: src/lib.rs

```
// --snip--

impl Worker {
    fn new(id: usize, receiver: Arc<Mutex<
{
    let thread = thread::spawn(move ||
        loop {
            let job = receiver.lock().
                .unwrap();

            println!("Worker {} got a job", id);

            (*job)();
        }
    });

    Worker {
        id,
        thread,
    }
}
}
```

Listing 20-20: Receiving and executing the jobs in a worker

Here, we first call `lock` on the `receiver` to acquire the lock. We then call `unwrap` to panic on any errors. Acquiring a lock can fail if the lock is already held by another thread, which can happen if some other thread panics while holding the lock. In this situation, calling `unwrap` panics. The `lock` method returns `Err` if the receiving side shuts down.

If we get the lock on the mutex, we call `recv` to receive a job. The final `unwrap` moves past any errors here as well. The `recv` method returns `Err` if the receiving side shuts down.

The call to `recv` blocks, so if there is no job yet, the thread will wait until a job becomes available. The `Mutex<T>` ensures that only one thread can be trying to request a job.

Theoretically, this code should compile. Unfortunately, it doesn't, and we get this error:

```

error[E0161]: cannot move a value of type
std::marker::Send: the size of std::ops::Fn
cannot be
statically determined
--> src/lib.rs:63:17
   |
63 |                 (*job)();
   |                 ^^^^^^^

```

This error is fairly cryptic because the problem is a closure that is stored in a `Box<T>` (which is what needs to move itself *out* of the `Box<T>` because when we call it. In general, Rust doesn't allow us because Rust doesn't know how big the value in Chapter 15 that we used `Box<T>` precisely because size that we wanted to store in a `Box<T>` to get.

As you saw in Listing 17-15, we can write method `self: Box<dyn Self>`, which allows the method stored in a `Box<T>`. That's exactly what we want won't let us: the part of Rust that implements be implemented using `self: Box<dyn Self>`. So Rust could use `self: Box<dyn Self>` in this situation move the closure out of the `Box<T>`.

Rust is still a work in progress with places where in the future, the code in Listing 20-20 should work working to fix this and other issues! After you've you to join in.

But for now, let's work around this problem using explicitly that in this case we can take ownership `self: Box<dyn Self>`; then, once we have own This involves defining a new trait `FnBox` with the `self: Box<dyn Self>` in its signature, defining `FnOnce()`, changing our type alias to use the new the `call_box` method. These changes are shown

Filename: src/lib.rs

```

trait FnBox {
    fn call_box(self: Box<Self>);
}

impl<F: FnOnce()> FnBox for F {
    fn call_box(self: Box<F>) {
        (*self)()
    }
}

type Job = Box<dyn FnBox + Send + 'static>

// --snip--

impl Worker {
    fn new(id: usize, receiver: Arc<Mutex<
{
    let thread = thread::spawn(move ||
        loop {
            let job = receiver.lock().
                println!("Worker {} got a
                    job.call_box();
        }
    });

    Worker {
        id,
        thread,
    }
}
}

```

Listing 20-21: Adding a new trait `FnBox` to work with `Box<dyn FnOnce()>`

First, we create a new trait named `FnBox`. This trait is similar to the `call` methods on the other traits, which take ownership of `self: Box<dyn Self>` to take ownership of `self` and call the closure. The implementation of `call_box` uses the `Box<T>` and call the closure.

Next, we implement the `FnBox` trait for any type that implements the `FnOnce()` trait. Effectively, this means that any `FnOnce()` type can be boxed and called. The implementation of `call_box` uses the `Box<T>` and call the closure.

We now need our `Job` type alias to be a `Box` of `FnBox`. This will allow us to use `call_box` in

instead of invoking the closure directly. Implementing `FnOnce()` closure means we don't have to change the closure, we're sending down the channel. Now Rust is able to do it fine.

This trick is very sneaky and complicated. Don't worry, someday, it will be completely unnecessary.

With the implementation of this trick, our thread `cargo run` and make some requests:

```
$ cargo run
   Compiling hello v0.1.0 (file:///project)
warning: field is never used: `workers`
--> src/lib.rs:7:5
|
7 |         workers: Vec<Worker>,
|         ^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^
|
= note: #[warn(dead_code)] on by default

warning: field is never used: `id`
--> src/lib.rs:61:5
61 |         id: usize,
|         ^^^^^^^^^
|
= note: #[warn(dead_code)] on by default

warning: field is never used: `thread`
--> src/lib.rs:62:5
62 |         thread: thread::JoinHandle<()>,
|         ^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^
|
= note: #[warn(dead_code)] on by default

    Finished dev [unoptimized + debuginfo]
    Running `target/debug/hello`
Worker 0 got a job; executing.
Worker 2 got a job; executing.
Worker 1 got a job; executing.
Worker 3 got a job; executing.
Worker 0 got a job; executing.
Worker 2 got a job; executing.
Worker 1 got a job; executing.
Worker 3 got a job; executing.
Worker 0 got a job; executing.
Worker 2 got a job; executing.
```



Success! We now have a thread pool that executes requests. There are never more than four threads created if the server receives a lot of requests. If we make requests fast enough, we can be able to serve other requests by having another thread available.

Note that if you open `/sleep` in multiple browser windows and load 5 seconds apart from each other, because of the thread pool, you will see instances of the same request sequentially for a while, which is caused by our web server.

After learning about the `while let` loop in Chapter 19, we didn't write the worker thread code as shown in Listing 20-22.

Filename: `src/lib.rs`

```
// --snip--

impl Worker {
    fn new(id: usize, receiver: Arc<Mutex<Vec<Job>>>)
    {
        let thread = thread::spawn(move || {
            while let Ok(job) = receiver.recv() {
                println!("Worker {} got a job!", id);
                job.call_box();
            }
        });

        Worker {
            id,
            thread,
        }
    }
}
```

Listing 20-22: An alternative implementation of the worker thread

This code compiles and runs but doesn't result in a thread pool. A slow request will still cause other requests to wait. There is a somewhat subtle issue: the `Mutex` struct has no public `lock` method. The ownership of the lock is based on the lifetime of the `LockResult<MutexGuard<T>>` that the `lock` method returns. The borrow checker can then enforce the rule that a resource can't be accessed unless we hold the lock. But this implementation of the lock being held longer than intended if we don't drop the `MutexGuard<T>`. Because the values in the `Vec<Job>` are dropped at the duration of the block, the lock remains held for the duration of the block.

`job.call_box()` , meaning other workers cannot

By using `loop` instead and acquiring the lock `lock` outside it, the `MutexGuard` returned from the `lock` the `let job` statement ends. This ensures that `recv` , but it is released before the call to `job.call_box()` to be serviced concurrently.

## Graceful Shutdown and Clean

The code in Listing 20-21 is responding to requests of a thread pool, as we intended. We get some `thread` fields that we're not using in a direct way up anything. When we use the less elegant `ctrl-c` other threads are stopped immediately as well, not a request.

Now we'll implement the `Drop` trait to call `join` they can finish the requests they're working on before way to tell the threads they should stop accepting see this code in action, we'll modify our server to gracefully shutting down its thread pool.

