

Introduction to Concurrent Programming in C

Filip Strömbäck
Linköping University

2026-03-19
(commit: 3936b37)

Preface

The goal of this book is to provide a detailed introduction to concurrent programming. As such, the goal is to cover the fundamentals required to write and maintain programs that are correct according to the language specification. That is, the main focus is on *correctness* rather than on *maximizing performance*. That being said, the book will of course cover obvious performance issues with certain approaches.¹ However, the book will *not* cover advanced techniques for maximizing performance by, for example, using lock-free data structures. Rather, the book focuses on the basics required to understand the problems that concurrent programming introduces, and the tools required to address those issues. This is sufficient for at least 90% of all cases, but it also provides a solid starting point for studying the advanced techniques.

This book will try to take a hands-on approach to a subject that is otherwise quite theoretical in nature. It does so by illustrating the ideas with examples and exercises. In particular, the book contains many boxes labeled *Practice* that contains small exercises related to the text so that you can experiment with the material while you read the chapter.

The book uses the C programming language, but the concepts covered are general enough to apply to most imperative and object-oriented languages that rely on the *shared-memory model* of concurrency, such as C++, Java, Rust, Python, etc. Compared to C, these languages provide more powerful facilities for building abstractions. This means that they make it possible to use the concepts from this book in a more ergonomic way. However, these abstractions often hide details that are important to consider when working with concurrent programs. As such, the simpler abstraction facilities available in C is actually beneficial for novices since these important details become more explicit.

The book assumes familiarity with programming in some C-like language. It is, however, not necessary to be deeply familiar with C in order to be able to follow along.

¹After all, a typical goal with concurrent programming is to improve performance.

Acknowledgements

I would like to thank the following people for their contributions (both direct and indirect) to the contents of this book.

Klas Arvidsson for developing the material for the course that I later started teaching. This material is still an important part of the course to this day, and as such it has been a source of inspiration for many of the examples in this book.

Gunnar Wolf for his interest in Progvis, his help in getting Progvis packaged for Debian, and contributions to address oversights in the book.

Contents

Preface	i
Acknowledgements	ii
Contents	iii
1 Tools and Software	1
1.1 Examples and Exercises from the Book	1
1.2 C Compiler	1
1.3 Progviz	4
2 Programming in C	7
2.1 Functions	7
2.2 Variables	9
2.3 Output	9
2.4 Types	11
2.5 Pointers	12
2.6 Arrays	15
2.7 Memory Management	19
2.8 Abstract Data Types	22
3 Concurrent Programming	26
3.1 Execution Model	26
3.2 Threads, Processes and Programs	30
3.3 Starting Threads	33
3.4 Scheduling	37
3.5 Problems in Concurrent Programs	39
3.6 The Concurrent Execution Model	48
3.7 Exercises	51
4 Semaphores	52
4.1 Semantics	52
4.2 Semantics in Sequential Programs	54
4.3 Semantics in Concurrent Programs	56
4.4 Waiting for Completion	59
4.5 Multiple Semaphores	65

4.6	Mutual Exclusion	71
4.7	Exercises	76
5	Locks	77
5.1	Semantics	77
5.2	Relation to Semaphores	79
5.3	Using Locks for Mutual Exclusion	81
5.4	Benefits of Error Checking	83
5.5	Exercises	87
6	Critical Sections	88
6.1	Why Critical Sections?	88
6.2	Critical Sections and Conditionals	91
6.3	Shared Data in Multiple Functions	94
6.4	Synchronization Granularity	96
6.5	Working with Critical Sections	113
6.6	Exercises	115
7	Deadlocks	116
7.1	What is a Deadlock?	118
7.2	Preventing Deadlocks	123
7.3	Deadlock Avoidance in Programs	126
7.4	Exercises	128
8	Condition Variables	130
8.1	Semantics	130
8.2	Condition Variables and Semaphores	140
8.3	Waiting Until a Condition is True	142
8.4	A More Complex Condition	147
8.5	Exercises	155
9	Abstractions and Concurrency	156
9.1	Functions	157
9.2	Abstract Data Types	167
9.3	Summary: Sequential Abstractions in Concurrent Systems	182
9.4	Concurrent Abstract Data Types	184
9.5	Summary: Thread Safety	202
9.6	Exercises	204
10	Atomics	205
10.1	Implementing a Lock	205
10.2	Atomic Operations	208
10.3	Lock-Free Access to Shared Data	220
10.4	Performance	226
10.5	Advanced Usage	231
10.6	Exercises	235
A	Threading Library Reference	236
A.1	Thread Management	236
A.2	Synchronization Primitives	239
A.3	Atomic Operations	242

CONTENTS

v

B Limitations in Progis

244

Tools and Software

This book uses two main tools for examples and exercises: a C compiler and the visualization tool Progvis. The C compiler along with the threading library will be used to compile and run small programs to observe their behavior in a “real” setting. Progvis is a complement to the C compiler. It is also able to compile and run simple C programs. The difference is that Progvis visualizes how the program is being executed, and even lets you influence the execution to find concurrency issues.

1.1 Examples and Exercises from the Book

All the examples in the book are also available as source code in an archive available at <https://storm-lang.org/progvis-book/>. You simply need to download and extract the contents of the archive somewhere on your system to be able to follow along with the examples in the book.

The archive contains two directories for each chapter in the book. The directories are named `xx-practice` and `xx-exercises`, where `xx` is the chapter number. The first one contains the code discussed in the chapter. As such, it is where you will find the source code for the problems in the *Practice* boxes in the text. The second one (`xx-exercises`) contains the code for the exercises at the end of the chapter. There are also example solutions for some of the problems. They are located separately in the directory `solutions/xx`.

1.2 C Compiler

The code in this book has been developed on a GNU/Linux system using GCC. While the code should work on other systems as well, you might encounter cases where “incorrect” examples behave differently on your system.¹ While this usually only serves to illustrate the point of the book better, it might lead to unnecessary confusion. As such, we advise you to use a GNU/Linux system if possible. If you are running Windows, you can use Windows Subsystem for Linux (WSL) to conveniently run Linux tooling from your Windows installation. On MacOS, the compiler from XCode should be similar enough, but a

¹This is since they contain *undefined behavior*, so anything may happen.

Linux Virtual Machine is required to run Progvis anyway, so you might as well use that for the C compiler.

1.2.1 Installation

The code in this book only depends on the system's C libraries,² and Make. In many cases they are installed by default. If not, you can install them using commands similar to the ones below:

```
Debian, Ubuntu:  sudo apt install gcc make
Arch Linux:      sudo pacman -S gcc make
```

Apart from a C compiler, you also need the threading library provided with this book. This library does not need to be installed since it is small enough to be compiled alongside each example. This is done automatically as long as you are working with any of the examples provided with the book. See Chapter A for documentation for the library.

If you wish to use the library for other programs, you need to create a make-file that includes the library. How this is done is shown below.

1.2.2 Usage

The tooling provided with this book is designed for small programs that consist of a single source file. To compile a program, open a terminal and change to the directory where the file you wish to compile is located. Then simply type:

```
make <name>
```

Where <name> is the name of the file you wish to compile, without the .c extension. So for example, to compile the file `test.c`, you type `make test` and hit enter. This process creates the file <name> (e.g., `make test` produces the file `test`). You can then run the program by typing `./<name>` (e.g., `./test`).

Often, you end up wanting to compile a program and then directly execute it. To make this more convenient, you can chain the two commands with the `&&` operator, like:

```
make <name> && ./<name>
```

The makefile also provides the option to remove all compiled files by running the command:

```
make clean
```

1.2.3 Optimizations

By default, the makefile does not enable compiler optimizations. This is often exactly what we want when working with the examples in this book, since it turns out that the optimizations in the compiler often breaks incorrectly

²More specifically, pthreads for Linux

synchronized concurrent programs in interesting ways. It is possible to enable optimizations by adding `OPT=1` to the `make` command:

```
make OPT=1 <name>
```

This also forces the program to be re-compiled, so it is not necessary to run `make clean` before enabling optimizations. This is not true in the other direction, so if you have compiled a file with optimizations (e.g., `make OPT=1 test`) and then wish to compile it without optimizations (e.g., `make test`), `make` will simply say “Nothing to be done...” and not re-compile the program. In this case, you need to either delete the compiled file first (`rm test`) or run `make clean` to remove all compiled programs first.

The text in the book will explicitly mention when it is useful to enable optimizations.

1.2.4 Creating Custom Examples

If you wish to create your own examples to experiment with the concepts in this book, you have two options:

- You can add your own `.c` file to any of the directories that already contain examples from the book. This makes it possible to compile your code the same way as you compile and run the examples.
- If you want a separate directory for your own experiments, you also need to create a file named `Makefile` in that directory (note the capital M) with the following contents:

```
BOOK_LIB_PATH := <path_to_lib>  
include $(BOOK_LIB_PATH)/Make.template
```

Replace the text `<path_to_lib>` with the path to the `lib` directory you extracted from the archive that contained the examples to this book. After this, you should be able to compile and run your code as outlined above.

1.3 Progvis

It is also possible to run the examples in this book in Progvis. This lets you see their behavior in more detail, and even affect the execution of the program. The programs in this book require Progvis version 0.6.9 or later.

1.3.1 Installation

How to install and run Progvis depends on your operating system as described below:

Debian, Ubuntu Progvis is available as a package through the system's package manager. You can install it by running `sudo apt install progvis` in a terminal. You can then start Progvis from the system menu (look under the *Education* category), or by typing `progvis` in a terminal.

Make sure that the version of Progvis installed is 0.6.9 or later (the release cycles of Debian and Ubuntu are quite long). You can check the version by opening Progvis and selecting *Help* \Rightarrow *About...* in the menu. If Progvis is too old, you can follow the instructions for Linux below instead.³

Linux For other distributions of Linux, first try to use the Storm files available on <https://storm-lang.org/>. They are known to work on Debian, Ubuntu, and Arch Linux at least. Download the archive for your system from the *Downloads* page, extract the files to some directory using `tar xf storm*.tar.gz` or any graphical tool available in your system.⁴ The archive contains a file named `progvis.sh`. Run that file in a terminal or by clicking it in your graphical environment.

If Progvis does not start, you likely need to compile it from source. This is described on <https://storm-lang.org/> under *Getting Started* \Rightarrow *Developing in Storm* \Rightarrow *Compile from Source*. When you are done, you can launch Progvis using the command `mymake Main -- -f progvis.main` in the directory where you compiled Storm.

Windows Even if you are using WSL to compile C code, it is recommended to use the native Windows version of Progvis. Progvis includes its own C-compiler, so you don't need to install any other C compilers to use it.

Download Storm from the Downloads page of <https://storm-lang.org/>. Extract the zip-file to some location. This can be done by opening the zip file in Explorer and dragging the files to some other directory. After this is done, you can launch Progvis by clicking the file `Progvis.bat`.

Depending on how you downloaded and extracted the files, you may see a message with a title like *Windows has protected your computer* the first time you launch Progvis. If this happens, click the *More information* link. This shows a button labeled *Run anyway*. Click that button to launch Progvis. This only happens the first time you launch Progvis.

³Note: version 0.6.3 is enough for almost all examples. You can update when you encounter compilation errors.

⁴Chrome sometimes uncompresses the downloaded file when it is downloaded. If `tar` complains, try to rename the file from `.tar.gz` to just `.tar`.

1.3.2 Usage

Once you have installed and started Progvis as described above, you will see the main window of Progvis. The first step to use Progvis is to open a program. To do this, select *File* \Rightarrow *Open file...*, and select a program. Note that Progvis wants you to select the source code of the program you wish to run. You *don't* need to compile the program beforehand. Progvis takes care of the entire compilation process.

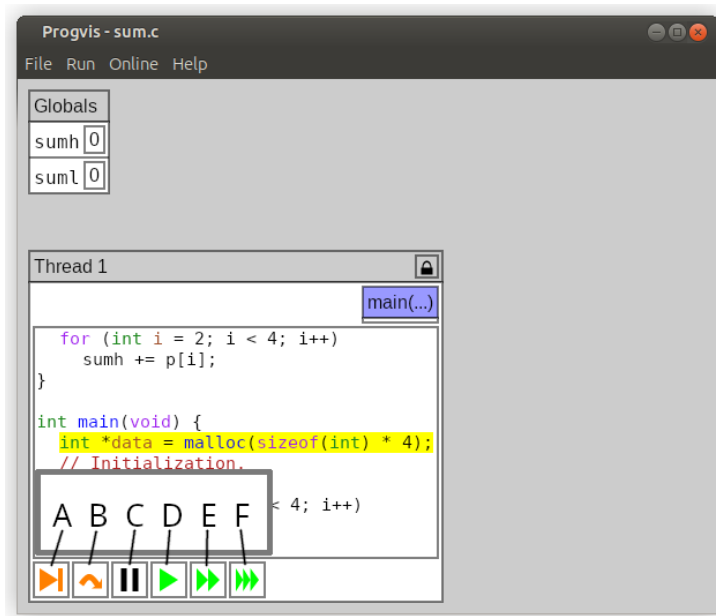


Figure 1.1: Progvis displaying the starting point of a program. In this case, the program `sum.c` from the built-in examples.

When you have selected a program, Progvis will automatically compile the program for you. If the compilation succeeded, Progvis will display the initial state of the program in its main window, similarly to what is shown in Fig. 1.1. The box labeled *Thread 1* shows the next line that the main thread of the program will execute next highlighted in yellow. Any local variables will also be shown in the blue box labeled `main(...)`. Global variables are shown in the box labeled *Globals*. The remainder of the gray area will be used for heap-allocated data as the program runs. Program execution can be controlled by pressing the buttons labeled A–F in *Thread 1*. They do the following:

- A: Let the current thread run to the next statement.
- B: Let the current thread run to the next *barrier*. These steps are typically much longer than stepping a single statement. The exact meaning will be covered later in the book.
- C: Pause automatic stepping.
- D: Start stepping automatically. Equivalent to clicking button A once every second.

- E: Start stepping automatically. Equivalent to clicking button A twice every second.
- F: Start stepping automatically. Equivalent to clicking button A four times every second.

All boxes in the main window of Progvis can be moved by dragging the title of the box. It is also possible to move everything by dragging the background. A few boxes have a small padlock in the upper-right corner. They are locked to the edges of the main window by default. If you drag them around, they will become detached and behave like the other boxes. You can lock them to the edge again by clicking the padlock.

If any errors were encountered during compilation, Progvis will display an dialog that highlights the error. Progvis still remembers which file you were attempting to open, so after you have fixed the problem (using your preferred text editor) you can simply select *File* \Rightarrow *Reload program* to open the same program again.

This ability is also useful when experimenting. If you made changes to your program and wish to see the new behavior of the program, you can simply select *File* \Rightarrow *Reload program* instead of locating the same file again.

As an additional convenience, Progvis is also able to open the current program in a text editor using *File* \Rightarrow *Open in editor*. This command attempts to open the default editor in your system. If this is not the editor you prefer, you can override the default in *File* \Rightarrow *Settings*. Note that on Linux, you can simply type the name of a command in the box *Editor to use to open code for Progvis*.

Finally, it is worth noting that Progvis supports a subset of the full C language. This is mainly to make C a bit more type-safe so that Progvis is able to visualize and analyze the program properly. The limitations are described in further detail in Chapter B.