## Implementing the `Drop` Trait on `ThreadPool`

Let's start with implementing `Drop` on our `ThreadPool` our threads should all join to make sure they finish first attempt at a `Drop` implementation; this code

Filename: `src/lib.rs`

```
impl Drop for ThreadPool {
    fn drop(&mut self) {
        for worker in &mut self.workers {
            println!("Shutting down worker");

            worker.thread.join().unwrap();
        }
    }
}
```

Listing 20-23: Joining each thread when the thread

First, we loop through each of the thread pool `w` because `self` is a mutable reference, and we all For each worker, we print a message saying that down, and then we call `join` on that worker's thread `unwrap` to make Rust panic and go into an ungraceful shutdown.

Here is the error we get when we compile this code:

```
error[E0507]: cannot move out of borrowed content
  --> src/lib.rs:65:13
   |
65 |         worker.thread.join().unwrap()
   |         ^^^^^^^ cannot move out of borrowed content
```

The error tells us we can't call `join` because we're borrowing `worker` and `join` takes ownership of its argument. To move the thread out of the `Worker` instance that we're borrowing, we need to consume the thread. We did this in Listing 17-15 by using `Option<thread::JoinHandle<>>` instead, we can use `Option` to move the value out of the `Some` variable. In other words, a `Worker` that is running a thread and when we want to clean up a `Worker`, we'll remove the thread. `Worker` doesn't have a thread to run.

So we know we want to update the definition of `Worker` to:

Filename: src/lib.rs

```
struct Worker {
    id: usize,
    thread: Option<thread::JoinHandle<>>
}
```

Now let's lean on the compiler to find the other errors. When we compile this code, we get two errors:

```

error[E0599]: no method named `join` found
  `std::option::Option<std::thread::JoinHandle<T>>`
  --> src/lib.rs:65:27
    |
65  |         worker.thread.join().unwr
    |                        ^^^^^
    |
error[E0308]: mismatched types
  --> src/lib.rs:89:13
    |
89  |         thread,
    |         ^^^^^^
    |
    |         expected enum `std::option::Option<std::thread::JoinHandle<T>>`
    |         help: try using a variant of the enum
    |
    = note: expected type `std::option::Option<std::thread::JoinHandle<T>>`
    found type `std::thread::JoinHandle<T>`

```

Let's address the second error, which points to the `thread` value. We need to wrap the `thread` value in `Some` when we create it. The following changes will fix this error:

Filename: src/lib.rs

```

impl Worker {
    fn new(id: usize, receiver: Arc<Mutex<Vec<String>>>)
    {
        // --snip--

        Worker {
            id,
            thread: Some(thread),
        }
    }
}

```

The first error is in our `Drop` implementation. We need to call `take` on the `Option` value to move `thread` out of the `Option`. The following changes will do so:

Filename: src/lib.rs

```
impl Drop for ThreadPool {
    fn drop(&mut self) {
        for worker in &mut self.workers {
            println!("Shutting down worker");

            if let Some(thread) = worker.thread {
                thread.join().unwrap();
            }
        }
    }
}
```

As discussed in Chapter 17, the `take` method on `Vec` removes the element and leaves `None` in its place. We're using `if let` to extract the `thread` from `worker`; then we call `join` on the thread. If a worker knows that worker has already had its thread cleaned up, it's a no-op case.

## Signaling to the Threads to Stop Listen

With all the changes we've made, our code compiles, but there's a bad news is this code doesn't function the way we want. The closures run by the threads of the `Worker` in `ThreadPool` will `join`, but that won't shut down the threads because they have jobs. If we try to drop our `ThreadPool` with our `main` thread will block forever waiting for the first thread to finish.

To fix this problem, we'll modify the threads so that they can signal that they should stop listening and exit their threads. In our `Worker` instances, our channel will send one of these two messages:

Filename: src/lib.rs

```
enum Message {
    NewJob(Job),
    Terminate,
}
```

This `Message` enum will either be a `NewJob` variant that tells a thread to run, or it will be a `Terminate` variant that tells a thread to stop.

We need to adjust the channel to use values of type `Message` as shown in Listing 20-24.

Filename: src/lib.rs

```

pub struct ThreadPool {
    workers: Vec<Worker>,
    sender: mpsc::Sender<Message>,
}

// --snip--

impl ThreadPool {
    // --snip--

    pub fn execute<F>(&self, f: F)
        where
            F: FnOnce() + Send + 'static
    {
        let job = Box::new(f);

        self.sender.send(Message::NewJob(j
    }

// --snip--

impl Worker {
    fn new(id: usize, receiver: Arc<Mutex<
    Worker {

        let thread = thread::spawn(move ||
            loop {
                let message = receiver.loc

                match message {
                    Message::NewJob(job) =>
                        println!("Worker {

                        job.call_box();
                    },
                    Message::Terminate =>
                        println!("Worker {

                        break;
                },
            }
        });

        Worker {
            id,
            thread: Some(thread),
        }
    }
}

```

Listing 20-24: Sending and receiving `Message` via `Message::Terminate`

To incorporate the `Message` enum, we need to change the definition of `ThreadPool` and the signature of the `send` method of `ThreadPool` needs to send jobs wrapped in a `Message`. Then, in `Worker::new` where a `Message` is received, we need to process it if the `NewJob` variant is received, and if the `Terminate` variant is received.

With these changes, the code will compile and run as it did after Listing 20-21. But we'll get a warning about the messages of the `Terminate` variety. Let's fix this implementation to look like Listing 20-25.

Filename: `src/lib.rs`

```
impl Drop for ThreadPool {
    fn drop(&mut self) {
        println!("Sending terminate message to all workers");

        for _ in &mut self.workers {
            self.sender.send(Message::Terminate).unwrap();
        }

        println!("Shutting down all workers");

        for worker in &mut self.workers {
            println!("Shutting down worker");

            if let Some(thread) = worker.thread.take() {
                thread.join().unwrap();
            }
        }
    }
}
```

Listing 20-25: Sending `Message::Terminate` to terminate each worker thread

We're now iterating over the workers twice: once to send the `Terminate` message to each worker and once to call `join` on each worker thread. The worker that receives the `Terminate` message and `join` immediately in the same location is the worker in the current iteration would be the one that was just created.

To better understand why we need two separate



workers. If we used a single loop to iterate through a terminate message would be sent down the channel to the first worker's thread. If that first worker was busy processing a request, the second worker would pick up the terminate message and terminate. We would be left waiting on the first worker to finish because the second thread picked up the terminate message.

To prevent this scenario, we first put all of our `take` calls in one loop; then we join on all the threads in another loop. This way, we are receiving requests on the channel once it gets a message, and we are sure that if we send the same number of terminate messages, each worker will receive a terminate message before it finishes its work.

To see this code in action, let's modify `main` to gracefully shut down the server, as shown in Listing 20-26.

Filename: `src/bin/main.rs`

```
fn main() {
    let listener = TcpListener::bind("127.0.0.1:8080");
    let pool = ThreadPool::new(4);

    for stream in listener.incoming().take(100) {
        let stream = stream.unwrap();

        pool.execute(|| {
            handle_connection(stream);
        });
    }

    println!("Shutting down.");
}
```

Listing 20-26: Shut down the server after serving requests

You wouldn't want a real-world web server to shut down after serving a few requests. This code just demonstrates that the graceful shutdown is possible in the right working order.

The `take` method is defined in the `Iterator` trait. It takes a closure that returns a `Result` of `T` or `Err`. The `ThreadPool` will go out of scope, and the `drop` implementation will run.

Start the server with `cargo run`, and make three requests. In your terminal you should see output like the following:

```

$ cargo run
   Compiling hello v0.1.0 (file:///project)
   Finished dev [unoptimized + debuginfo]
   Running `target/debug/hello`
Worker 0 got a job; executing.
Worker 3 got a job; executing.
Shutting down.
Sending terminate message to all workers.
Shutting down all workers.
Shutting down worker 0
Worker 1 was told to terminate.
Worker 2 was told to terminate.
Worker 0 was told to terminate.
Worker 3 was told to terminate.
Shutting down worker 1
Shutting down worker 2
Shutting down worker 3

```

You might see a different ordering of workers at this code works from the messages: workers 0 and 3 then on the third request, the server stopped and `ThreadPool` goes out of scope at the end of `main` and the pool tells all workers to terminate. The workers then see the terminate message, and then the thread for each worker thread.

Notice one interesting aspect of this particular example: the terminate messages go down the channel, and before the first message, we tried to join worker 0. Worker 0 hasn't received the message, so the main thread blocked waiting for it. Once each of the workers received the termination message, the main thread waited for the rest of the workers to receive the termination message and were able to finish.

Congrats! We've now completed our project; we've used a thread pool to respond asynchronously. We're also using a `ThreadPool` of the server, which cleans up all the threads in the background.

Here's the full code for reference:

Filename: `src/bin/main.rs`

```

extern crate hello;
use hello::ThreadPool;

use std::io::prelude::*;
use std::net::TcpListener;
use std::net::TcpStream;
use std::fs;
use std::thread;
use std::time::Duration;

fn main() {
    let listener = TcpListener::bind("127.0.0.1:8080");
    let pool = ThreadPool::new(4);

    for stream in listener.incoming().take(1000) {
        let stream = stream.unwrap();

        pool.execute(|| {
            handle_connection(stream);
        });
    }

    println!("Shutting down.");
}

fn handle_connection(mut stream: TcpStream) {
    let mut buffer = [0; 512];
    stream.read(&mut buffer).unwrap();

    let get = b"GET / HTTP/1.1\r\n";
    let sleep = b"GET /sleep HTTP/1.1\r\n";

    let (status_line, filename) = if buffer.starts_with(get) {
        ("HTTP/1.1 200 OK\r\n\r\n", "hello.txt")
    } else if buffer.starts_with(sleep) {
        thread::sleep(Duration::from_secs(5));
        ("HTTP/1.1 200 OK\r\n\r\n", "hello.txt")
    } else {
        ("HTTP/1.1 404 NOT FOUND\r\n\r\n", "")
    };

    let contents = fs::read_to_string(filename).unwrap();

    let response = format!("{}", status_line, contents);

    stream.write(response.as_bytes()).unwrap();
    stream.flush().unwrap();
}

```

Filename: src/lib.rs

```

use std::thread;
use std::sync::mpsc;
use std::sync::Arc;
use std::sync::Mutex;

enum Message {
    NewJob(Job),
    Terminate,
}

pub struct ThreadPool {
    workers: Vec<Worker>,
    sender: mpsc::Sender<Message>,
}

trait FnBox {
    fn call_box(self: Box<Self>);
}

impl<F: FnOnce()> FnBox for F {
    fn call_box(self: Box<F>) {
        (*self)()
    }
}

type Job = Box<dyn FnBox + Send + 'static>

impl ThreadPool {
    /// Create a new ThreadPool.
    ///
    /// The size is the number of threads
    ///
    /// # Panics
    ///
    /// The `new` function will panic if 1
    pub fn new(size: usize) -> ThreadPool {
        assert!(size > 0);

        let (sender, receiver) = mpsc::channel();

        let receiver = Arc::new(Mutex::new(receiver));

        let mut workers = Vec::with_capacity(size);

        for id in 0..size {
            workers.push(Worker::new(id, Arc::clone(&receiver)));
        }

        ThreadPool {
            workers,
            sender,
        }
    }
}

```

```

    }
}

pub fn execute<F>(&self, f: F)
    where
        F: FnOnce() + Send + 'static
    {
        let job = Box::new(f);

        self.sender.send(Message::NewJob(j
    }
}

impl Drop for ThreadPool {
    fn drop(&mut self) {
        println!("Sending terminate message

        for _ in &mut self.workers {
            self.sender.send(Message::Termin

        }

        println!("Shutting down all worker

        for worker in &mut self.workers {
            println!("Shutting down worker

                if let Some(thread) = worker.t
                    thread.join().unwrap();
                }
            }
        }
    }
}

struct Worker {
    id: usize,
    thread: Option<thread::JoinHandle<()>>
}

impl Worker {
    fn new(id: usize, receiver: Arc<Mutex<
        Worker {

            let thread = thread::spawn(move ||
                loop {
                    let message = receiver.loc

                    match message {
                        Message::NewJob(job) =
                            println!("Worker {

                                job.call_box();
                            },

```

```

        Message::Terminate =>
            println!("Worker {
                id,
            }
            break;
        },
    },
}

});

Worker {
    id,
    thread: Some(thread),
}
}
}

```

We could do more here! If you want to continue some ideas:

- Add more documentation to `ThreadPool`.
- Add tests of the library's functionality.
- Change calls to `unwrap` to more robust error handling.
- Use `ThreadPool` to perform some task other than printing.
- Find a thread pool crate on <https://crates.io> and use it instead. Then compare its API to the one we implemented.

## Summary

Well done! You've made it to the end of the book on this tour of Rust. You're now ready to implement your own projects. Keep in mind that there are other Rustaceans who would love to help you with your Rust journey.

## Appendix

The following sections contain reference material for your Rust journey.

### Appendix A: Keywords

The following list contains keywords that are reserved in the Rust language. As such, they cannot be used as identifiers, including names of functions, variables, parameters, constants, macros, static values, attributes, type names, etc.

## Keywords Currently in Use

The following keywords currently have the following functions:

- `as` - perform primitive casting, disambiguate an item, or rename items in `use` and `extern`
- `break` - exit a loop immediately
- `const` - define constant items or constant functions
- `continue` - continue to the next loop iteration
- `crate` - link an external crate or a macro version, which the macro is defined
- `else` - fallback for `if` and `if let` control flow
- `enum` - define an enumeration
- `extern` - link an external crate, function, or module
- `false` - Boolean false literal
- `fn` - define a function or the function pointer type
- `for` - loop over items from an iterator, implement a bounded lifetime
- `if` - branch based on the result of a condition
- `impl` - implement inherent or trait functions
- `in` - part of `for` loop syntax
- `let` - bind a variable
- `loop` - loop unconditionally
- `match` - match a value to patterns
- `mod` - define a module
- `move` - make a closure take ownership of a variable
- `mut` - denote mutability in references, raw pointers
- `pub` - denote public visibility in struct fields, modules
- `ref` - bind by reference
- `return` - return from function
- `Self` - a type alias for the type implementing a trait
- `self` - method subject or current module
- `static` - global variable or lifetime lasting the entire program
- `struct` - define a structure
- `super` - parent module of the current module
- `trait` - define a trait

- `true` - Boolean true literal
- `type` - define a type alias or associated type
- `unsafe` - denote unsafe code, functions, traits
- `use` - import symbols into scope
- `where` - denote clauses that constrain a type
- `while` - loop conditionally based on the result

## Keywords Reserved for Future Use

The following keywords do not have any function and have potential future use.