Programming in C

This chapter provides a brief introduction to the C language as it is used in the remainder of this book. The reader is assumed to already be familiar with some other imperative programming language (e.g., C++). Readers who are already familiar with C can largely skip this chapter and use as a reference while reading the remainder of the book. Section 2.8 might, however, still be useful to read as it outlines the conventions used to create abstract data types in the book.

The examples covered in this chapter can be visualized in Progvis to illustrate their behavior in more detail if desired. To do this, simply select *File* ⇒ *Open...* in Progvis and select the example in the `02-c-programming` directory. Then press the leftmost button in the box labeled *Thread 1* to single-step the program and see the behavior of the constructs.

2.1 Functions

C requires that all program code appears inside a function. Functions are declared as below:

```
<result> <name>(<parameters>) {  
    <code>  
}
```

Where `<name>` is the name of the function, `<result>` is the type of the value returned by the function, and `<parameters>` is a comma-separated list of parameters to the function. For functions that do not return anything, `<result>` should be `void`. Similarly, but perhaps surprisingly, `<parameters>` should be `void` if the function does not accept any parameters.¹ As such, a function `f` that accepts zero parameters and does not return any result looks like below:

```
void f(void) {  
    /* ... */  
}
```

¹Empty parentheses (i.e., `f()`) actually means that the parameter list is *unknown* and will be specified later. This will change in an upcoming C standard, however.

The parameter list (`<parameters>`) is a comma-separated list of parameters. Each element has the form `<type> <name>` and defines a single parameter. For example, a function `f` that returns an integer and accepts two integer parameters looks as follows:

```
int f(int a, int b) {
    /* ... */
}
```

The function `main` is special. It dictates the start of program execution and will be automatically called by the system when the program is started. The `main` function is expected to return an integer, which will be returned to the system so that whoever started your program can assess whether your program succeeded in doing its job or not. The `main` function may also accept command-line parameters from the system.

In the context of this book (for compatibility with Progvis), the `main` function will always look like below:

```
int main(void) {
    return 0; // For success.
}
```

Note that the `return` statement is required by Progvis, even though it is optional according to the C standard.

Parameters to functions are always passed *by value*. This means that the values passed to a function through parameters are always *copied* into the function. Any changes the called function (the *callee*) makes to its parameters are thereby not visible in the caller. This is illustrated by the program `by-value.c`:

```
int f(int x) {
    x = 20;
    return 12;
}

int main(void) {
    int a = 10;
    int b = f(a);
    printf("a=%d\n", a);
    printf("b=%d\n", b);
    return 0;
}
```

The program will print `a=10` and `b=12`. Even though `a` is passed as the first parameter to `f`, and `f` sets `x` to 20, the modification is not reflected in `a` since `f` creates its own copy of the variable. What happens can be illustrated by the program below, where `f` has been inserted inside `main`, and the copy made explicit (also available in `by-value-expanded.c`):

```
int main(void) {
    int a = 10;
    int b;
    {
        int x = a;
        x = 20;
        b = 12;
    }
    printf("a=%d\n", a);
    printf("b=%d\n", b);
    return 0;
}
```

Note that C does not support references (i.e., `&` in C++). It is, however, possible to use pointers to explicitly let a called function modify variables in the calling function.

2.2 Variables

As we saw above, variables are declared as `<type> <name>`, just like parameters. For example, an integer variable named `x` is declared as follows:

```
int x;
```

Declarations like the above can appear either on their own, at global scope, inside functions, or inside data structures. A variable that at global scope defines a global variable. A variable inside a function defines a variable that is local to each function call. As we saw above, it is possible to add additional blocks inside functions to limit the scope of local variables as desired. Variables defined inside a struct become members of the struct.

By default, global variables are initialized to 0, while the contents of local variables are undefined.² To avoid difficulties in debugging programs, it is therefore a good habit to *always initialize variables explicitly*. In C, initialization looks like an assignment to the newly declared variable:

```
int x = 0;
```

Of course we can initialize local variables using an arbitrary expression. We are, however, more limited when we try to initialize global variables.

Note that C allows multiple variables to be declared at the same time (e.g., `int x, y;`). This notation is not supported by Progviz, and will not be used in this book.

2.3 Output

Formatted output in C is most commonly done through the `printf` function (from the header `<stdio.h>`). The first parameter to `printf` is a *format string* that describes the text that should be printed. We can use `printf` to print a simple message as follows:

²Progviz shows them as 0 in an effort to be deterministic.

```
printf("message\n");
```

Note that the string ends with a newline character (i.e., `\n`). The `printf` function does not automatically add a newline character. As such, if we omit it from the format string, any subsequent calls to `printf` would appear on the same line. Additionally, the C standard library typically buffers entire lines of output. This means that without the explicit newline character, the message may not be visible until much later, when another call to `printf` outputs a newline character.

As the name *format string* implies, we can not only use it to output plain text. We can also include instructions on how to format the contents of variables. For `printf` this is done by adding a *format specifier* to the format string. A format specifier consists of a percent sign (`%`) followed by a description of what to output. This causes `printf` to replace the format specifier with a formatted version of the next parameter to `printf`.

In this book, we will use the following format specifiers:

- `%s` outputs a string (i.e., type `const char *`).
- `%c` outputs a character (i.e., type `char`).
- `%d` outputs an integer (i.e., type `int`).
- `%%` outputs a single percent sign.

Strings and integers can optionally be padded to a specified width by adding a number between the `%` and the `s` or `d`. For example, `%10d` will output an integer, but padded with spaces until the resulting number is at least 10 characters wide. If the number is positive (e.g., `%10d`) the spaces are added to the left, and if the number is negative (e.g., `%-10d`) they are added to the right. Integers can optionally be padded with zeros if the number starts with a zero (e.g., `%010d`).

Usage of `printf` is illustrated by the program `printf.c`, which is also depicted below:

```
int main(void) {
    int a = 1;
    int b = 20;
    const char *s = "test";

    printf("%s\n", s);           // Outputs: "test\n"
    printf("%s%d\n", s, a);     // Outputs: "test1\n"
    printf("%s %d\n", s, b);    // Outputs: "test 20\n"
    printf("%s: %3d\n", s, a);  // Outputs: "test:  1\n"
    printf("%s: %3d\n", s, b);  // Outputs: "test: 20\n"
    return 0;
}
```

Note that format specifiers do not have to be delimited by spaces. If there are spaces or other characters between format specifiers they will appear in the output, just as when passing normal strings to `printf`. Also note that it is possible for a single format string to contain multiple format specifiers. In this case, the first format specifier uses the first parameter after the format string, the second uses the second, and so on.

2.4 Types

C provides a number of built-in integer types. This book uses the following types:

- `int`: a signed integer. In general it is only safe to assume that an integer is at least 16 bits wide, but on most modern desktop systems it is safe to assume that an `int` is at least 32 bits wide. This is the case in Progviz, and what is assumed in this book.
- `char`: a single character. It is usually safe to assume that a `char` is 8 bits wide, since the POSIX standard (which e.g., Linux follows) requires 8 bits in a `char`.
- `bool`: a boolean truth value (i.e., `true` or `false`). Note that `bool` is not a built-in type in C. It is implemented in the header `stdbool.h`, so it is necessary to `#include <stdbool.h>` to use booleans. The headers in the threading library for this book does this for you, so you usually do not have to worry about including it separately.

All types can be prefixed with `const` to refer to a constant value of the type. Most notably is the absence of a dedicated string type. We will see later in this chapter how strings are represented in C.

It is also possible to define custom data structures using the `struct` keyword. We will therefore refer to them as *structs*. A struct is a collection of variables that are stored one after another in memory. The variables in a struct are furthermore handled as a unit, so if we create an instance of a struct, we will create an instance of each of the variables that are a member of the struct. Just like with other variables, it is possible to copy structs with the assignment operator (`=`).

The program `struct.c` below illustrates how structs can be defined and used in C:

```
struct data {
    int x;
    int y;
};

int main(void) {
    struct data d;
    d.x = 10;
    d.y = 20;
    struct data e = d;
    printf("e.x=%d, e.y=%d\n", e.x, e.y);
    return 0;
}
```

Notice that we need to write `struct data` to name the struct. Simply writing `data` (like in C++) does not work, since the name of structs are in a different namespace by default.³

³It is possible to get around this using `typedef`, but this is not supported by Progviz and therefore not used in this book.

It is also worth pointing out that a struct may *only* contain variables. In contrast to a class in an object oriented languages, structs may *not* contain functions. We will see in Section 2.8 how we can use structs and functions to create abstract data types that are similar to “classes”.

2.5 Pointers

Pointers are a central part of C. A pointer is essentially an integer variable that we know contains the *address* of some other data in memory. Because they refer to some other data, it is usually convenient to represent them as an arrow that points to some other data. The idea that a pointer is just a numeric value is, however, important to keep in mind. It helps reasoning about the behavior of pointers, both regarding assignments of pointers and regarding the semantics of pointer variables in a concurrent program.

C provides special *pointer types* to denote which variables are pointers, and what type of data they refer to. This makes it more difficult to accidentally treat a pointer as an integer, or to accidentally interpret the data pointed to as an integer when it is actually something else. A pointer is expressed using an asterisk (*) after the name of a type (e.g., `int *` is a pointer to an `int`).

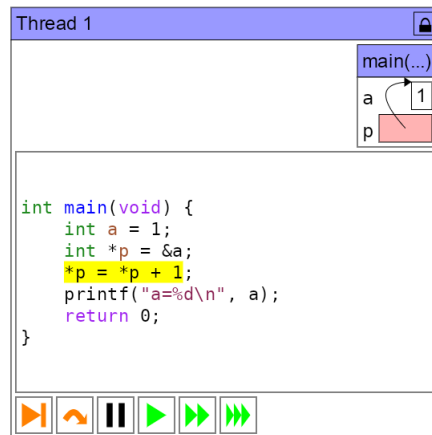
In addition to pointer types, C provides two operators that create pointer values and traverse pointers, respectively. These are illustrated in the program `pointers-1.c` below:

```
int main(void) {
    int a = 1;
    int *p = &a;
    *p = *p + 1;
    printf("a=%d\n", a);
    return 0;
}
```

The program first creates the integer variable `a` and initializes it to 1. The next line then uses the *address-of operator*, `&`, to compute the address of the variable `a` (i.e., `&a`). One way to think of this operation is that the `&` operator creates a pointer value that refers to `a`. The type of this pointer value is the type of `a` (i.e., `int` in this case) with an additional asterisk added at the end. In this case, this produces the type `int *` which matches with the declared type of the variable `p`.

At this point it is useful to see how Progvis represents pointers. Load the program `pointers-1.c` in Progvis by selecting *File* \Rightarrow *Open...* Then click the leftmost button in the *Thread 1* box a few times, and you will see the representation in Fig. 2.1. In particular, notice the smaller box labeled *main(...)* in the top-left corner. It contains a list of the local variables inside the `main` function. The first line contains the variable `a` and shows that it contains the value 1. The second line shows the variable `p`. Its content is the address where the variable `a` is located. To make it easy to see what happens, Progvis represents this as an arrow that points to the data in the variable `a`.

The next line in the program (the one highlighted in yellow in Fig. 2.1) contains the next pointer-specific operator, the asterisk (*). This operator is used to *dereference* a pointer, meaning that we examine the address in the pointer, go to that location in memory and read whatever value is there. In

Figure 2.1: Progvis' representation of pointers in the program `pointers-1.c`.

the case of `*p` in Fig. 2.1 it means that we look at the arrow stored in `p`, and instead of operating on `p` itself, we operate on the data that the arrow points to instead. The type of the expression `*p` is the type of `p` (i.e., `int *`) but with the rightmost asterisk removed, which in this case is just an `int`.

Practice: To see the difference between `p` and `*p`, change the line `*p = *p + 1` to `p = p + 1` and see what happens in Progvis. Remember that `p` refers to the pointer itself (i.e., the integer address that a pointer is), while `*p` dereferences the pointer and operates on that data instead.

At this point it is useful to spend a moment to compare the syntax used to denote pointer types with the operators that operate on pointers. We saw that the address-of operator (`&`) essentially creates pointers. Therefore it might be surprising that the same operator is not used to denote a pointer type. The reason for this is that types in C are designed to reflect how the variable is *used*. That is, one way to think of the declaration `int *p` is: “We have a variable `p`, which has a type so that the expression `*p` is an integer.” This is the reason why it might at first feel like the roles of `*` and `&` are swapped in type declarations compared to their usage in code. This is also one reason why this book uses the notation `int *x` rather than `int* x`.⁴

Since parameters in C are always passed by-value, pointers are often used to simulate by-reference semantics. This both allows passing large structs without copying them, but also allows functions to modify data provided by the caller easily. This is illustrated by the program `pointers-2.c` which is shown below. The program contains two functions, `main` and `f`. The `main` function creates a variable `a` that it wishes to let `f` modify. Therefore, it passes the address of `a` to `f` by using the address-of operator (`&`). The function `f` then receives the address as an `int *` parameter, and modifies the variable `a` from `main` by using the dereference operator (`*`). As such, the program prints `a=10`.

⁴The idea that variable declarations reflect the usage of the variable is also apparent if multiple variables are defined at once. For example, `int *x, y;` defines two variables. `x` is of type `int *` while `y` is of type `int`! Please don't do this.

```

void f(int *x) {
    *x = 10;
}

int main(void) {
    int a = 0;
    f(&a);
    printf("a=%d\n", a);
    return 0;
}

```

Again, it is useful to see how Progvis represents this situation. As such, open the program in Progvis and single-step the program until you reach the state in Fig. 2.2. The representation is very similar to Fig. 2.1. The difference is that *two* boxes are visible in the top-right corner. The topmost contains the local variables in `f` and the one that is partially covered contains the local variables in `main`. With this in mind, we can easily see that the pointer variable `x` refers to a variable in `main`. We do not see the name of the variable from this particular figure,⁵ but we can easily deduce that it has to be the variable `a`.

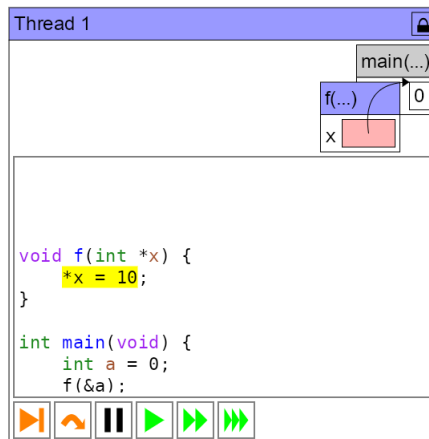


Figure 2.2: Progvis' representation of pointers in the program `pointers-2.c`.

Practice: Use Progvis to see what happens if you change `f` to accept a `int x` instead of `int *x`. What other parts of the program do you need to change, and what will the program print when you run it?

⁵This is actually intentional. The name `a` is not accessible from `f`, so the only way to access the variable is if we have a pointer to it!

One problem with the dereference operator (*) is that it is applied after the member access operator (.) operator. This is sub-optimal when working with pointers to structs, as illustrated by the program `pointers-3.c`:

```

struct data {
    int a;
    int b;
};

int main(void) {
    struct data d;
    d.a = 1;
    d.b = 2;

    struct data *p = &d;
    printf("a: %d\n", (*p).a);
    printf("b: %d\n", p->b);

    return 0;
}

```

Note that the program defines the pointer `p` that refers to `struct data`. If we wish to access the variable `a` inside `struct data`, we need to write `(*p).a`. We can not just write `p.a` since the member access operator (.) expects its left hand side to be a struct, not a pointer to a struct. This is why we need the dereference operator (*). Adding a dereference operator like `*p.a` will not solve the problem since it means `*(p.a)`. This is why we need an explicit parenthesis as in the first `printf` statement: `(*p).a`. The dereference operator first “converts” the pointer into a value, so that the member access operator receives a struct instead of a pointer as its left hand side.

This is, however, not very convenient. For this reason C provides the `->` operator as a convenience. The expression `x->y` is equivalent to `(*x).y`. As such, instead of writing `(*p).a` we can simply write `p->a`. This is what `pointers-3.c` does to access the variable `b` in the second `printf` statement.

2.6 Arrays

An array in C is a collection of elements of the same type that are stored after each other. Note that unlike some other languages (e.g., Java or `vectors` in C++) C stores *no* additional information about the array. In particular, this means that it is up to the programmer to keep track about the size of the array and to make sure to not access elements out of bounds. While this is sometimes problematic, it allows a convenient analogy between arrays and pointers. We use the program `arrays.c` to illustrate this.

```

int main(void) {
    int arr[3] = { 1, 2, 3 };
}

```

The listing above shows the first two lines of `arrays.c`. The program starts by defining a variable `arr` that contains array that is large enough to contain 3 elements. We can once more see that declarations in C are designed to resemble how the variable is used. As we will see, we can access elements in the array

using square brackets (e.g., `arr[0]`), which is why the size of the array appears *after* the variable name rather than as a part of the type (i.e., not `int[3] arr` as is the case in Java). The program also initializes the array to contain the numbers 1, 2, 3. As for other variables in C, initialization is optional. Note that the size of the array needs to be a constant value (either a number or the name of a constant).⁶ We will see how we can allocate dynamic arrays in Section 2.7.

```

{
    int first = arr[0];
    printf("First: %d\n", first);
}

```

The next part of the program reads the first element of the array using the square bracket operator (`arr[0]`) and prints it. Of course the 0 could be replaced with an arbitrary expression.

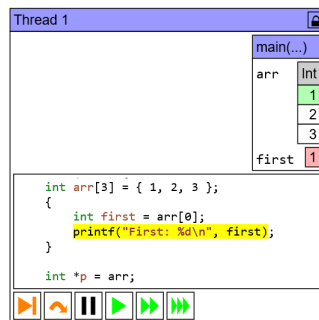


Figure 2.3: Progvis' illustration of the program `arrays.c` at the first call to `printf`.

Figure 2.3 shows how Progvis visualizes the program at this point. In the box labeled `main(...)` in the top-right part of the picture, we see that two variables (`arr` and `first`) are currently in scope in `main`. We can also see that the array (`arr`) is represented as a sequence of integer elements. Progvis labels the array with the type of the array elements, similarly to how it represents structs. The similarity between arrays and structs is no accident. Both are laid out sequentially in memory. A struct that contains 3 integer variables is actually identical to an array that contains 3 elements in memory.

```

int *p = arr;
{
    int first = *p;
    printf("First: %d\n", first);
}

```

The first line in the next part of the program shows an interesting property of arrays in C, namely that they are almost always automatically converted to

⁶Some compilers do, however, support dynamically-sized arrays as local variables as an extension to the standard. There are a number of problems with this extension, in particular that it is easy to accidentally overflow the stack without the ability to easily avoid it.

pointers when they are used.⁷ In this case, the program utilizes the automatic conversion to get a pointer that is stored in `p`. As we can see in Fig. 2.4, the automatic conversion produces a pointer to the first element of the array, which is how arrays are often represented in C. The next statement in the program then uses the dereference operator to access and print the first element of the array, as shown in Fig. 2.4. Again, note that `p` contains *no* information about the size of the array. It is therefore up to us as programmers to keep track of the size separately.

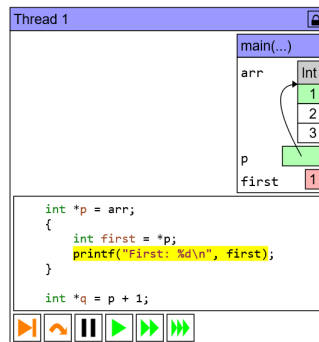


Figure 2.4: Progvis’ illustration of the program `arrays.c` at the second call to `printf`. Note that `p` is a pointer to the first element of the array.

```
int *q = p + 1;
{
  int second = *q;
  printf("Second: %d\n", second);
}
```

The next part of the program starts by doing something that is perhaps surprising. It creates a variable `q` and initializes it to `p + 1`. To understand why this makes sense, remember that pointers are really just integers that we interpret as the address of some other data in memory. In this interpretation, we can interpret `p + 1` to mean that we add 1 to the address stored in the pointer value. This is, however, not exactly what happens. To make it easier to work with arrays through pointers, C helps us a bit. Since we have a pointer to an integer, C will helpfully interpret `p + 1` as “compute the address of the next integer in an array”. As such, instead of just adding 1 to the address, it adds `1 * sizeof(int)` instead. As shown in Fig. 2.5, this gives us a pointer to the second element in the array which we can access as `*q`. Subtraction (both between two pointers and between a pointer and an integer) works similarly.

```
{
  int second = p[1];
  printf("Second: %d\n", second);
}
return 0;
}
```

⁷This is called that arrays *decay* into pointers. More or less the only situation where this does *not* happen is when using the `sizeof` operator to get the total size of the array.

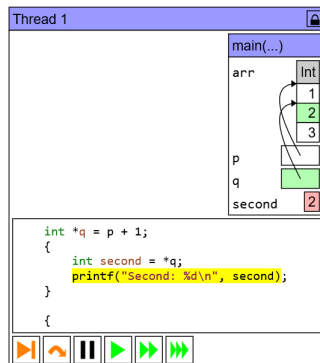


Figure 2.5: Progvis’ illustration of the program `arrays.c` at the third call to `printf`. Note that `q` contains the expression `p + 1`, which is a pointer to the second element of the array.

The final part of `arrays.c`, perhaps surprisingly, uses the square bracket operator to access the second element of the pointer `p`, even though `p` is *not* an array. This works because the expression `x[y]` is equivalent to `*(x + y)`.⁸ As such, `p[1]` simply means `*(p + 1)`, which as we saw earlier is equivalent to accessing the second element in the array.

This shows the deep connection between pointers and arrays. We previously saw that array variables are automatically converted into pointers when they are used. We have also seen that this conversion does not really matter, since we can treat a pointer as an array anyway, since the square brackets can be used regardless. In fact, the expression `arr[0]` we used to access an element in the array is no different from what we are doing here. The array `arr` is first implicitly converted to a pointer, and then the square bracket operator adds zero and dereferences the pointer in order to access the first element.

2.6.1 Strings

As we noted before C does not have a special type for representing strings. Instead, C represents strings as an array of characters (i.e., the `char` type). Since it is usually not possible to get the length of an array, C strings end with the NUL character (i.e., `'\0'`) by convention. This idea is illustrated by the program `strings.c`:

```
int main(void) {
  const char *str = "test";
  printf("%s\n", str);
  printf("%c\n", str[0]);
  return 0;
}
```

The program creates a variable `str` of the type `const char *` — a pointer to an array of constant characters. The `const` is needed since we are not allowed

⁸A fun consequence of this is that since `*(x + y)` is the same as `*(y + x)`, it means that `x[y]` is the same as `y[x]`. This means that `p[1]` is the same as `1[p]`, which is a very confusing way to access elements in an array. Please don’t do that.

to modify the contents of the string literal (it is stored in a portion of memory that is usually write-protected).

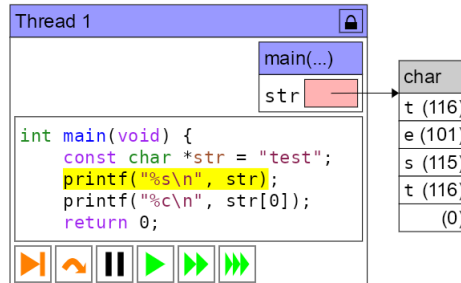


Figure 2.6: The program `strings.c` in Progvis.

Figure 2.6 shows how strings are visualized by Progvis. From the figure, we can see that `str` is a pointer to an array of characters. Each element is an 8-bit number as shown in parenthesis. To aid readability, Progvis also shows which character the number corresponds. Of particular note is that even though the string `test` only consists of 4 characters, the array has 5 elements so that there is room for a final `'\0'` at the end to signal the end of the string.

The C standard library provides a few utility functions to manipulate strings. The ones relevant to this book are:

```
#include <string.h>
```

The utilities are available in the header `string.h`. Perhaps notably, the constant `NULL` is also defined in this header.

```
int strlen(const char *);
```

Computes the number of characters in a string. Note that the final NUL character is *not* counted. As such `strlen("test")` returns 4, not 5.

```
char *strdup(const char *);
```

Creates a copy of the string by allocating memory for the string using `malloc`. The copied string therefore needs to be `free`'d later.

```
int strcmp(const char *, const char *);
```

Compares the contents of the two supplied strings. If they are equal, 0 is returned. Otherwise a positive or a negative number is returned to indicate whether the first string is lexicographically before or after the second string.

2.7 Memory Management

So far we have only used global and stack-allocated data. C also supports dynamic memory management to make it possible to work with dynamic data structures, such as dynamic arrays. This is achieved through the functions `malloc` and `free` that are defined in the `<stdlib.h>` header.

```
void *malloc(size_t size);
```

Allocates `size` bytes of memory and returns a pointer to it. If the allocation fails (e.g., because there is not enough available memory), `NULL` is returned. Note that `size_t` is an integer type similar to `int`, but it is unsigned and sometimes larger than an `int`. The contents of the allocated memory is undefined.

Typically the size supplied to `malloc` is computed using the `sizeof` operator.⁹ Either as `malloc(sizeof(T))` or `malloc(sizeof(T) * x)`.

```
void free(void *memory);
```

Deallocates memory that was previously allocated by `malloc`. This makes the memory available for future allocations by `malloc`. After memory is `free`'d, it is an error to use that piece of memory.

The program `malloc-1.c` illustrates how `malloc` and `free` can be used to allocate structs that represent nodes in a singly linked list.

```
struct node {
    int value;
    struct node *next;
};

int main(void) {
    struct node *n = malloc(sizeof(struct node));
    n->value = 0;

    n->next = malloc(sizeof(struct node));
    n->next->value = 1;

    free(n->next);
    free(n);
    return 0;
}
```

One thing particularly worth noting is that it is *not* necessary to convert the `void *` returned by `malloc` into `struct node *`. C treats `void *` as a pointer to some data, but trusts the programmer to keep track of what kind of data. As such, in assignments and when calling functions, C allows implicitly converting a `void *` into any other pointer type (it does not even produce warnings). This makes usage of `malloc` quite convenient.¹⁰

Progvis illustrates heap-allocated memory as shown in Fig. 2.7. Note that memory allocated by `malloc` (also known as *heap-allocated data*) is drawn outside of the box labeled *Thread 1*. That is, all stack-allocated data is inside the box of a thread (in the top-right corner, that represents the stack) while all heap-allocated data is drawn outside the boxes of threads.

Before terminating, the program frees the memory previously allocated using `malloc` by calling `free`. Figure 2.8 shows how Progvis illustrates the situation. First and foremost, the rightmost node has a red cross drawn on top of

⁹In fact, Progvis requires that `sizeof` is used, otherwise it does not know the type of the data that is allocated.

¹⁰Another reason for not including the explicit type cast is that Progvis does not support `void *` or type-casts in an effort to be type-safe.

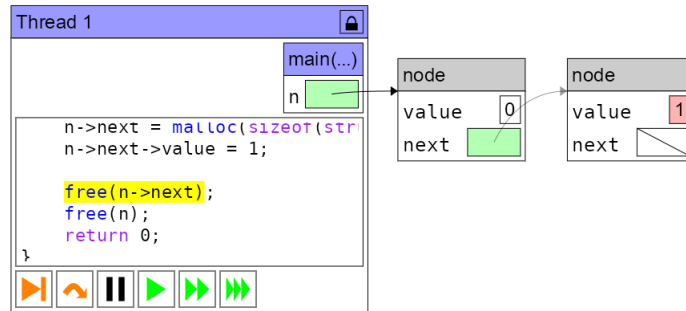


Figure 2.7: The program malloc-1.c visualized by Progvis.

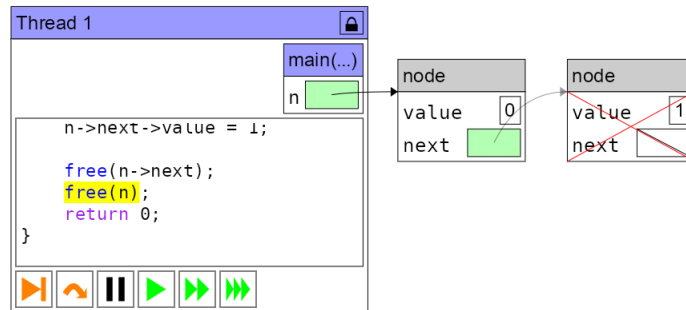


Figure 2.8: The program malloc-1.c after freeing the second node.

it to indicate that the memory is freed. We can also notice that `n->next` still points to the memory that was freed. Since `free` is a function just like any other, it can not modify the parameter we pass to it, and therefore it is unable to set the pointer to `NULL` for us. Even though we *could* still try to access data in `next`, such as `n->next->value`, doing so is undefined. Progvis will notify us if this happens.

It is also possible to use `malloc` to allocate dynamic arrays by simply allocating space for more than one element at once. This is illustrated by the program `malloc-2.c`:

```

struct elem {
    int a;
    int b;
};

int main(void) {
    struct elem *array = malloc(sizeof(struct elem) * 3);
    array[0].a = 0;
    array[0].b = 1;
    array[1].a = 2;
    array[1].b = 3;
    free(array);
    return 0;
}

```

Notice that the program calls `malloc` with the size `sizeof(struct elem) * 3`, which means that we requires memory that is large enough to contain 3 instances of `struct elem` back-to-back. Since pointers can be viewed as arrays, we can then use the pointer returned from `malloc` as any other array.

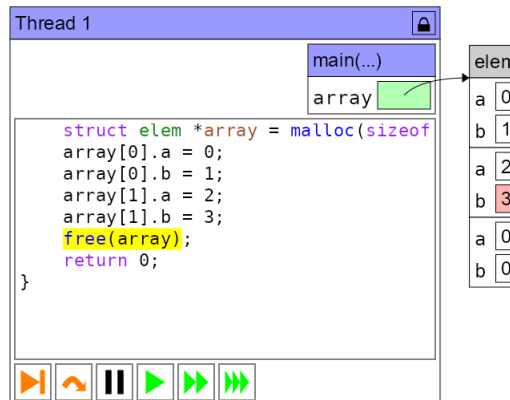


Figure 2.9: Progvis illustrating dynamic arrays allocated by `malloc-2.c`.

In Fig. 2.9, we can see how Progvis illustrates the array. As before, we can see that the array is outside the *Thread 1* box since it is heap-allocated. The representation of the array itself is otherwise identical to stack-allocated arrays. In this case the array is different from what we saw before since we allocated an array of structs, rather than an array of integers or characters.

2.8 Abstract Data Types

Structs in C can only contain variables. In particular, it is not possible to create member functions like in other object-oriented languages. This does not mean that it is impossible to create abstract data types in C, they just have to look a bit different compared to other languages. There are different ways to do this. This chapter introduces the conventions used throughout this book.