- `abstract`
- `alignof`
- `become`
- `box`
- `do`
- `final`
- `macro`
- `offsetof`
- `override`
- `priv`
- `proc`
- `pure`
- `sizeof`
- `typeof`
- `unsized`
- `virtual`
- `yield`

## Raw identifiers

Raw identifiers let you use keywords where they are not allowed by prefixing them with `r#`.

For example, `match` is a keyword. If you try to call

```
fn match(needle: &str, haystack: &str) -> bool {
    haystack.contains(needle)
}
```



You'll get this error:

```
error: expected identifier, found keyword
--> src/main.rs:4:4
   |
4 | fn match(needle: &str, haystack: &str) -
   |      ^^^^^ expected identifier, found ke
```

You can write this with a raw identifier:

```
fn r#match(needle: &str, haystack: &str) -
    haystack.contains(needle)
}

fn main() {
    assert!(r#match("foo", "foobar"));
}
```

Note the `r#` prefix on both the function name and

## Motivation

This feature is useful for a few reasons, but the most important are in the following situations. For example, `try` is not a keyword in Rust 2018. So if you have a library that is written in C++ and you want to call it in Rust 2018, you'll need to use the raw identifier syntax.

## Appendix B: Operators and Syntax

This appendix contains a glossary of Rust's syntax symbols that appear by themselves or in the context of macros, attributes, comments, tuples, and brackets.

### Operators

Table B-1 contains the operators in Rust, an example of how they appear in context, a short explanation, and whether an operator is overloadable, the relevant trait to implement, and whether it is listed.

Table B-1: Operators

Operator	Example	Expl
!	ident!(...), ident!{...}, ident![...]	Macro expansion
!	!expr	Bitwise or logical complement
!=	var != expr	Nonequality
%	expr % expr	Arithmetic remainder
%=	var %= expr	Arithmetic remainder assignment
&	&expr, &mut expr	Borrow
&	&type, &mut type, &'a type, &'a mut type	Borrowed pointer
&	expr & expr	Bitwise AND
&=	var &= expr	Bitwise AND assignment
&&	expr && expr	Logical AND
*	expr * expr	Arithmetic multiplication
*=	var *= expr	Arithmetic multiplication and assignment
*	*expr	Dereferencing
*	*const type, *mut type	Raw pointer
+	trait + trait, 'a + trait	Compound constraint
+	expr + expr	Arithmetic addition
+=	var += expr	Arithmetic addition assignment
,	expr, expr	Argument and separator
-	- expr	Arithmetic negation
-	expr - expr	Arithmetic subtraction

Operator	Example	Expl
<code>--</code>	<code>var -- expr</code>	Arithmetic and assignment
<code>-&gt;</code>	<code>fn(...) -&gt; type</code> <code>,  ...  -&gt; type</code>	Function arrow return type
<code>.</code>	<code>expr.ident</code>	Member access
<code>..</code>	<code>.., expr..,</code> <code>..expr,</code> <code>expr..expr</code>	Right-exclusive range literal
<code>..=</code>	<code>..=expr,</code> <code>expr..=expr</code>	Right-inclusive range literal
<code>..</code>	<code>..expr</code>	Struct literal syntax
<code>..</code>	<code>variant(x, ..),</code> <code>struct_type {</code> <code>  x, .. }</code>	"And the rebinding"
<code>...</code>	<code>expr...expr</code>	In a pattern range pattern
<code>/</code>	<code>expr / expr</code>	Arithmetic division
<code>/=</code>	<code>var /= expr</code>	Arithmetic division assignment
<code>:</code>	<code>pat: type,</code> <code>ident: type</code>	Constraints
<code>:</code>	<code>ident: expr</code>	Struct field
<code>:</code>	<code>'a: loop {...}</code>	Loop label
<code>;</code>	<code>expr;</code>	Statement terminator
<code>;</code>	<code>[...; len]</code>	Part of fixed-size array syntax
<code>&lt;&lt;</code>	<code>expr &lt;&lt; expr</code>	Left-shift
<code>&lt;&lt;=</code>	<code>var &lt;&lt;= expr</code>	Left-shift assignment
<code>&lt;</code>	<code>expr &lt; expr</code>	Less than comparison
<code>&lt;=</code>	<code>expr &lt;= expr</code>	Less than or equal comparison
<code>=</code>	<code>var = expr,</code> <code>ident = type</code>	Assignment

Operator	Example	Expl
<code>==</code>	<code>expr == expr</code>	Equality comparison
<code>=&gt;</code>	<code>pat =&gt; expr</code>	Part of match expression
<code>&gt;</code>	<code>expr &gt; expr</code>	Greater than comparison
<code>&gt;=</code>	<code>expr &gt;= expr</code>	Greater than or equal to comparison
<code>&gt;&gt;</code>	<code>expr &gt;&gt; expr</code>	Right-shift
<code>&gt;&gt;=</code>	<code>var &gt;&gt;= expr</code>	Right-shift and assignment
<code>@</code>	<code>ident @ pat</code>	Pattern binding
<code>^</code>	<code>expr ^ expr</code>	Bitwise XOR
<code>^=</code>	<code>var ^= expr</code>	Bitwise XOR assignment
<code> </code>	<code>pat   pat</code>	Pattern alternation
<code> </code>	<code>expr   expr</code>	Bitwise OR
<code> =</code>	<code>var  = expr</code>	Bitwise OR assignment
<code>  </code>	<code>expr    expr</code>	Logical OR
<code>?</code>	<code>expr?</code>	Error propagation

## Non-operator Symbols

The following list contains all non-letters that do not behave like a function or method call.

Table B-2 shows symbols that appear on their own in source code locations.

Table B-2: Stand-Alone Syntax

Symbol	
<code>'ident</code>	Named lifetime
<code>...u8</code> , <code>...i32</code> , <code>...f64</code> , <code>...usize</code> , etc.	Numeric literals
<code>"..."</code>	String literals

Symbol	
<code>r"..."</code> , <code>r#"..."#</code> , <code>r##"..."##</code> , etc.	Raw string processed
<code>b"..."</code>	Byte string of a string
<code>br"..."</code> , <code>br#"..."#</code> , <code>br##"..."##</code> , etc.	Raw byte s and byte s
<code>'...'</code>	Character
<code>b'...'</code>	ASCII byte
<code> ...  expr</code>	Closure
<code>!</code>	Always em functions
<code>-</code>	"Ignored"   integer lite

Table B-3 shows symbols that appear in the con hierarchy to an item.

Table B-3: Path-Related Syntax

Symbol	
<code>ident::ident</code>	Namespac
<code>::path</code>	Path relati explicitly a
<code>self::path</code>	Path relati explicitly r
<code>super::path</code>	Path relati module
<code>type::ident</code> , <code>&lt;type as trait&gt;::ident</code>	Associat
<code>&lt;type&gt;::...</code>	Associat directly na <code>&lt;[T]&gt;::...</code>
<code>trait::method(...)</code>	Disambigu trait that c
<code>type::method(...)</code>	Disambigu type for w

Symbol	
<code>&lt;type as trait&gt;::method(...)</code>	Disambiguates trait and type

Table B-4 shows symbols that appear in the context of generic functions.