To illustrate these conventions, we will implement a simple abstract data type `array` that represents a simple dynamic array (i.e., a simpler version of `vector` from C++). In terms of object oriented programming, the goal is to create a class (`array`) with private data members and public member functions.

To represent this idea in C, we start by creating a struct that represent the abstract data type itself. As shown below, the struct contains an array of elements (`data`) and the number of elements in the array (`count`).

```
struct array {
    int *data;
    int count;
};
```

The next point is to introduce a convention to determine which functions are a part of the abstract data type (i.e., which functions do we consider to be member functions in the “class”). All other functions should treat the data

type as a black box, and should not access the variables inside `struct array`.¹¹ The convention used in this book is that all functions we consider to be a part of an abstract data type starts with the name of the data type followed by an underscore (i.e., `array_` in this case). Furthermore, they should all accept a pointer to an instance of the data structure as their first parameter (this corresponds to the implicit `this` parameter in many OOP languages).

Using this convention, we can implement all functions that manipulate the contents of our simple dynamic array. We leave construction and destruction of the abstract data type for later, as they have special semantics. In particular, we implement functions to get the number of elements (`array_count`), to get an element by its index (`array_get`), to set an element by its index (`array_set`) and to print the contents of the array (`array_print`) as follows:

```
int array_count(struct array *a) {
    return a->count;
}

int array_get(struct array *a, int pos) {
    return a->data[pos];
}

void array_set(struct array *a, int pos, int value) {
    a->data[pos] = value;
}

void array_print(struct array *a) {
    for (int i = 0; i < array_count(a); i++)
        printf("%2d: %2d\n", i, array_get(a, i));
}
```

What remains is to create functions to create and destroy the data structure. There are two ways to structure this, each with its own benefits and drawbacks. The difference is whether we let the function that creates the data structure allocate memory for the data structure, or leave it to the user of the data structure. We therefore refer to them as *external allocation* and *internal allocation*.

2.8.1 External Allocation

The first option is to let the user of the abstract data type allocate memory for the struct. The memory allocation is therefore *external* to the abstract data type itself. Because of this, the function that creates an instance of the abstract data type only needs to initialize the struct to a known state. As such, we call the function `<type>_init` (i.e., `array_init` in this case), and let it accept a pointer to the memory that should be initialized as its first parameter. In the case of `struct array`, the `array_init` function allocates an array of suitable size so that it can initialize `data` and `count` as below:

¹¹In this book, we do not leverage C to enforce this. There are ways to hide the definition of `struct array` so that it is not *possible* for code outside of the abstract data type to access variables in `struct array`. This is, however, out of scope for this book.

```

void array_init(struct array *a, int count) {
    a->data = malloc(sizeof(int) * count);
    a->count = count;
}

```

Since the abstract data type contains dynamically allocated memory, we also need to provide a function for destroying the struct. By convention, we call this function `<type>_destroy` (i.e., `array_destroy` in this case). It accepts a pointer to the instance that should be destroyed, and is responsible for freeing any resources allocated by the abstract data type. Note, however, since the `array_init` function did *not* allocate memory for the struct itself, `array_destroy` will not free the struct. Both functions leave memory management of `array` up to the user of the abstract data type. In this case, `array_destroy` can be implemented as below:

```

void array_destroy(struct array *a) {
    free(a->data);
}

```

Since `array_init` and `array_destroy` do not manage the memory of `struct array`, it is the responsibility of the user to manage the memory of `struct array`. This gives the user flexibility to allocate the memory wherever is convenient in each situation. For example, the `struct array` could be allocated on the stack, as a global variable, or in another struct. The simplest case is perhaps to allocate it on the stack as in the code below:

```

{
    struct array a;
    array_init(&a, 2);
    array_set(&a, 0, 1);
    array_set(&a, 1, 2);
    array_print(&a);
    array_destroy(&a);
}

```

The code above first allocates memory for a `struct array` in the form of a local variable. It then initializes the memory by calling `array_init` (and passing a pointer to the data structure). It then sets element 0 to 1 and element 2 to 1 before printing the contents of the array. Finally, it destroys the abstract data type by calling `array_destroy`.

2.8.2 Internal Allocation

The other second option is to let the abstract data type itself handle allocation of the struct. This removes the ability of the user to choose *how* the struct is allocated. It does, however, make it possible to hide the contents of the struct from the user, thereby making it impossible to accidentally access the contents of the struct from functions that do not belong to the abstract data type.

Since the abstract data type is in charge of managing memory for its struct, the function that creates the abstract data type can simply return the created instance (in contrast to `array_init` above). By convention, we create this function `<type>_create` (i.e., `array_create` in this case). It does not need to

accept additional parameters, but is expected to return a pointer to the newly created instance. It can be implemented as follows:

```
struct array *array_create(int count) {
    struct array *v = malloc(sizeof(int) * count);
    v->data = malloc(sizeof(int) * count);
    v->count = count;
    return v;
}
```

Note that the implementation above contains an additional call to `malloc` compared to `array_init`. Apart from that, `array_init` and `array_create` are very similar.

Similarly to before, we also need to provide a function for deallocating the abstract data type. In this case, this function also frees the memory for the struct, since the memory was allocated by the `<type>_create` function. Therefore, we call this function `<type>_free`. For the `array` type, it can be implemented as follows:

```
void array_free(struct array *v) {
    free(v->data);
    free(v);
}
```

Again, note that `array_free` has an additional call to `free` compared to `array_destroy`.

The differences between `array_init` and `array_create` as well as `array_destroy` and `array_free` also affects how the abstract data type is used. In particular, since the abstract data type is in charge of managing memory for the struct, the user of the data type just needs to store the address of the struct in a pointer. This can look like below:

```
{
    struct array *a = array_create(2);
    array_set(a, 0, 1);
    array_set(a, 1, 2);
    array_print(a);
    array_free(a);
}
```

Concurrent Programming

Many operating systems allow programs to start multiple *threads* that run concurrently. These programs are sometimes called *multithreaded programs* or *concurrent programs*. We will use the term *concurrent programs* in this book.

As we shall see, writing correct concurrent programs is not as simple as starting additional threads. There are a number of additional concerns that need to be taken into account when multiple threads need to co-exist and share resources within a single process. To properly understand these concerns, we need a fairly detailed model of program execution. This chapter thus focuses on this model. We start by examining the model in terms of normal, sequential, programs and then extend it to also cover concurrent programs. We will then use the model to illustrate some of the problems that arise when multiple threads need to co-exist in a single program.

3.1 Execution Model

We will start by looking closer at how sequential programs are executed. This helps understanding why the execution model for concurrent programs is more complex compared to the model for sequential programs.

As the name implies, sequential programs can be considered to be executed in sequence, one statement after the other. To illustrate this, consider the program in Listing 3.1. The program contains two global variables, `a` and `b`. The C language guarantees that they are initialized to zero at the start of the program since they are global variables.¹ The `main` function manipulates the variables and then prints the value of them. In particular, it first stores the value 10 in the variable `a` and the value 20 in the variable `b`. After that, it reads the value in `a`, multiplies it with 5 and stores the result in the variable `a`. After that, it reads `a` and `b` in some order, adds the values together and stores the result in `b`. At the end of all of this, `a` contains 50 and `b` contains 70 as we can see by running the program.

The largest benefit of this sequential model of program execution is that it is easy to understand. This is one of the reasons why Progvis uses the model

¹As you probably know, this is *not* the case for stack-allocated variables in functions.

```

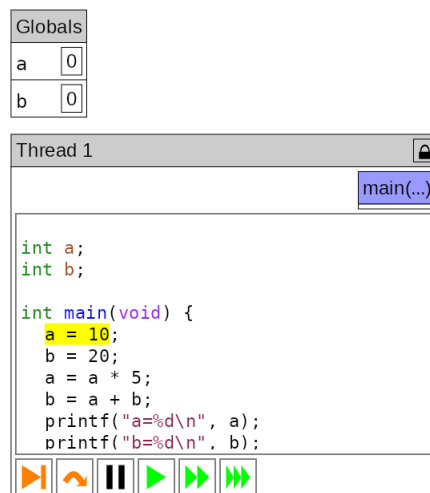
int a;
int b;

int main(void) {
    a = 10;
    b = 20;
    a = a * 5;
    b = a + b;
    printf("a=%d\n", a);
    printf("b=%d\n", b);
    return 0;
}

```

Listing 3.1: A simple program that manipulates global variables.

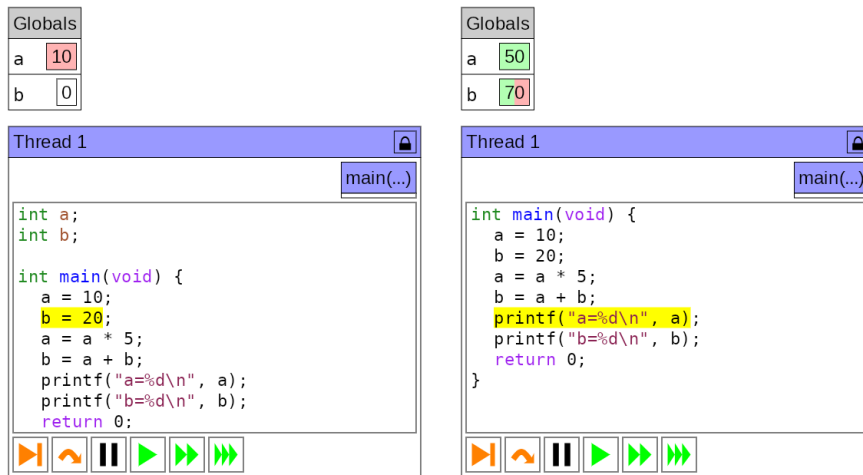
to visualize programs, and why the model that describes concurrent execution will be based upon it later on.

Figure 3.1: Progvis after loading the program `seq-a.c`.

To see how Progvis visualizes the program, we can open it by starting Progvis, selecting *File* \Rightarrow *Open program...*, and opening the file. After doing this, Progvis should display the contents of Fig. 3.1. In particular, we can see that the global variables `a` and `b` are both zero, and that Progvis paused the program just before executing the line `a = 10`, since that line is highlighted in yellow.

At this point, we want to turn off some more advanced parts of the visualization. Select *Run* \Rightarrow *Report data races* \Rightarrow *Disabled*. After that, we can inspect each step of the program execution by clicking the leftmost button in the box *Thread 1*. Each click executes the statement that is highlighted in yellow and pauses it again before the next statement. At each step, Progvis also highlights the memory locations the program accessed. Reads are colored in green and writes are colored red. For example, after executing the line `a`

= 10, Progvis shows that `a` was written by coloring the background in red as can be seen in Fig. 3.2a. In cases where a single variable was both read from and written to in the same statement, such as after executing `b = a + b`, it is highlighted in both colors as is shown in Fig. 3.2b.

(a) After running `c = a * 5`(b) After running `b = a + b`Figure 3.2: Progvis executing the program `seq-a.c`.

We can use this step-by-step model to reason about the program. This program has two distinct stages. First it manipulates data in memory (i.e., the variables `a` and `b`) and then it prints them using `printf`. In normal program execution, we are only able to observe the second stage, the output from `printf`. Because of this, the important part of the program is, in some sense, the call to `printf`. Most of the time we don't care how the computations leading up to `printf` were performed, as long as they produce the same result as the program we have written.

This idea is general enough that it is applicable to *any* program. For the purposes of this discussion, we will only consider programs running in user-mode.² Since programs in user-mode are isolated from each other, their behavior is only visible to the rest of the system through the system calls they perform. As such, we can think of program execution as the pseudocode below:

```
int main() {
    while (true) {
        void *syscall_args = computations();
        system_call(syscall_args);
    }
}
```

That is, a program can be thought of as performing some amount of computations that affect the contents of memory, followed by a system call. This process is repeated until the result from the computations decide to call the

²The reasoning works for kernel-mode and embedded systems as well. But since we can not think about system calls in those contexts, we need another definition of observability.

system call `exit` to terminate the program. To illustrate this idea, we can think of the program `seq-a.c` as three iterations of the loop above. The first call to `computations()` performs the assignments to the variables `a`, `b`, and `c` and returns the arguments for the call to `printf`.³ The second call to `computations()` only needs to return the arguments for the second call to `printf`. The second system call would thereby `printf` again. The final call to `computations()` determines that the program is done and returns the parameters to call `exit`, which becomes the third and final system call performed by the program.

This way of thinking about programs may initially seem strange and convoluted. It does, however, give us one important insight: namely that it does *not* matter *how* the computations are performed. The only thing that matters is that `computations()` modifies global state and returns `syscall_args` so that the system call behaves properly.⁴ Since it is only possible to observe the state at each system call, it does not matter in what order `computations()` is performed the computations, as long as the end result looks *as if* the program was executed line by line. Similarly, it does not matter that all computations were actually performed, again as long as the end result is *as if* they were performed.

<pre>int a; int b; int main(void) { a = 10; a = a * 5; b = 20; b = a + b; printf("a=%d\n", a); printf("b=%d\n", b); return 0; }</pre>	<pre>int a; int b; int main(void) { a = 50; b = 70; printf("a=%d\n", a); printf("b=%d\n", b); return 0; }</pre>
(a) <code>seq-b.c</code>	(b) <code>seq-c.c</code>

Figure 3.3: Two equivalent implementations of `seq-a.c`.

To illustrate the implications of this observation, consider the two versions of `seq-a.c` in Fig. 3.3. The difference between the original version and the left-most version Fig. 3.3a is that the two lines `b = 20` and `a = a * 5` are swapped. Since the line `a = a * 5` does not use the variable `b` at all, we can quite easily see that this change does not affect the behavior of the program in terms of the final outputs from the program.

Similarly, but perhaps surprisingly, the program in Fig. 3.3b is also equivalent to both `seq-a.c` and `seq-b.c` in terms of output: if we were to observe the value of the global variables `a` and `b` at each system call, we would not

³For the purposes of this discussion, we can consider `printf` to be a system call, even though it is mostly implemented in usermode. It does, however, eventually invoke the system call `write` to output characters to the screen.

⁴The parameters to the system call may contain pointers to any memory, this is why we also need to consider at least a subset of the global state to be observable when a system call is performed.

be able to tell them apart. This is true even though `seq-c.c` has removed all computations that `seq-a.c` and `seq-b.c` were doing and replaced them with the results of those computations.

This notion of equivalence is important, as it is the underlying reason for why compilers can perform any form of meaningful optimizations at all. We consider two programs as equivalent if they have the same *observable* behavior. In this case, we can think of *observable* in terms of the state that is observably through the system calls they perform (i.e., which system call, and any data passed to the system calls either directly or indirectly via pointers). This also gives rise to the *as-if rule* in C and C++: the compiler is allowed to re-structure your code as much as it likes, as long as the modified version behaves *as-if* the original version was executed. As such, all versions of `seq-a.c` above are legal optimizations that a compiler may choose to do. The third one (`seq-c.c`) in particular is a common optimization called *constant-folding*. If you compile `seq-a.c` with anything but the most basic compiler with optimizations turned on, the compiler will certainly produce a program that looks very much like `seq-c.c`, even if you compiled `seq-a.c`.