Table B-4: Generics

Symbol	
<code>path&lt;...&gt;</code>	Specifies parameter type (e.g., <code>Vec&lt;u8&gt;</code> )
<code>path::&lt;...&gt;, method::&lt;...&gt;</code>	Specifies parameter type and method in an expression (e.g., <code>Vec::new()</code> )
<code>fn ident&lt;...&gt; ...</code>	Define generic function
<code>struct ident&lt;...&gt; ...</code>	Define generic struct
<code>enum ident&lt;...&gt; ...</code>	Define generic enum
<code>impl&lt;...&gt; ...</code>	Define generic implementation
<code>for&lt;...&gt; type</code>	Higher-ranked lifetime
<code>type&lt;ident=type&gt;</code>	A generic type with associated types

Table B-5 shows symbols that appear in the context of generic parameters with trait bounds.

Table B-5: Trait Bound Constraints

Symbol	
<code>T: U</code>	Generic parameter <code>T</code> must implement <code>U</code>
<code>T: 'a</code>	Generic type <code>T</code> must have a lifetime that cannot transitively outlive <code>'a</code>
<code>T: 'static</code>	Generic type <code>T</code> must have a lifetime that is at least as long as <code>'static</code>
<code>'b: 'a</code>	Generic lifetime <code>'b</code> must be at least as long as <code>'a</code>
<code>T: ?Sized</code>	Allow generic type parameter to be unsized

Symbol	Expr
'a + trait , trait + trait	Compound type constr

Table B-6 shows symbols that appear in the context of specifying attributes on an item.

Table B-6: Macros and Attributes

Symbol	Expr
#[meta]	Outer macro
#![meta]	Inner macro
\$ident	Macro identifier
\$ident:kind	Macro identifier and kind
\$(...)...	Macro invocation

Table B-7 shows symbols that create comments.

Table B-7: Comments

Symbol	Expr
//	Line comment
//!	Inner line comment
///	Outer line comment
/*...*/	Block comment
/*!...*/	Inner block comment
/**...*/	Outer block comment

Table B-8 shows symbols that appear in the context of specifying tuples.

Table B-8: Tuples

Symbol	Expr
()	Empty tuple
(expr)	Parenthesized expression
(expr,)	Single-element tuple
(type,)	Single-element tuple type
(expr, ...)	Tuple expression

Symbol	
<code>(type, ...)</code>	Tuple type
<code>expr(expr, ...)</code>	Function call tuple <code>struct</code>
<code>ident!(...),</code> <code>ident!{...},</code> <code>ident![...]</code>	Macro invoc
<code>expr.0, expr.1, etc.</code>	Tuple indexi

Table B-9 shows the contexts in which curly brack

Table B-9: Curly Brackets

Context	Ex
<code>{...}</code>	Block
Type <code>{...}</code>	struct

Table B-10 shows the contexts in which square b

Table B-10: Square Brackets

Context	
<code>[...]</code>	Array literal
<code>[expr; len]</code>	Array literal co
<code>[type; len]</code>	Array type con
<code>expr[expr]</code>	Collection inde <code>IndexMut</code> )
<code>expr[..], expr[a..],</code> <code>expr[..b], expr[a..b]</code>	Collection inde slicing, using <code>RangeFull</code> as

## Appendix C: Derivable Traits

In various places in the book, we’ve discussed th  
apply to a struct or enum definition. The `derive`  
implement a trait with its own default implemen  
with the `derive` syntax.



In this appendix, we provide a reference of all the traits you can use with `derive`. Each section covers:

- What operators and methods deriving this trait provide
- What the implementation of the trait signifies about the type
- What implementing the trait signifies about the type
- The conditions in which you're allowed or required to implement them
- Examples of operations that require the trait

If you want different behavior than that provided in the [standard library documentation](#) for each trait, you can implement them.

The rest of the traits defined in the standard library are `derive` types using `derive`. These traits don't have semantics, so you to implement them in the way that makes sense to accomplish.

An example of a trait that can't be derived is `Display`. It's for end users. You should always consider the approach from the end user. What parts of the type should an end user find relevant? What format of the data should they see? The Rust compiler doesn't have this insight, so it's up to you to define the behavior for you.

The list of derivable traits provided in this appendix can implement `derive` for their own traits, making `derive` with truly open-ended. Implementing `derive` macro, which is covered in Appendix D.

## `Debug` for Programmer Output

The `Debug` trait enables debug formatting in for adding `:?` within `{}` placeholders.

The `Debug` trait allows you to print instances of your type and other programmers using your type can point in a program's execution.

The `Debug` trait is required, for example, in use `println!` prints the values of instances given as argument so that programmers can see why the two instances we

## PartialEq and Eq for Equality Comparison

The `PartialEq` trait allows you to compare instances and enables use of the `==` and `!=` operators.

Deriving `PartialEq` implements the `eq` method. On structs, two instances are equal only if *all* fields are equal; on enums, two instances are equal only if any fields are not equal. When derived on a type, `PartialEq` compares a value to itself and not equal to the other variants.

The `PartialEq` trait is required, for example, with `HashMap`, which needs to be able to compare two instances of a key.

The `Eq` trait has no methods. Its purpose is to signal that a type is annotated type, the value is equal to itself. The `Eq` trait is required to implement `PartialEq`, although not all types can implement `Eq`. One example of this is floating point numbers: a `NaN` value is not equal to each other.

An example of when `Eq` is required is for keys in `HashMap<K, V>`. `HashMap` can tell whether two keys are the same.

## PartialOrd and Ord for Ordering Comparison

The `PartialOrd` trait allows you to compare instances. A type that implements `PartialOrd` can be used with the `partial_cmp` operators. You can only apply the `PartialOrd` trait to types that also implement `PartialEq`.

Deriving `PartialOrd` implements the `partial_cmp` method that will be `None` when the values are not comparable. An example of a value that doesn't provide an ordering is a floating point number. An example of a value that doesn't provide an ordering is a floating point number. Calling `partial_cmp` with any floating point number will return `None`.

When derived on structs, `PartialOrd` compares the value in each field in the order in which the fields are listed. When derived on enums, variants of the enum are considered less than the variants listed later.

The `PartialOrd` trait is required, for example, for

`rand` crate that generates a random value in the high value.

The `Ord` trait allows you to know that for any two valid ordering will exist. The `Ord` trait implementer `Ordering` rather than an `Option<Ordering>` be possible. You can only apply the `Ord` trait to type and `Eq` (and `Eq` requires `PartialEq`). When derived behaves the same way as the derived implementation `PartialOrd`.

An example of when `Ord` is required is when storing a structure that stores data based on the sort order.

## `Clone` and `Copy` for Duplicating Values

The `Clone` trait allows you to explicitly create a duplication process might involve running arbitrary code. See the “Ways Variables and Data Interact: Clone” section for information on `Clone`.

Deriving `Clone` implements the `clone` method. For a whole type, calls `clone` on each of the parts of the type. Values in the type must also implement `Clone`.

An example of when `Clone` is required is when storing a slice. The slice doesn't own the type instances it contains. `to_vec` will need to own its instances, so `to_vec` type stored in the slice must implement `Clone`.

The `Copy` trait allows you to duplicate a value by value; no arbitrary code is necessary. See the “Stack vs. Heap” Chapter 4 for more information on `Copy`.

The `Copy` trait doesn't define any methods to provide, so overloading those methods and violating the assumption of being run. That way, all programmers can assume it's fast.

You can derive `Copy` on any type whose parts all implement `Copy`. You can apply the `Copy` trait to types that also implement `Clone`. `Copy` has a trivial implementation compared to `Clone`.