While we consider the three programs `seq-a.c`, `seq-b.c`, and `seq-c.c` to be equivalent in terms of their observable behavior, they behave quite differently internally. For example, it is quite easy to tell them apart if we are able to observe all memory accesses to `a` and `b`. If we run `seq-a.c`, we will see that the value 50 is stored in `a` *after* 20 has been stored in `b`. If we run `seq-b.c` instead, the order will be reversed. Similarly, if we run `seq-c.c` we will observe that both `a` and `b` are written to once, rather than twice in the other two programs. As we saw above, these differences are usually not a problem, since the program is still executed *as if* the original program was executed and produces the same result in all cases. In fact, we *expect* our compiler to make these types of changes, since we want the compiler to optimize our code. It turns out that this kind of changes are safe to do for sequential programs in general, since it is not possible to observe the differences between the versions just by observing the program's behavior.⁵ This is *not* true for concurrent programs. As we shall see, concurrent programs *are able to* observe these transformations (both by the compiler and by the hardware), and it turns out that some of these transformations will break concurrent programs. The difficult part about concurrent programming is therefore not to start threads, but rather to avoid the *problematic* transformations from making our programs behave incorrectly.

3.2 Threads, Processes and Programs

As mentioned in the beginning of this chapter, the difference between a sequential program and a concurrent program is that a concurrent program “contains” more than one thread. However, to properly understand what this means we need to take a brief detour and define some more specific terms:

Program The term *program* refers to the overall behavior of some piece of software. It is often useful to think of this in terms of the *source code* of the program. The source code specifies how the program should behave.

⁵We exclude execution time from what we consider to be observable. We *want* the compiler to reduce execution time as much as possible!

We can then compile the source code into *machine code* that is an alternative representation of the same program. Note, however, that *program* refers to a passive entity. A program by itself in this context does not *do* anything, it is just a specification of some computation that we can execute.

Process If we execute a program, we get a *process*. In contrast to a *program*, a *process* is therefore an *active* entity that is actively performing some form of computation. More concretely, a *process* is a piece of memory in the operating system that the process may use as it sees fit. One part of this memory contains the *machine code* of the *program* that the process is executing. The memory also contains other state, such as the value of global variables (e.g., `a` and `b` from the programs above), open files, etc. Generally different processes do not share memory,⁶ and they are therefore isolated enough to not interfere with each other. Furthermore, if we start the same *program* more than once, each invocation of the program becomes its own *process* with its own memory space. The different processes therefore have *their own copy* of all state, even global variables. This means that even if we run the program `seq-a.c` in multiple processes at the same time, the different processes will not interfere with each other since each process has its own copy of the global variables.

Thread Each process contains one or more *threads*. Each thread can be thought of as a virtual CPU core that executes code in the process. As such, threads are responsible for actually performing the computations described in the *program*, using the memory provided to them by the *process* they are associated with. As such, a process that contains zero threads is not able to do anything, not even launch new threads. Because of this the operating system automatically creates a thread that executes the `main` function when it creates a process for a program. This thread is sometimes referred to the *main thread* of the program. The *main thread* is usually a bit special, since the operating system terminates the entire process when the main thread terminates.

One important aspect of threads is that they belong to exactly one *process*. Each thread is able to access the memory of that particular process only. This means that two threads that belong to the same process share memory with each other, while two threads that belong to different processes do not.⁷

The relation between these terms are also depicted in Fig. 3.4. To the left is the program `multi-a` (a version of `seq-a` which we will use shortly). As before, we can think of this entity as the abstract behavior of the program, represented by the code stored on disk. The right part of the figure shows two processes that are both running the program `multi-a`. As mentioned previously, the processes are independent even though they are running the same program. This is illustrated in `fig:program-process-threads` by both processes having their own copy of the code in the program, as well as their own copy of global variables.

⁶Processes may actively *ask* the operating system to share a piece of their memory, but that is outside the scope of this book.

⁷See the note about explicitly shared memory above, however.

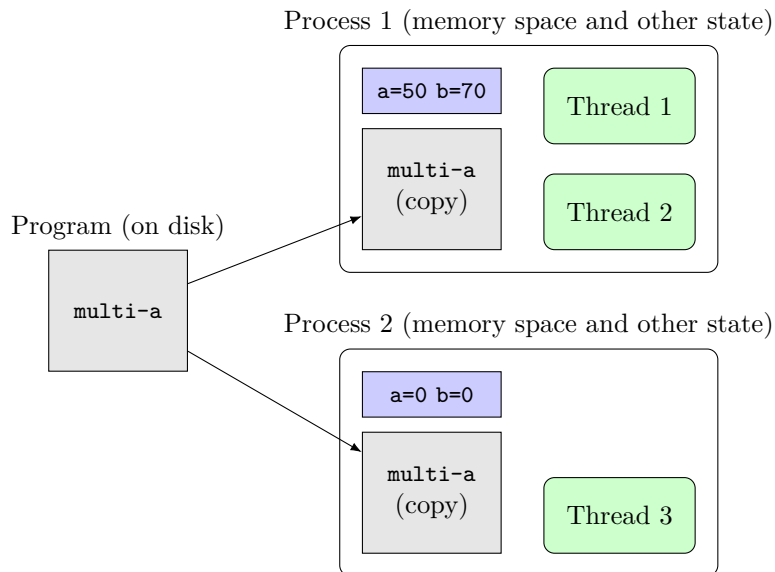


Figure 3.4: The relation between the terms *program*, *process*, and *thread*.

The figure also shows that a total of three threads are running. Threads 1 and 2 belong to process 1, while thread 3 belongs to process 2 (perhaps process 2 has not yet had time to start its second thread). Since threads 1 and 2 execute code in the context of the same process, they share memory with each other. In particular, both have access to all memory allocated to process 1. Thread 3 does not share anything with threads 1 and 2 since it belongs to a different process.

With a clearer view of these terms, we can more clearly define what we mean with concurrent programming. This book focuses on concurrent programming in the *shared-memory* model. That is, concurrent programming with multiple threads that communicate through memory accessible to all threads. This is the case for threads 1 and 2 in the picture above. We would therefore consider process 1 to be a process that uses concurrency. This is not the case for process 2 since it has not (yet) started additional threads. The program `multi-a` as a whole is, however, a concurrent program since it launches additional threads at least in some situations (since process 1 contains multiple threads). It is worth noting that even if we launch multiple process of a sequential program (such as `seq-a`), we would not consider the program to be concurrent since the two threads belonging to the two processes would not share memory.⁸

⁸There are, however, other forms of concurrent programming which covers this type of applications, such as message-passing.

3.3 Starting Threads

The threading library introduced in Chapter 1 with this book provides the function `thread_new` that starts a new thread in the process of the calling thread.⁹ The function expects a pointer to the function the new thread should execute, followed by zero or more pointers that are passed as parameters to the function in the new thread.

```
int a;
int b;

void thread_fn(void) {
    a = 10;
    b = 20;
    a = a * 5;
    b = a + b;
    printf("a=%d\n", a);
    printf("b=%d\n", b);
}

int main(void) {
    thread_new(&thread_fn);
    thread_fn();
    return 0;
}
```

Listing 3.2: Modified version of `seq-a`.

The program `multi-a.c` (in Listing 3.2) is a modified version of `seq-a.c` that uses `thread_new` to start an additional thread. As can be seen from the figure, the code that was previously in the `main` function has been moved to a function named `thread_fn`. The `main` function first starts a new thread using `thread_new` and then calls `thread_fn`. Since `thread_fn` is passed as the first parameter to `thread_main`, the new thread will also execute `thread_fn`. The end result is therefore that we have two threads that both run `thread_fn` concurrently.

To explore exactly what this means, and what implications this has on program behavior, we use Progvis. Open Progvis and open the program `multi-a.c`. Once again, select `Select Run ⇒ Report data races ⇒ Disabled` to tell Progvis that we do not want to be informed about the concurrency issues in this program. As we shall see, there are many of them, so we start by exploring the program ourselves without having Progvis helping us identify the issues.

After loading the program, Progvis will show the starting point of the program like in Fig. 3.5. As before, Progvis has paused the program right before the first statement in the `main` function. That is, program execution was paused *before* the program had the chance to call `thread_new`, and therefore only one thread, the *main thread* is displayed currently.

If we ask Progvis to let the thread continue to the next statement, it will execute `thread_new`, which creates a new thread. Progvis illustrates the new thread as a new box labeled *Thread 2*. The contents of this box is similar to the box for thread 1, but with some notable differences. In particular, they

⁹On Linux, the underlying library call is `pthread_create`, and on Windows it is `CreateThread`.

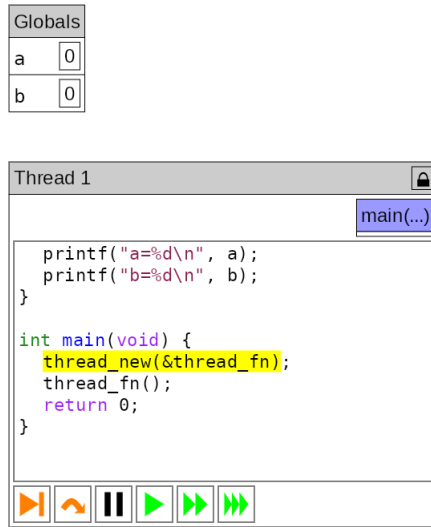


Figure 3.5: Progvis after loading the program multi-a.c.

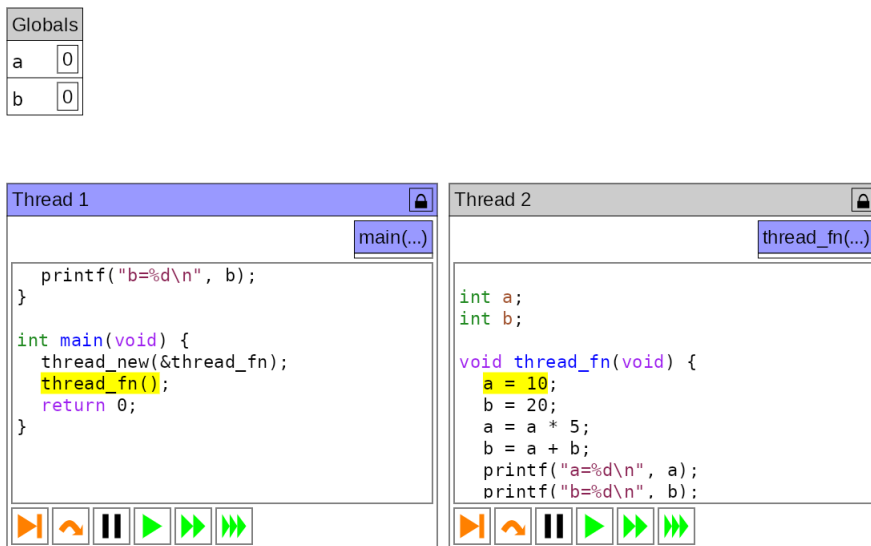


Figure 3.6: Progvis after letting thread 1 create an additional thread.

are about to execute different pieces of the program. Thread 1 is about to call `thread_fn`, while thread 2 is already inside `thread_fn` and is about to execute the line `a = 10`. Another difference is that the call stack is different. Progvis displays the call stack in the top-right corner of each box, right below the padlock. The representation in Fig. 3.6 is not very interesting at since neither `main` nor `thread_fn` have any local variables. We can, however, see the name of the currently executing function. If we ask Progvis to step thread 1 once by clicking the leftmost button in the box labeled Thread 1, it will jump into the call to `thread_fn`, and we will see the representation in Fig. 3.7.

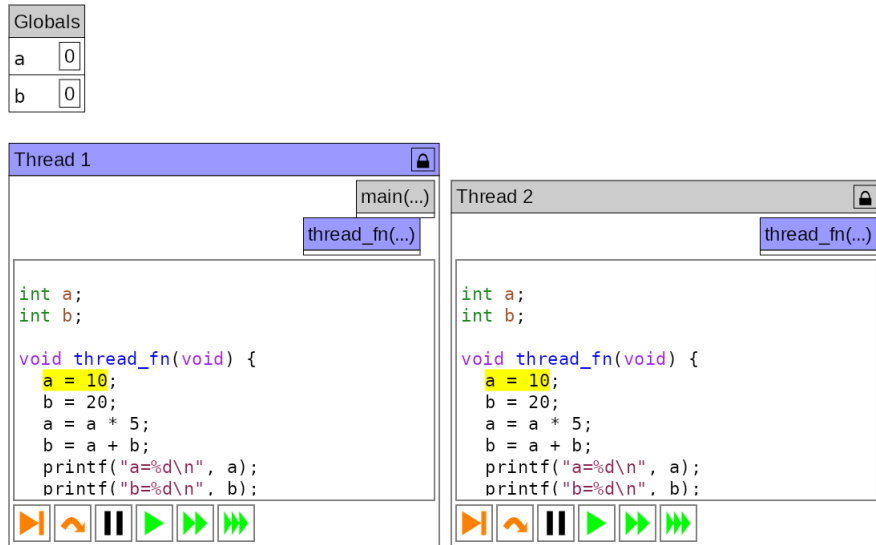
Figure 3.7: Progvis after letting thread 1 jump into `thread_fn`.

Figure 3.7 shows the representation of the stack trace in a bit more detail. In the top right corner of Thread 1's box, there are now two squares. The bottom one that is highlighted in blue is labeled `thread_fn`, and the top one is labeled `main`. This means that once thread 1 has finished executing `thread_fn`, program execution will resume inside `main`. The box labeled `main` is actually drawn behind the `thread_fn` box. This will be more apparent in future examples, where the boxes will also contain any local variables in the functions. In contrast to thread 1, thread 2 only contains one box labeled `thread_fn` in its upper right corner. This means that once thread 2 finishes executing `thread_fn`, the thread will terminate. Thread 2 will therefore never return back to `main`.

One key observation we can make at this point is that each thread has their own set of controls in Progvis. So, instead of choosing to step thread 1 before, we could also have chosen to step thread 2, which would of course have caused the program to end up in a different state. This is how Progvis represents *concurrent execution*. When two or more threads execute concurrently, the user is able to determine in which order the threads should be allowed to execute, and how long each thread should be allowed to execute before other threads get a chance to execute. For example, we can choose to let thread 1 run to completion before thread 2 gets a chance to run. We could equally well choose to let thread 2 run to completion before letting thread 1 continue, or anything in between.

This might seem odd to let the user choose between such a wide range of possible executions. However, it turns out that this matches what is allowed to happen when multiple threads coexist. All choices you can make when using Progvis are choices the scheduler in your operating system can make when your program is running. As such, the goal to keep in mind when writing concurrent programs is that the program should behave correctly for *all* possible executions. That is, regardless of which order you press the *step* button for

different threads in a program, the program should behave correctly. We call one such order for an *interleaving*, since each such order corresponds to one way to interleave the statements of the different threads. As you might imagine, ensuring that the program behaves correctly in spite of this non-deterministic execution is what is difficult about concurrent programming.

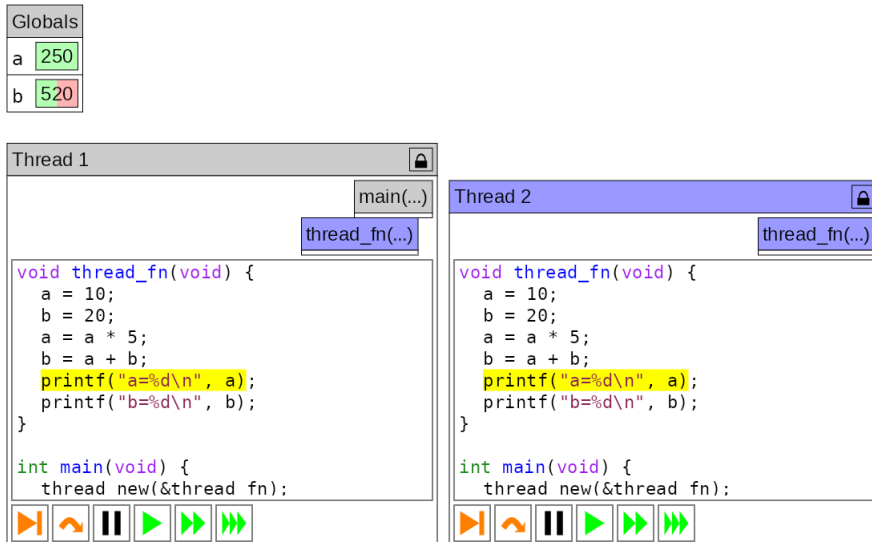


Figure 3.8: Progvis after loading the program multi-a.c.