The `Copy` trait is rarely required; types that implement it are available, meaning you don't have to call `clone`.

Everything possible with `Copy` you can also accomplish without it, but it might be slower or have to use `clone` in places.

## `Hash` for Mapping a Value to a Value of Fixed Size

The `Hash` trait allows you to take an instance of a type and map it to a value of fixed size using a hash function. The `hash` method. The derived implementation of `hash` calls `hash` on each of the parts of the type, so you must implement `Hash` to derive `Hash`.

An example of when `Hash` is required is in storing data efficiently.

## `Default` for Default Values

The `Default` trait allows you to create a default value for a type. It implements the `default` function. The derived implementation calls the `default` function on each part of the type. Values in the type must also implement `Default`.

The `Default::default` function is commonly used to create a default value. It is discussed in the "Creating Instance with Update Syntax" section in Chapter 5. You can create a default value and use a default value for the rest of the type with `..Default::default()`.

The `Default` trait is required when you use the `Option<T>` instances, for example. If the `Option::unwrap_or_default` will return the result of `Default::default()` in the `Option<T>`.

## Appendix D: Macros

We've used macros like `println!` throughout the book. This appendix explains what a macro is and how it works. This appendix

- What macros are and how they differ from
- How to define a declarative macro to do m
- How to define a procedural macro to creat

We're covering the details of macros in an appendix in Rust. Macros have changed and, in the near future, the rest of the language and standard library will likely to become out-of-date than the rest of the guarantees, the code shown here will continue to work. There may be additional capabilities or easier ways to do things at the time of this publication. Bear that in mind when you read from this appendix.

## The Difference Between Macros and Functions

Fundamentally, macros are a way of writing code that generates code known as *metaprogramming*. In Appendix C, we use `vec!` to generate an implementation of various traits for `Vec` and `vec!` macros throughout the book. All of the code generated is less than the code you've written manually.

Metaprogramming is useful for reducing the amount of code to maintain, which is also one of the roles of functions. Functions have additional powers that functions don't have.

A function signature must declare the number of arguments it has. Macros, on the other hand, can take a variable number of arguments. For example, you can call `println!("hello")` with one argument or two arguments. Also, macros are expanded before the meaning of the code, so a macro can, for example, generate code that a function can't, because it gets called at runtime instead of at compile time.

The downside to implementing a macro instead of a function is that they are more complex than function definitions because they write Rust code. Due to this indirection, macros are harder to read, understand, and maintain than functions.

Another difference between macros and functions is that macros are namespaced within modules like function definitions. When you have a name clash when using external crates, you have to use the scope of your project at the same time as you use the crate using the `#[macro_use]` annotation. The following

macros defined in the `serde` crate into the scope

```
#[macro_use]
extern crate serde;
```

If `extern crate` was able to bring macros into scope with the `#[macro_use]` annotation, you would be prevented from using macros with the same name. In practice, this could be a problem if you use more crates you use, the more likely it is.

There is one last important difference between `use` and `extern crate`. You can't define or bring macros into scope *before* you call `use`. You can define functions anywhere and call them anywhere.

## Declarative Macros with `macro_rules!` Metaprogramming

The most widely used form of macros in Rust are *declarative macros*, sometimes referred to as *macros by example*, `macro_rules!` macros. At their core, declarative macros allow you to define a macro that takes a Rust `match` expression. As discussed in Chapter 10, `match` expressions are structures that take an expression, compare the expression to a series of patterns, and then run the code associated with the pattern that matches. In a `match` expression, you compare a value to patterns that have code associated with them. The value being compared is the literal Rust source code passed to the `match` expression. The patterns are compared with the structure of that source code. The code associated with a pattern is the code that replaces the code passed to the `match` expression during compilation.

To define a macro, you use the `macro_rules!` macro. In the `vec!` macro, we can use the `vec!` macro to create a new `Vec`. For example, the following macro creates a new `Vec`:

```
let v: Vec<u32> = vec![1, 2, 3];
```

We could also use the `vec!` macro to make a `Vec` of string slices. We wouldn't be able to use a function to create a `Vec` because we wouldn't know the number or type of values up

Let's look at a slightly simplified definition of the

```

#[macro_export]
macro_rules! vec {
    ( $( $x:expr ),* ) => {
        {
            let mut temp_vec = Vec::new();
            $(
                temp_vec.push($x);
            )*
            temp_vec
        }
    };
}

```

Listing D-1: A simplified version of the `vec!` macro

---

Note: The actual definition of the `vec!` macro is more complex, including code to preallocate the correct amount of memory for the vector. We omit this optimization here to make the code more readable.

---

The `#[macro_export]` annotation indicates that the macro is exported from the crate in which we're defining it. This means that the macro is available to other crates that depend on this crate, even if someone depending on this crate doesn't use the macro. Without this annotation, the macro wouldn't be brought into scope by the `use` statement.

We then start the macro definition with `macro_rules!`, which is followed by the name of the macro we're defining *without* the exclamation mark. This is followed by curly brackets denoting the body of the macro definition.

The structure in the `vec!` body is similar to the structure of a `match` statement. Here we have one arm with the pattern `( $( $x:expr ),* )` followed by a block of code associated with this pattern. If the pattern matches, the block of code will be emitted. Given that this is the only valid way to match; any other will be an error. A macro can have more than one arm.

Valid pattern syntax in macro definitions is different from the pattern syntax in Chapter 18 because macro patterns are matched against tokens rather than values. Let's walk through what the pattern `( $( $x:expr ),* )` means; for the full macro pattern syntax, see [the Rust Reference](#).

First, a set of parentheses encompasses the whole pattern. This is followed by a set of parentheses, which captures the expression within the parentheses for use in the replacement code. The `$x` is a placeholder which matches any Rust expression and gives the

The comma following `$()` indicates that a literal `1` optionally appear after the code that matches the pattern `1`. Following the comma specifies that the pattern `1` precedes the `*`.

When we call this macro with `vec![1, 2, 3];`, the code generated that matches the three expressions `1`, `2`, and `3`.

Now let's look at the pattern in the body of the `vec!` macro. The `temp_vec.push()` code within the `$()*` part is `$()` in the pattern, zero or more times depending on the number of matches. The `$x` is replaced with each expression with `vec![1, 2, 3];`, the code generated that matches the following:

```
let mut temp_vec = Vec::new();
temp_vec.push(1);
temp_vec.push(2);
temp_vec.push(3);
temp_vec
```

We've defined a macro that can take any number of arguments and generate code to create a vector containing the arguments.

Given that most Rust programmers will use macros, we won't discuss `macro_rules!` any further. To learn more about macros, see the online documentation or other resources, such as ["Macros"](#).

## Procedural Macros for Custom `derive`

The second form of macros is called *procedural macros* (which are a type of procedure). Procedural macros take code as an input, operate on that code, and produce new code. Unlike `macro_rules!`, procedural macros do more than matching against patterns and replacing them. They can do anything that macros do. At the time of this writing, you can only use procedural macros to implement traits on a type by specifying a `derive` annotation.

We'll create a crate named `hello_macro` that defines one associated function named `hello_macro`. We'll implement the `HelloMacro` trait for each of the types we want to use the macro so users can annotate their type with `#[derive]`.



implementation of the `hello_macro` function. The `Hello, Macro! My name is TypeName!` where this trait has been defined. In other words, another programmer to write code like Listing D

Filename: src/main.rs

```
extern crate hello_macro;
#[macro_use]
extern crate hello_macro_derive;

use hello_macro::HelloMacro;

#[derive(HelloMacro)]
struct Pancakes;

fn main() {
    Pancakes::hello_macro();
}
```

Listing D-2: The code a user of our crate will be a procedural macro

This code will print `Hello, Macro! My name is` step is to make a new library crate, like this:

```
$ cargo new hello_macro --lib
```

Next, we'll define the `HelloMacro` trait and its a:

Filename: src/lib.rs

```
pub trait HelloMacro {
    fn hello_macro();
}
```

We have a trait and its function. At this point, our to achieve the desired functionality, like so:

```

extern crate hello_macro;

use hello_macro::HelloMacro;

struct Pancakes;

impl HelloMacro for Pancakes {
    fn hello_macro() {
        println!("Hello, Macro! My name is Pancakes!");
    }
}

fn main() {
    Pancakes::hello_macro();
}

```

However, they would need to write the implementation of the `hello_macro` function that we wanted to use with `hello_macro`; we want to split this into two crates that will work.

Additionally, we can't yet provide a default implementation of the `hello_macro` function that will print the name of the type that has the `HelloMacro` trait. Rust doesn't have reflection capabilities, so it can't look up the macro to generate code at compile time.

The next step is to define the procedural macro. Procedural macros need to be in their own crate. Eventually, there is a convention for structuring crates and macro crates. For example, if you have a custom derive procedural macro crate, it is usually named `foo_derive`. In our case, the crate is called `hello_macro_derive` inside our `hello_macro` crate.

```
$ cargo new hello_macro_derive --lib
```

Our two crates are tightly related, so we create a subdirectory of our `hello_macro` crate. If we change the name of the `hello_macro` crate, we'll have to change the implementation of the `hello_macro` function in `hello_macro_derive` as well. The two crates will be used by other crates, and programmers using these crates will need to bring them both into scope. We could instead have a single crate that contains both the `hello_macro` and `hello_macro_derive` as a dependency and reexport the `hello_macro` trait. The way we've structured the project makes it possible to use `hello_macro` even if they don't want the `hello_macro_derive` crate.

We need to declare the `hello_macro_derive` crate as a dependency in the `hello_macro` crate. We also need functionality from the `syn` and `quote` crates to generate code. We need to add them as dependencies. Add the

hello\_macro\_derive:

Filename: hello\_macro\_derive/Cargo.toml

```
[lib]
proc-macro = true

[dependencies]
syn = "0.14.4"
quote = "0.6.3"
```

To start defining the procedural macro, place the `src/lib.rs` file for the `hello_macro_derive` crate. until we add a definition for the `impl_hello_macro`

Filename: hello\_macro\_derive/src/lib.rs

```
extern crate proc_macro;
extern crate syn;
#[macro_use]
extern crate quote;

use proc_macro::TokenStream;

#[proc_macro_derive(HelloMacro)]
pub fn hello_macro_derive(input: TokenStream)
    // Construct a representation of Rust code
    // that we can manipulate
    let ast = syn::parse(input).unwrap();

    // Build the trait implementation
    impl_hello_macro(&ast)
}
```

Listing D-3: Code that most procedural macro crates use in Rust code

Notice the way we've split the functions in D-3; the `hello_macro_derive` is the procedural macro crate you see or create, because it's the one that defines the `HelloMacro` trait. The `impl_hello_macro` function is called will be different depending on the macro's purpose.

We've introduced three new crates: `proc_macro` crate comes with Rust, so we didn't need to add it to `Cargo.toml`. The `proc_macro` crate allows us to create a crate containing that Rust code. The `syn` crate parses Rust code into a structure that we can perform operations on. The

structures and turns them back into Rust code. To parse any sort of Rust code we might want to have, this task is no simple task.

The `hello_macro_derive` function will get called with `#[derive(HelloMacro)]` on a type. The reason it's called `hello_macro_derive` function here with `proc_macro_derive(HelloMacro)`, which matches our trait name; that's how macros follow.

This function first converts the `input` from a `TokenStream` to a `Parser`. We can then interpret and perform operations on the `Parser`. The `parse` function in `syn` takes a `TokenStream` representing the parsed Rust code. The following is the `DeriveInput` struct we get from parsing the

```
DeriveInput {  
    // --snip--  
  
    ident: Ident(  
        "Pancakes"  
    ),  
    body: Struct(  
        Unit  
    )  
}
```

The fields of this struct show that the Rust code `ident` (identifier, meaning the name) of `Pancakes` struct for describing all sorts of Rust code; check `DeriveInput` for more information.

At this point, we haven't defined the `impl_hello_macro` to build the new Rust code we want to include. But we also have a `TokenStream` which is added to the code. When they compile their crate, they'll get extra functions.

You might have noticed that we're calling `unwrap` on `syn::parse` function fails here. Panicking on error is not good because `proc_macro_derive` functions must return `Result` to conform to the procedural macro API. In our example by using `unwrap`; in production code, you should return error messages about what went wrong by using `panic!`.

Now that we have the code to turn the annotated

a `DeriveInput` instance, let's generate the code trait on the annotated type:

Filename: `hello_macro_derive/src/lib.rs`

```
fn impl_hello_macro(ast: &syn::DeriveInput) {
    let name = &ast.ident;
    let gen = quote! {
        impl HelloMacro for #name {
            fn hello_macro() {
                println!("Hello, Macro! My name is {}",
                    (#name));
            }
        }
    };
    gen.into()
}
```

We get an `Ident` struct instance containing the type using `ast.ident`. The code in Listing D-2 shows `Ident("Pancakes")`.

The `quote!` macro lets us write the Rust code that the result of its execution is not what is expected by the compiler, but is converted to a `TokenStream` by calling the `into` method, which returns an intermediate representation and returns a value of type `TokenStream`.

This macro also provides some very cool templates. The `quote!` will replace it with the value in the variable. We can do some repetition similar to the way regular macros work. See [crate's docs](#) for a thorough introduction.

We want our procedural macro to generate an implementation of the `HelloMacro` trait for the type the user annotated, which we call `hello_macro`. The implementation has one function, `hello_macro`, which provides the functionality we want to provide: printing `Hello` and the name of the annotated type.

The `stringify!` macro used here is built into Rust. It takes an expression like `1 + 2`, and at compile time turns the expression into a string. This is different than `format!` or `println!`, which turn the result into a `String`. There is a possibility of a macro expression to print literally, so we use `stringify!` to avoid string allocation by converting `#name` to a string literal.

At this point, `cargo build` should complete successfully.

`hello_macro_derive` . Let's hook up these crates to see the procedural macro in action! Create a new binary using `cargo new pancakes` . We need to add `hello_macro` as dependencies in the `pancakes` crate's `Cargo.toml` of `hello_macro` and `hello_macro_derive` to `hello_macro` dependencies; if not, you can specify them as `path =`

```
[dependencies]
hello_macro = { path = "../hello_macro" }
hello_macro_derive = { path = "../hello_macro_derive"
```

Put the code from Listing D-2 into `src/main.rs`, and run `cargo run`. The output is `Hello, Macro! My name is Pancakes!`. The implementation from the procedural macro was included without having to implement it; the `#[derive(HelloMacro)]` attribute does the work.