Practice: There are many different interleavings of the few lines in multi-a.c. Some of them cause both threads to output the expected result (a=50 and b=70), while others produce more unexpected results. Find an interleaving of the two threads that leads to the state shown in Fig. 3.8, where a is 250 and b is 520. Use the menu item *Run* \Rightarrow *Restart program* (Ctrl+R) to quickly restart the program from the beginning and try a new interleaving.

One important observation from the task above is that the possible interleavings are highly dependent on the code that is actually executed. For example, we previously concluded that we would not be able to distinguish between the programs seq-a.c, seq-b.c, and seq-c.c by running them since they have the same observable behavior. This is *not* true if we introduce multiple threads in our programs. If we were to explore the programs multi-a.c, multi-b.c, and multi-c.c (which are modified versions of seq-a.c, seq-b.c, and seq-c.c respectively), we would see that the possible outcomes differ, even though we previously concluded that the programs were equivalent.

Practice: Examine the programs `multi-b.c` and `multi-c.c` like you examined `multi-a.c` above. Not all programs are able to print the same values. For example, it is not possible to get `multi-c.c` to print `a=250` and `b=520`. However, both `multi-a.c` and `multi-c.c` are able to do so. Which of the programs (`multi-a.c`, `multi-b.c`, and `multi-c.c`) are able to produce the following outputs?

- `a=50` and `b=70`
- `a=50` and `b=80`
- `a=250` and `b=320`

The reason why we are able to differentiate between the programs `multi-a.c`, `multi-b.c`, and `multi-c.c`, even though we previously concluded that they were equivalent, is that our assumptions regarding *observability* no longer hold. A core assumption we made was that it is only possible to inspect the program's internal state (e.g., global variables) at each system call. This assumption is *no longer true* since the two threads in the program share memory and are therefore able to observe the program state at any time. This is the reason why we are able to produce different results from the different versions of `multi-*.c`.

3.4 Scheduling

Before we move on to see the full implications of the strange observations above, we will take a brief detour to examine how different threads are scheduled in practice. This helps understanding *why* the model for concurrent program execution looks the way it does (in particular, why we can make few assumptions in general).

Modern systems tend to have multiple physical CPU cores. Each CPU core can be thought of as an independent CPU that is able to execute its own stream of instructions in sequence. This maps very well to how threads work. We can therefore think of each CPU core as a hardware representation of a thread in our program. For example, a CPU with 4 cores can therefore execute 4 threads by allocating each thread to one of the CPU cores. In this case, we say that the threads execute *in parallel*, meaning that the threads execute instructions more or less at the same time. If the program `multi-a.c` would be executed in this way, it is possible that the two threads execute the line `a = a * 5` at the *same* time, not one after the other. As you might imagine, this can lead to interesting results, even more so than what we saw in Progviz before.¹⁰

Most systems have more threads active than the number of available CPU cores, so it is not possible to dedicate an entire CPU core to each thread. To solve this issue, the operating system employs a technique called *time-sharing* to allow multiple threads to share a single CPU core. To illustrate the idea, imagine that we run the program `multi-a.c` on a system that only has a single CPU core. Initially, the operating system can let the main thread of `multi-a.c` uninterrupted until it starts the second thread. At this point the operating

¹⁰Progviz does not model this form of parallel execution. As we shall see, this does not matter for its ability to detect problems. It only means that Progviz is not able to reproduce some states reachable by *incorrect* programs that might happen on a real CPU.

system has two threads but only a single CPU core. At this point, the operating system has a few options. It could let the main thread continue executing on the CPU core until it is done, and then switch to the second thread and let it resume. Another option would be to run the newly created thread first and let the main thread wait. A third option is to let one of the threads execute for a while (a few milliseconds, perhaps), then switch to another thread and let that execute for a while, and so on until the threads terminate. If we switch between threads fast enough, this third option gives the illusion that the system is able to execute more threads than what the hardware actually supports. However, threads do not actually execute *in parallel* in this model. We do, however, consider them to execute *concurrently* with each other. Note that threads that run *in parallel* are also considered to run *concurrently*. The term *concurrent execution* is thereby broader than *parallel execution*.

One interesting observation is that time-sharing concurrent execution is what Progvis simulates. One way of thinking of Progvis' visualization is therefore that *you* are the scheduler in an operating system with one CPU core. You are therefore in charge of deciding which thread should be executed at each time. This gives you the power to explore how different scheduling decisions affect the behavior of concurrent programs.

As you might imagine, many details of this task switching are up to the operating system. For example, how long should we let one thread execute before switching to the next one? If we switch too often, the overhead of switching between threads will reduce the overall system performance. On the other hand, if we switch too infrequently the illusion of concurrent execution will be broken, and it will instead look like we just run one thread after another. Furthermore, if the system contains more than two threads, the system also needs to decide *which* thread to run based on some criteria. All of this gets more complicated in systems that have more than one CPU core. Then the operating system also needs to determine which CPU cores should execute which threads so that threads get their fair share of CPU time.

All of these design decisions and the decisions made by the scheduler in response to current system load will affect the interleavings experienced by a concurrent program. For example, executing `multi-a.c` on a system with 4 CPU cores that are mostly idle will look very different from executing it on a system that has a single CPU core and many other threads that wish to execute. Since we want our concurrent programs to work correctly regardless of the number of CPU cores and system load, it is important that our programs work correctly for *all possible* interleavings. In Progvis, this corresponds to enumerating all combinations in which you can step the threads, trying all those combinations and verifying that the program behaves correctly in all of them. As you probably imagine, this is not practical to do for all but the simplest programs. Because of this, we will need to reason about the program's behavior at a higher level. Progvis also has a *model checker* that essentially automates checking all combinations that we will use later to verify that our programs are actually correct.

3.5 Problems in Concurrent Programs

Before generalizing our observations from the previous sections into a model for how concurrent programs execute, we will take a moment to closer investigate how the issues we found previously show themselves in programs running on real hardware. As such, in this section we will compile and run programs outside of Progviz, using the system's C compiler.

Before we proceed, it is worth mentioning that all of the programs used in this section exhibit *undefined behavior*. That means that essentially anything may happen when they are executed. While the programs are designed to most likely show the behavior described in this book, there is a chance that you will see slight differences depending on the hardware and software in your system.

```
int result;

void thread_fn(void) {
    for (int i = 0; i < 100000; i++) {
        result += 2;
    }
}

int main(void) {
    thread_new(&thread_fn);
    printf("result=%d\n", result);
    return 0;
}
```

Listing 3.3: The program add-1.c.

The first of the programs we will explore in this section is add-1.c as shown in Listing 3.3. The main thread first starts a new thread that executes `thread_fn` and then prints the contents of the variable `result`. The thread that executes `thread_fn` adds 2 to `result` 100 000 times, so the final contents of `result` will be 200 000 when the loop finishes.

Practice: Based on what we have observed from the previous section, predict the output of the program add-1.c. When you have done so, compile and run the program using:

```
make add-1
./add-1
```

Does the actual output match your observations? Does the program produce the same output every time you run it?

When running the program, you will most likely have seen that it does not print `result=200000` as we hoped that it would. Instead it prints some smaller number (usually between 10 000 and 20 000 on my system, but this varies greatly between different hardware). The reason for this is that the main thread does not wait for the started thread to finish executing. As soon as the main thread starts the second thread, both threads are able to execute

concurrently. This means that the scheduler may allocate them to a core at any time from that point onwards. The reason that the program produces different outputs is that the scheduler makes slightly different decisions each time the program is run.

Practice: Load the program `add-1-progvis.c` in Progvis. This program is the same as `add-1.c` with the exception that the loop only adds 2 to `result` 4 times instead of 100 000. As before, select *Run* \Rightarrow *Report data races* \Rightarrow *Disable* to tell Progvis to not complain about the issues in the program for now. Find interleavings that cause the program to print `result=x` for $x = 0, 2, 4, 6, 8$. What is the difference between these interleavings?

As you have likely seen from above, the scheduling decisions greatly affect the behavior of this program. In one extreme, the scheduler chooses to let the main thread continue to execute immediately after `thread_new` while making the other thread wait. This makes the program print `result=0`. In the other extreme, the other thread gets to execute to completion before the main thread gets a chance to continue after `thread_new`. In this case the program prints 200 000 in the case of `add-1.c` or 8 in the case of `add-1-progvis.c`. The scheduler may also choose to let the other thread run for a while before letting the main thread resume. This produces some value in between the two extremes. All of these options are legal for the scheduler to make. As such, we can not make *any* assumptions about *when* different threads execute, or even how *fast* they appear to execute in relation to each other (the other thread usually manages to execute a few thousand iterations of the loop before the main thread manages to print the result, for example).

```
int result;
bool done;

void thread_fn(void) {
    for (int i = 0; i < 100000; i++) {
        result += 2;
    }
    done = true;
}

int main(void) {
    thread_new(&thread_fn);
    while (!done)
        ;
    printf("result=%d\n", result);
    return 0;
}
```

Listing 3.4: The program `add-2.c` with an attempted solution to the problems in `add-1.c`.

We could try to solve this issue as shown in Listing 3.4 (`add-2.c`), by adding a boolean variable (`done`) that the started thread uses to tell the main thread when it is done. The variable is `false` initially. After completing the loop, `thread_fn` sets it to `true`. Finally, the main thread waits for it to become `true` using a `while`-loop.

Practice: Compile and run the program using:

```
make add-2
./add-2
```

Do you see any problems with this approach? Consider both problems related to the discussion about program equivalence from Section 3.1, and any other problems you can think of.

If you wish to explore the program using Progvis, you can use the program `add-2-progvis.c`. This program is the same as `add-2.c`, but again with a reduced number of loop iterations.

As you have likely found above, the program *appears* to work correctly. This appears to be the case even if we compile the program with optimizations (`make OPT=1 add-2`). The program does, however, exhibit undefined behavior and may do anything. This is a case where the program *happened* to work even though it is faulty.

To illustrate why, consider the program transformations that the compiler and the hardware are able to do based on the notion of equivalence covered in Section 3.1. One valid transformation of `thread_fn` is as follows:

```
void thread_fn(void) {
    done = true;
    for (int i = 0; i < 100000; i++) {
        result += 2;
    }
}
```

This is valid since it does not affect the behavior of `thread_fn` according to what is observable through system calls. Both the version above and the original version sets `done` to `true` and `result` to 200 000. However, we can quite easily see that setting `done` to `true` before finishing the computation of `result` would make the program behave just as poorly as `add-1.c`.

The reason we don't see this behavior is that the compiler performs an additional optimization when we have turned on optimizations. Since `thread_fn` does is simple enough, the compiler can see that `result` will always be 200 000 at the end. Therefore, it transforms the function to the code below:

```
void thread_fn(void) {
    result = 200000;
    done = true;
}
```

The order of the lines above is arbitrary, but do not matter in this case since `thread_fn` usually finishes before the main thread gets a chance to continue execution. This is why it appears as though our changes work in spite of

enabling compiler optimizations. However, this optimization actually hides another more serious problem in the `main` function.

If we add a call to the library function `prevent_optimization` inside the loop like below, we get the code in the program `add-3.c`. The call to `prevent_optimization` only prevents the compiler from performing the two optimizations above. It has no other side effects.

```
void thread_fn(void) {
    for (int i = 0; i < 100000; i++) {
        result += 2;
        prevent_optimization();
    }
    done = true;
}
```

Now that the compiler is unable to optimize `thread_fn` as much, the issue in `main` becomes visible. We can see the issue by compiling and running `add-3.c` as follows:

```
make add-3
./add-3
```

The program appears to never finish. If you look at the CPU usage in your system (e.g., using `htop`, press `q` to quit) you will find that the program is not idle, but is using all CPU time it can get. You can press `Ctrl+C` in the terminal to quit it.

What happens is that the compiler has transformed the `while (!done)` loop that the main thread uses to wait for the other thread into one of the two versions below.¹¹

<pre>bool tmp = !done; while (tmp) ;</pre>	<pre>if (!done) { while (true) ; }</pre>
---	--

As we can see, the compiler has helpfully avoided reading the global variable `done` in our loop to avoid the cost of accessing memory. The left version reads `done` once and places the negation of it in the local variable `tmp`.¹² The right version also reads `done` once, but this time by extracting the loop condition into a separate `if`-statement before the loop.

Both of these transformations are valid for sequential programs, as they preserve the behavior of the original program. If `done` is `true` then the program continues. If `done` is `false` the program enters an infinite loop. Again, this equivalence only holds under the assumption that the program is sequential. Since we use the original loop to wait for another thread to change the variable, these transformations break our program since they only check `done` *once*. What happens when we run `add-3.c` is therefore that the main thread checks `done` *once*, and when it realizes that the other thread is not yet done, it enters an infinite loop and never checks `done` again. This is why the program never

¹¹If you are curious, you can use a debugger to inspect which version your compiled produced. With `gdb`, you can type `disas main` to see the assembler version of the code.

¹²Which is then placed in a register to make it cheaper to access.

terminates. This also illustrates a problem with using a `while` loop to wait for another thread to complete, doing so wastes CPU-time doing nothing.

```

int result;
bool done;

void thread_fn(void) {
    for (int i = 0; i < 100000; i++) {
        result += 2;
    }
    done = true;
}

int main(void) {
    thread_new(&thread_fn);
    for (int i = 0; i < 100000; i++) {
        result += 5;
    }
    while (!done)
        ;
    printf("result=%d\n", result);
    return 0;
}

```

Listing 3.5: The program `add-4.c`.

All of the problems above are caused by *shared data* between the two threads. So far, we have only seen what happens when one thread (the main thread in this case) reads from shared data while another thread writes/modifies it. As illustrated by the program `add-4.c` in Listing 3.5, problems also arise when two threads attempt to modify the same data. The program `add-4.c` further expands `add-2.c`, which worked correctly without optimizations. This program adds a loop in the `main` function that adds 5 to `result` 100 000 times. They are otherwise identical.

Practice: Predict the output of the program `add-4.c`. Then compile and run the program using the commands below (i.e., *without* optimizations):

```

make add-4
./add-4

```

Does the actual output match your prediction? Does the program output the same result every time?

Even though we would have expected the program to output $200\,000 + 500\,000 = 700\,000$, we typically get a number that is smaller than that, even though we know that both loops run to completion before the result is printed.

The reason for this behavior is that the addition of 2 and 5 to `result` are not performed as a single operation. Even though we use the `+=` operator, the CPU eventually needs to perform three distinct operations: 1) read the value of `result`, 2) increment the value by 2 or 5, and 3) write the value back to `result`. In C, this would look like the code below:

```

void thread_fn(void) {
    for (int i = 0; i < 100000; i++) {
        int tmp = result;
        tmp += 2;
        result = tmp;
    }
    done = true;
}

```

As such, the reason why `add-4.c` does not always output 700 000 is that some increments of `result` get lost. The `+=` operators have been expanded as above in the program `add-5-progvis.c`, so that you can explore how it affects the program.

Practice: Load `add-5-progvis.c` in Progvis. As before, select *Run* \Rightarrow *Report data races* \Rightarrow *Disable* to tell Progvis to not complain about the issues in the program for now. Find an interleaving where the program prints a result that is lower than $(2 + 5) \cdot 4 = 28$.

From the task above, you will notice that the introduction of the temporary variable means that changes to `result` are not always immediately visible to the computations done by other threads. This means that threads may accidentally overwrite the results of each others' computations. As we saw previously, this happens even if we use constructs like `result += 2` that don't *look* like they introduce temporary variables. Since the hardware eventually needs to perform this kind of operation as three separate steps, there will always be some form of temporary involved, even if it is not visible in the program's source code, or even in the machine code (e.g., if an *increment* instruction is available, which is the case on x86). Note, however, that this behavior is not visible in Progvis, since it treats entire statements as a unit. As we will see soon Progvis will still notice the problem. It will just not give an example of *what* may happen.

To illustrate exactly what happens, we can load `add-5-progvis.c` in Progvis and step both threads until they have both read the value of `result` into `tmp`. In this case `result` is 0 since it is the first iteration of both loops. This is shown in Fig. 3.9. In this figure, we can also see how Progvis visualizes local variables. Notice that the loop variable `i` is displayed inside each thread's box, under the name of the function. This representation helps remind that local variables and parameters in functions are *not accessible* to other threads by default.

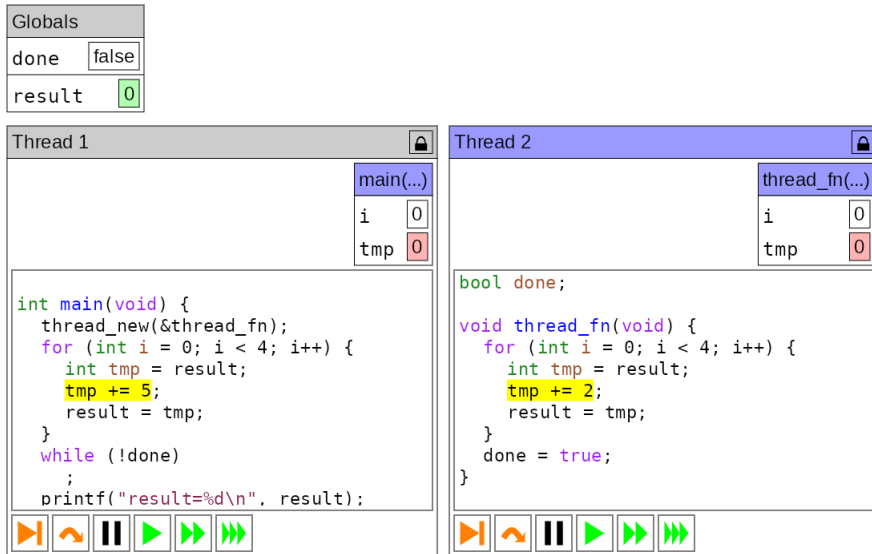


Figure 3.9: Progvis visualizing add-5-progvis.c.

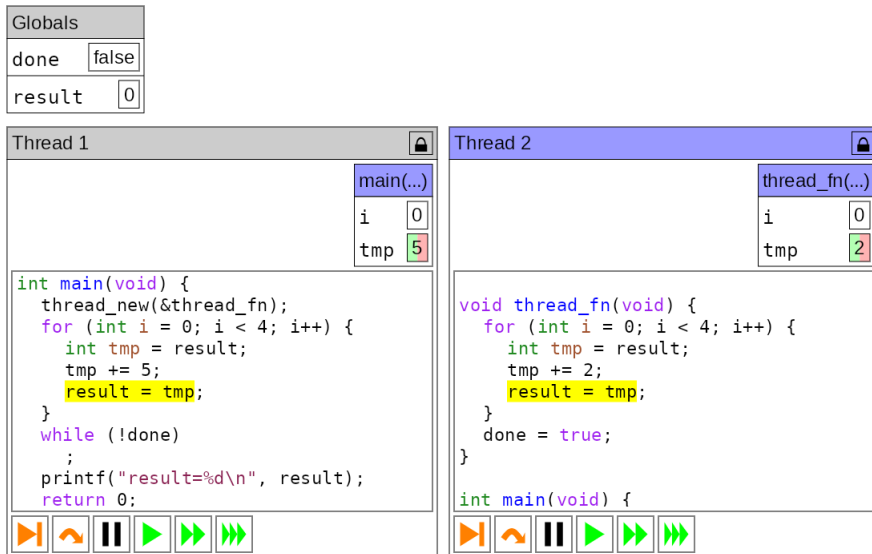


Figure 3.10: After advancing both threads one step.

Figure 3.10 shows the state of the program after we have advanced both threads to the next step. Here, we can see that both threads have updated their `tmp` variable to 5 and 2 respectively. Already here, we can imagine that the result of both threads writing their `tmp` back to `result` will not produce the desired results.

We can see that this is indeed true by advancing thread 1 to get the state in Fig. 3.11. Here, thread 1 has written its `tmp` variable to `result`, and `result` thereby contains 5. If we now let thread 2 advance one step, we will get the

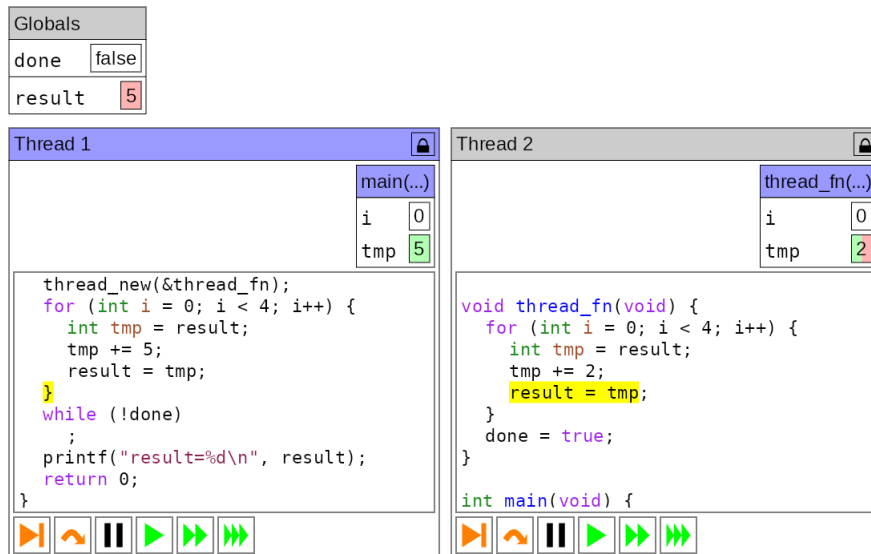


Figure 3.11: After advancing thread 1.

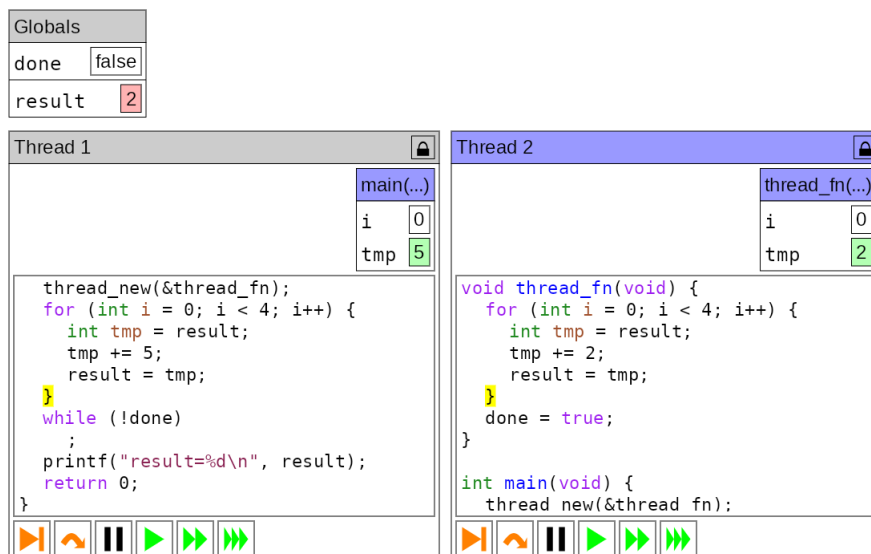


Figure 3.12: After advancing thread 2.

output in Fig. 3.12. We can see that `result` now contains 2 since thread 2 stored the value from its `tmp` variable into `result`, thereby overwriting the 5 that thread 1 stored there previously. This will eventually lead to `result` being too low, since we have “lost” the addition of 5 by thread 1.

At this point, we can notice that the above is *one* example of what can go wrong. We will “lose” more than one addition of 5 if we let thread 1 run for more iterations before stepping thread 2. Similarly, if we would have let thread 2 run first, we would have lost thread 2’s work instead. What happens

in practice when we run the program outside of Progvis is some combination of all of these incorrect situations.

3.5.1 Summary

In summary, the examples from above show that *shared data* between multiple threads need special attention in concurrent programs. This is for two main reasons:

- Sometimes we need to *wait* for another thread to complete their work and write their result to some shared variable. This is the problem we saw in `add-1.c` and attempted to solved in `add-2.c`.
- We also need to be careful when multiple threads *access* the same data at the same time. Since the compiler and the hardware assumes that programs are sequential by default, access to shared data does not behave as we would expect which leads to the problems we saw in `add-2.c`, `add-3.c`, and `add-4.c`.

Even though Progvis does not model exactly how programs may misbehave if shared data is accessed concurrently, it *is* able to detect *when* problems would happen and inform you about them. Up until now, we have explicitly disabled these messages since the programs we have used so far all have major problems. If you select *Run* \Rightarrow *Report data races* \Rightarrow *Full*, you re-enable these messages. For example, if you run `add-5-progvis.c` with this setting, Progvis will notice that the two threads access the same variable concurrently and report it. Furthermore, it will draw a red cross over the problematic variable as shown in Fig. 3.13.

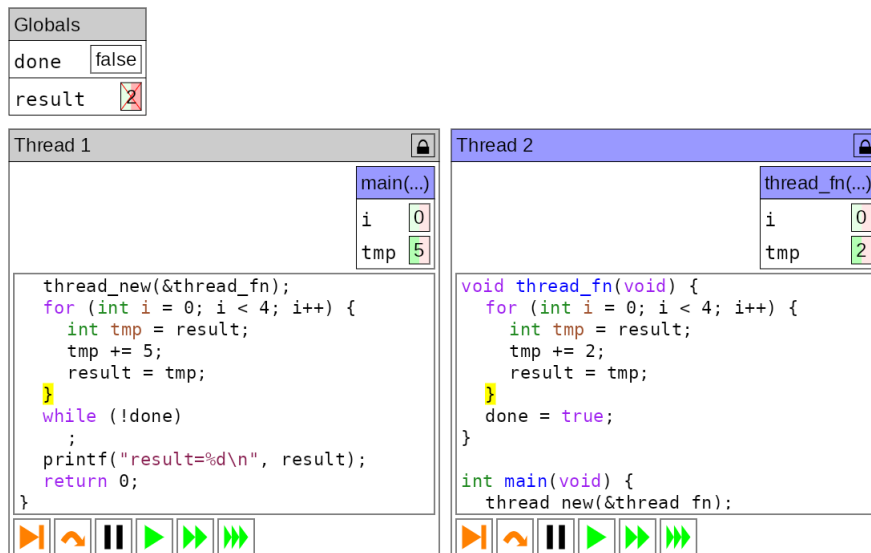


Figure 3.13: How Progvis illustrates which variables are accessed concurrently. Notice the red cross over the value of the global variable `result`.

3.6 The Concurrent Execution Model

Now that we have seen a number of examples of problems that arise when we introduce multiple threads to sequential programs, we are ready to summarize our findings into an *execution model*. Just as our sequential execution model, the goal is to provide a *model* that we can use to reason about the behavior of concurrent programs. We want this model to be *simple* but *accurate* so that we can understand it. It is worth noting that the model does *not* specify what happens at the compiler- or hardware level. Just as the model for sequential execution, the compiler and hardware may do whatever is efficient, as long as programs appear *as if* they are executed according to the model.

The model described here is based on the memory model used by C and C++. The memory models used by other popular languages (e.g., Java, Rust, etc.) are largely similar, but there are small differences. Luckily, the memory model used by C and C++ is usually stricter, so programs that are correct according to the model presented here are almost always correct according to the memory models in other languages as well.

The concurrent execution model can be summarized as follows. Each of the points will be discussed in more detail afterwards.

1. Programs that are free from *data races* behave *as if* they are executed in program order.
2. Multiple threads execute *concurrently* with each other. The relative execution speed and timing of multiple threads are *arbitrary* and *often vary* between different program executions. The only thing we can assume is that each thread will *eventually* execute *some* code.
3. A *data race* is a situation where one thread *writes* to data that another thread *accesses* (i.e., reads and/or writes) concurrently.
4. The behavior of a program where a *data race* is possible is *undefined*. If a data race is possible, we say that a program *contains a data race*.

The points above are intentionally short and to the point to make it easier to refer back to them later. This does, however, make them a bit hard to understand at first. Below is a more elaborate description that hopefully clarifies more of the nuances of each point:

1. This point is similar to the sequential execution model: programs are executed as if each line in the source code is executed line by line. The only difference is the point about *data races*. As we shall see, data races can not happen without multiple threads. This means that the concurrent execution model is the same as the sequential execution model for sequential programs, and as long as we make sure we don't have data races we can reason in the same way for concurrent programs.
2. This point covers the fact that concurrent execution is implemented differently on different systems (e.g., time-sharing vs. hardware parallelism), and that concurrent execution is often non-deterministic (i.e., two consecutive runs of the same program may produce different results). This means that the *only thing* we can assume is that threads will eventually

be able to execute some portion of their code. We can not make assumptions of how often this will happen, or how quickly one code executes code in relation to another thread. Therefore, if it is important that some piece of code executes before some other piece, we need to make sure that they are executed in the right order.

3. This point defines a *data race*. Intuitively, a data race occurs when two threads access the same data and at least one of them writes to the data. In the examples above, the shared data is in the form of global variables, but data may be shared via other means as well. From the examples above, we saw that all problems with concurrent executions were in some way related to shared data, and one thread modifying the data at an inconvenient time.
4. The point goes on to state that all programs where data races are *possible* are undefined. As in point 1, this essentially says that “if a data race is possible, anything may happen when the program is executed”. This means that if we want to be able to reason about the behavior of our program, it may not contain any data races. In combination with point 2, that we don’t know when threads execute in relation to each other, this means that data races should not be possible for *any* legal interleaving of the program. This means that we as programmers need to make sure that data races are not possible by *protecting* shared data in a suitable manner.

3.6.1 Implications of the Model

Now that we hopefully understand what the model *is*, it is also useful to understand *why* the model looks the way it does.

First and foremost, the model is similar to that of sequential program execution. This is nice, since the line-by-line execution model is simple to reason about, and most programmers are already familiar with it. Since sequential programs cannot contain data races, the concurrent execution model is the same as the sequential execution model for sequential programs. This is nice, since it lets us use *one* execution model regardless, and we can instead speak about *the* execution model.