## The Future of Macros

In the future, Rust will expand its declarative and procedural macro system with the `macro_rules!` macro system for more powerful tasks than `macro_rules!` is currently capable of. This section is under development at the time of this publication. See the [macro\\_rules! documentation](#) for the latest information.

## Appendix E: Translations of the Rust Book

For resources in languages other than English. See the [Translations label](#) to help or let us know about a translation.

- [Português \(BR\)](#)
- [Português \(PT\)](#)
- [Tiếng việt](#)
- [简体中文, alternate](#)
- [Українська](#)
- [Español](#)
- [Italiano](#)
- [Русский](#)
- [한국어](#)
- [日本語](#)
- [Français](#)

- [Polski](#)
- [עברית](#)
- [Cebuano](#)
- [Tagalog](#)

## Appendix F - How Rust “Nightly Rust”

This appendix is about how Rust is made and how the developer.

### Stability Without Stagnation

As a language, Rust cares a *lot* about the stability of the rock-solid foundation you can build on, and if that foundation would be impossible. At the same time, if we care about stability, we may not find out important flaws until after they have changed things.

Our solution to this problem is what we call “stable Rust”. The guiding principle is this: you should never have to upgrade Rust. Each upgrade should be painless, with fewer bugs, and faster compile times.

### Choo, Choo! Release Channels and Ridi

Rust development operates on a *train schedule*. The `master` branch of the Rust repository. Release channels, which has been used by Cisco IOS and other operating systems, three *release channels* for Rust:

- Nightly
- Beta
- Stable

Most Rust developers primarily use the stable channel. If you want to experiment with new features, you may use nightly or beta.

Here’s an example of how the development and release cycle of Rust works. The Rust team is working on the release of F

December of 2015, but it will provide us with releases. A new commit is added to Rust: a new commit lands on the `master` branch. A new nightly version of Rust is produced. Every day a new commit is created by our release infrastructure automatically. A new nightly release looks like this, once a night:

```
nightly: * - - * - - *
```

Every six weeks, it's time to prepare a new release. We branch repository branches off from the `master` branch to create new releases:

```
nightly: * - - * - - *
          |
beta:    *
```

Most Rust users do not use beta releases actively. The release system to help Rust discover possible regressions or bugs in the nightly release every night:

```
nightly: * - - * - - * - - * - - *
          |
beta:    *
```

Let's say a regression is found. Good thing we have a beta branch before the regression snuck into a stable release. The nightly is fixed, and then the fix is backported to the beta branch. A new beta release is produced:

```
nightly: * - - * - - * - - * - - * - - *
          |
beta:    * - - - - - - - - *
```

Six weeks after the first beta was created, it's time to create a new stable branch. A new stable branch is produced from the `beta` branch:

```
nightly: * - - * - - * - - * - - * - - * -
          |
beta:    * - - - - - - - - *
          |
stable:  *
```

Hooray! Rust 1.5 is done! However, we've forgotten to update the version number. We have gone by, we also need a new beta of the next version. We branch `stable` branches off of `beta`, the next version is



again:

```
nightly: * - - * - - * - - * - - * - - * -  
          |  
beta:    * - - - - - - - - *  
          |  
stable:  *  
          |
```

This is called the “train model” because every six weeks a new version of Rust is released, but still has to take a journey through the beta channel before reaching stable.

Rust releases every six weeks, like clockwork. If you know the date of the next release, you can know the date of the next one: having releases scheduled every six weeks is the feature happens to miss a particular release, the happening in a short time! This helps reduce pre-release features in close to the release deadline.

Thanks to this process, you can always check out the latest version of Rust and know that it’s easy to upgrade to: if a beta release has a bug, you can report it to the team and get it fixed before the next release. Breakage in a beta release is relatively rare, but bugs do exist.

## Unstable Features

There’s one more catch with this release model: the technique called “feature flags” to determine when a feature is ready for release. If a new feature is under active development, it’s available in nightly, but behind a *feature flag*. If you want to use a work-in-progress feature, you can, but you must annotate your source code with the appropriate flags.

If you’re using a beta or stable release of Rust, you can’t use the key that allows us to get practical use with nightly. Those who wish to opt into the beta channel can do so, but those who want a rock-solid experience can stick with stable. Stability without stagnation.

This book only contains information about stable features, but they’re still changing, and surely they’ll be different between versions when they get enabled in stable builds. You can find the latest features online.

## Rustup and the Role of Rust Nightly

Rustup makes it easy to change between different toolchains on a global or per-project basis. By default, you'll have the stable toolchain installed, but you can also install the nightly toolchain, for example:

```
$ rustup install nightly
```

You can see all of the *toolchains* (releases of Rust) that have been installed with `rustup` as well. Here's an example on a Windows computer:

```
> rustup toolchain list
stable-x86_64-pc-windows-msvc (default)
beta-x86_64-pc-windows-msvc
nightly-x86_64-pc-windows-msvc
```

As you can see, the stable toolchain is the default. You might want to use the stable toolchain most of the time. You might want to use the nightly toolchain for a project, because you care about a cutting-edge feature. You can use the `rustup override` command in that project's directory to set the toolchain that `rustup` should use when you're in that directory.

```
$ cd ~/projects/needs-nightly
$ rustup override set nightly
```

Now, every time you call `rustc` or `cargo` inside that directory, it will make sure that you are using the nightly Rust compiler. This comes in handy when you have a lot of Rust code that needs the latest features.

## The RFC Process and Teams

So how do you learn about these new features? The Rust team uses the *Request For Comments (RFC)* process. If you'd like to propose a new feature, you write up a proposal, called an RFC.

Anyone can write RFCs to improve Rust, and they are discussed by the Rust team, which is comprised of several teams listed on [Rust's website](#), which include language design, compiler implementation, infrastructure, and more. The appropriate team reads the proposal and discusses it. If they agree, they work on it of their own, and eventually, there's consensus to accept the feature.

If the feature is accepted, an issue is opened on the Rust issue tracker.

can implement it. The person who implements it is the person who proposed the feature in the first place! When the feature is ready, the `master` branch is moved behind a feature gate, as we saw in the previous section.

After some time, once Rust developers who use the new feature, team members will discuss the feature nightly, and decide if it should make it into stable. If it does, the feature gate is removed, and the feature moves forward, the feature gate is removed, and the feature is included in the next stable release of Rust. It rides the trains into a new stable release of Rust.

## G - Other useful tools

In this appendix, we'll talk about some additional tools that are useful when developing Rust code.

### Automatic formatting with `rustfmt`

`rustfmt` is a tool that can re-format your code automatically. All Rust projects use `rustfmt` to prevent arguments about formatting. Just do what the tool does!

`rustfmt` is not at 1.0 yet, but a preview is available. Please give it a try and let us know how it goes!

To install `rustfmt`:

```
$ rustup component add rustfmt-preview
```

This will give you both `rustfmt` and `cargo-fmt`. To take any Cargo project and format it, use `rustc` and `cargo`. To take any Cargo project and format it, use `rustc` and `cargo`.

```
$ cargo fmt
```

### IDE integration with the Rust project

To help IDE integration, the Rust project distributes the `rust-analyzer` as in <http://langserver.org/>. This can be used by the [rust-analyzer plugin for Visual Studio: Code](#).

The `rls` is not at 1.0 yet, but a preview is available! Please give it a try and let us know how it goes!

To install the `rls`:

```
$ rustup component add rls-preview
```

Then, install the language server support in your