As we saw earlier in this chapter, (almost) all problems in concurrent programs are related to shared data. Either because one thread needed to wait for another thread to produce data, or because both threads attempted to update data at the same time. Both of these situations exposed the non-deterministic nature of thread execution, or exposed compiler optimizations that re-ordered memory accesses. The execution model gets around these problems by asking the programmer to avoid data races. What this essentially means is that the programmer need to *synchronize* access to shared data, which has the effect of marking these accesses in the program. If the compiler can assume that all accesses to shared data are marked correctly, then it is able to optimize concurrent programs just as well as sequential programs, as long as it is “careful” around the marked accesses.

The big downside of just leaving the behavior of programs with data races undefined is, of course, that it is difficult to detect if a program contains data races or not. Some data races are rare in practice, so they are unlikely to

be noticed by just running the program and observing its output. It is also possible that the choices your compiler and hardware makes happen to match with what you intended the program to do. That way the program happens to work fine in spite of it containing undefined behavior. However, this may no longer be true if you upgrade your hardware and/or your compiler, at which point the program starts to fail and you don't know where to start looking. The takeaway from these observations is that it is nearly impossible to check whether or not a program is correct by just running it. It is necessary to carefully reason about the places in the program that deal with shared data!

3.6.2 The Execution Model in Progvis

Now that we know what the model says, we can also examine what Progvis visualizes in more detail. Progvis visualizes the program *as if* it is executed sequentially. According to point 1 in our model, this is in line with point 1 in the model. To help with identifying problems, Progvis also keeps track of all memory accesses made by the program. If it ever detects that a thread wrote to a variable that was accessed by another thread it knows that a data race has occurred, and reports it as such. Since data races are undefined by the model, anything may happen. As such, there is no single “accurate” way to show what might have happen in those cases. This is why we could not see some problems in Progvis — it selects *one* possible behavior for undefined programs, but some other behavior may happen in practice. However, for programs that are free from data races, Progvis matches the execution model.

It is worth noting that the observations Progvis use to detect data races are based on the interleavings selected by the user. As such, to be entirely certain that a program is free from data races, you would have to try *all possible* interleavings. As you can imagine, this is quite tedious. Therefore Progvis also has a *model checker* that tests them for you. You can start it by selecting *Run* \Rightarrow *Look for errors...* If Progvis finds any interleavings that lead to a data race, it will show them as an animation. Otherwise it will report that it was unable to find any as the program is used currently.

If you have used the mode *Run* \Rightarrow *Report data races* \Rightarrow *Full* you might have noticed that Progvis sometimes colors variables in a faded-out nuance of red and green. Just like the regular, stronger, colors this represents variables that have been written to and read from. The faded-out color represents that the variables were not accessed by the last statement, but further back in the program. In particular, they were accessed since the last *barrier* (e.g., starting a new thread, protecting a variable, we will define it in more detail later). It turns out that the memory model allows us to treat blocks between such barriers as a single unit. The second button in each thread (labeled B in Fig. 1.1) jumps to the next such synchronization event.

3.7 Exercises

The code used in the exercises can be compiled either with your system's C compiler, or visualized in Progvis. If you are using Progvis, you can turn off messages about data races using *Run* \Rightarrow *Report data races* \Rightarrow *None* and turn them back on using *Run* \Rightarrow *Report data races* \Rightarrow *Full*.

1. The file `01-equivalence.c` contains four versions of the function `update`. The function `update1` is the original version of the function, and `update2`, `update3`, and `update4` are modifications to the original. Which of the modifications are equivalent to the original according to the execution model discussed in this chapter?
2. The file `02-concurrent-equivalence.c` contains the same four versions of the `update` function as in the previous exercise. This time, the main function starts another thread before calling one version.
 - a) If you run the program in Progvis (turn off messages about data races) and use the original version (`update1`) you can get two different final values of `x`, which ones?
 - b) If you instead let the program run other versions in each thread, what final values can `x` take then? Only use the versions of `update` that you found to be equivalent to the original one in the previous exercise.
3. The files `03-global-a.c` and `03-global-b.c` contain two similar programs. What variables are shared between threads in each of the programs? If there is any shared data, does any of the programs contain a data race? You can verify your solution using Progvis with messages about data races turned on.
4. The file `04-ptr.c` contains a program that launches a thread and passes a pointer to it. Does the program contain any shared data? If so, what data is shared? Does the program contain a data race? You can verify your solution using Progvis with messages about data races turned on.
5. The file `05-local.c` contains a program that computes 2^x in two threads. Does the program contain any shared data? If so, what data is shared? Does the program contain a data race? You can verify your solution using Progvis with messages about data races turned on.

Semaphores

In the previous chapter we saw the problems that arise when we allow more than one thread to execute concurrently in the same process. Even though we found different types of problems, all of them originate from careless use of *shared data*. We also realized that we did not yet have the necessary tools to solve the problems we found. We tried to solve one of the problems that we found using a boolean variable and a while loop, and while the approach initially *appeared* to work, we saw that it stopped working as soon as we asked the compiler to optimize our code.¹ The problems we found can be grouped into two broad groups based on the type of solution that would be needed to solve the problem:

- One thread needs to *wait for* something to happen.
- We need to *avoid data races*, by making sure that two threads do not access the same data concurrently (i.e., ensure *mutual exclusion*).

In this chapter we will focus on solving the first problem (*waiting for* something to happen). We will, however, see that the second problem is just a variant of the first problem. It is, however, often useful to think of them as two separate problems.

4.1 Semantics

Since we are not able to solve this type of synchronization issues using clever programming techniques, we need some other tool. In particular, we need to use a *synchronization primitive*. These are implemented using specific low-level primitives that inform both the compiler and the hardware that they protect some shared data. Therefore, both the compiler and the hardware know that they need to be careful whenever some synchronization primitive is used.

A *semaphore* is one of a few synchronization primitives that is available in many programming languages. A semaphore can be considered to contain a counter variable. The counter variable is an integer that is not allowed to

¹On some CPU architectures, the solution may not even work without turning on optimizations.

become negative. The semaphore is an opaque data structure, so the counter can not be manipulated directly. It is necessary to use the three functions below to manipulate the counter.² The presentation below focuses on how the functions are *used*. The definitions are available in Chapter A for the curious reader.

```
struct semaphore sema;
```

The line above defines a variable named `sema` that contains a semaphore. As with other variables in C, semaphores can be declared either at global scope, inside data structures, or as local variables inside functions.

```
sema_init(&sema, 0);
```

Each semaphore variable needs to be initialized exactly once before it can be used. This is done using the `sema_init` function as shown above. The `sema_init` function takes two parameters. The first parameter is a pointer to the semaphore that should be initialized. As such, we pass `&sema` to get a pointer to the semaphore variable. The second parameter is an integer that is used to initialize the counter inside the semaphore. As such, the second parameter cannot be negative. In this case, we initialize the counter inside the semaphore `sema` to 0.

When the semaphore is initialized, it is safe to call the functions `sema_down` and `sema_up`. Furthermore, it is safe to call these functions from different threads concurrently, which means that we can use them to solve the concurrency problems we found previously!

```
sema_up(&sema);
```

The function `sema_up` accepts a single parameter, a pointer to the semaphore that should be updated. As the name implies, the function increases the value of the counter inside the semaphore by 1.

```
sema_down(&sema);
```

The function `sema_down` also accepts a pointer to the semaphore to update as its single parameter. Again, as the name implies `sema_down` attempts to decrease the counter inside the semaphore by 1. However, remember that the counter is not allowed to become negative. As such, if a thread, *T*, calls `sema_down` to decrease a semaphore with a counter that is already at 0, `sema_down` will let *T* sleep until another thread increases the counter above zero (by calling `sema_up`). When *T* wakes up, it decreases the counter and continues execution.

We can make a few important observations from the descriptions above. First and foremost, we can see that `sema_up` will always increase the counter inside the semaphore by 1, and that `sema_down` will always decrease it by 1 (potentially after sleeping for some time). This makes it possible to think of the counter as if it counting the availability of some resource in the program. We will see how this can help us to synchronize programs in a bit.

Another important observation is that only `sema_down` may cause a thread to sleep. For this reason, the function `sema_down` is sometimes called *wait*, and `sema_up` is sometimes called *signal*.

²You can think of these functions as member functions in a class named `semaphore`. This is how they are often implemented in object oriented languages.

4.2 Semantics in Sequential Programs

Now that we know what operations are available and how they work, it is time to examine closer how they behave in actual programs. To keep things simple initially, we start with the sequential program `simple-semantics.c`.

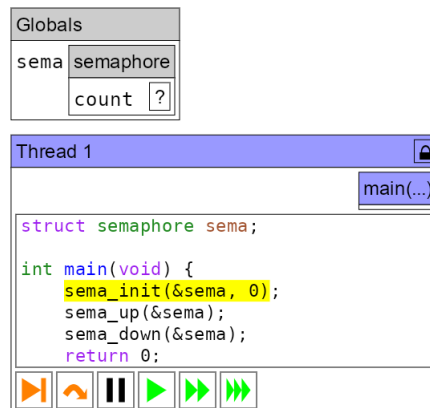


Figure 4.1: Progvis’ visualization of a semaphore at the start of the program `simple-semantics.c`.

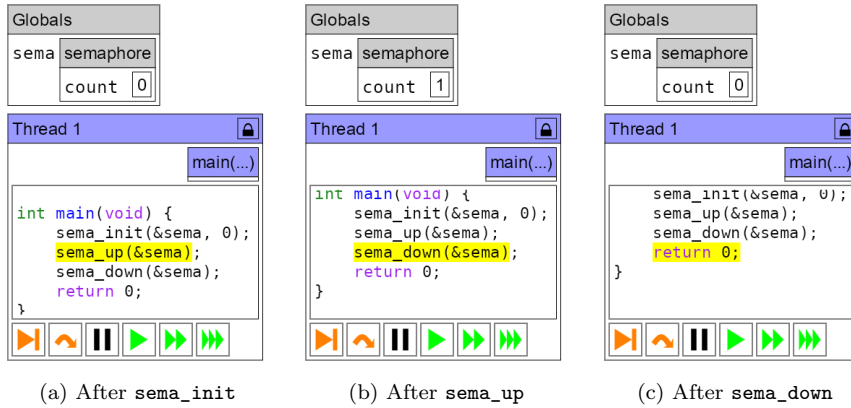
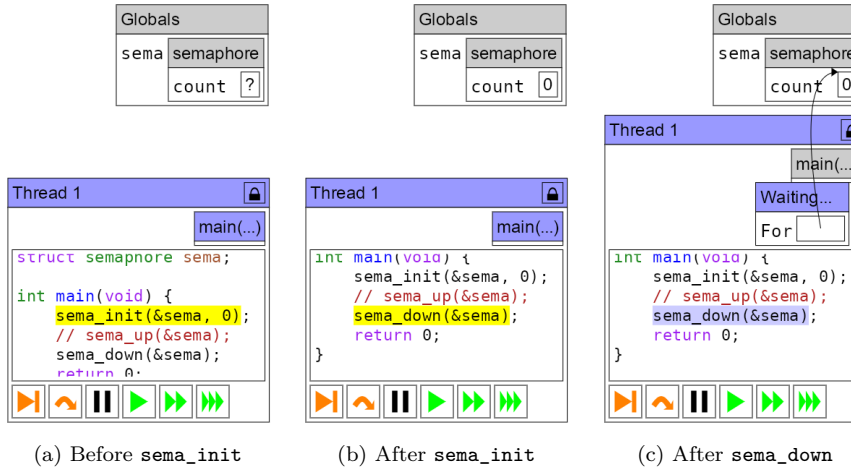
If you open `simple-semantics.c` in Progvis, you will see the representation in Fig. 4.1. Notice how Progvis visualizes the `sema` variable inside the *Globals* box. Since a semaphore is a complex data structure, it is shown as its own box. The title of the box is the name of the type (i.e., `semaphore` in this case). The type is shown to contain a single variable named `count`. This is how Progvis displays the current value of the counter inside the semaphore. Since the semaphore has not yet been initialized, the value of the counter is undefined and therefore displayed as a question mark. Notice that even though Progvis shows `count` as a member of the `semaphore` type, it is not possible to access it from the program. As mentioned earlier, it is only possible to manipulate the semaphore by using the three functions `sema_init`, `sema_up`, and `sema_down`, and none of them allows reading the value of the counter.³

If we continue to step thread 1 we will see the steps in Fig. 4.2. The first step (Fig. 4.2a) shows the state after running `sema_init`. Here, the counter in the semaphore is initialized to 0, since we passed 0 as the second parameter to `sema_init`. The second step (Fig. 4.2b) shows the semaphore after running `sema_up`. Since `sema_up` increases the counter by 1, the counter now contains 1. Finally (Fig. 4.2c), after running `sema_down` the counter is decreased to 0 again.

The program `simple-semantics.c` has not shown what makes semaphores special yet. The behavior we have seen so far could easily have been replaced by incrementing and decrementing an integer variable. To see what makes a semaphore special, remove the call to `sema_down` in `simple-semantics.c` (e.g., by commenting it out).

If we run the modified program and step thread 1, we will see the steps in Fig. 4.3. As we can see in Fig. 4.3b, the counter in the semaphore will be 0

³The reason for this is that manually inspecting the value of the counter rarely corresponds to correct ways of using the semaphore.

Figure 4.2: Progviz executing the remaining steps of `simple-semantics.c`.Figure 4.3: Progviz executing the remaining steps of `simple-semantics.c`.

when `sema_down` is called since we removed the call to `sema_up`. As mentioned before, the counter inside the semaphore is not allowed to become negative. To avoid decreasing the counter to a negative number, the semaphore puts the thread to sleep to wait for some other thread to increase the counter. In Fig. 4.3c we can see how Progviz shows a sleeping thread. First, another box is shown on top of the call to the `main` function with the label *Waiting...*, and that the current line is highlighted in purple instead of yellow. This means that the thread is waiting for something, rather than being ready to run. This means that the scheduler may *not* schedule the thread to run at this time. If you try to step the thread in this state, you will see that nothing happens. The *Waiting...* box also contains a line *For* that shows what the thread is waiting for. In this case, the thread is waiting for the counter in the semaphore to become positive, which is illustrated by an arrow that points to the counter in the semaphore.

In this case there are no other threads in the program. Therefore, no other

thread is able to increase the counter in the semaphore. This means that thread 1 will wait forever. Progviz detects this situation and reports it as a *deadlock*. We will examine deadlocks closer in a future chapter. For now, it is enough to know that if all threads in a program are waiting at the same time, Progviz will report it as a deadlock.

4.3 Semantics in Concurrent Programs

As we saw in the last example, using a semaphore in a sequential program is not very useful. Since semaphores are used to coordinate multiple threads, we also need to consider how semaphores affect the behavior of multiple threads. To illustrate this, we will examine the program `semantics.c` in more detail.

```
struct semaphore sema;

void thread_fn(void) {
    sema_up(&sema);
}

int main(void) {
    sema_init(&sema, 0);
    thread_new(&thread_fn);
    sema_down(&sema);
    return 0;
}
```

Listing 4.1: Source code of the program `semantics.c`.

As we can see in Listing 4.1, the program `semantics.c` is similar to the sequential program in the previous section. It defines a semaphore (`sema`) as a global variable and initializes it to zero in `main`. After initializing the semaphore, it starts a new thread that runs `thread_fn` and then calls `sema_down`. As such, if we just consider the contents of `main`, the program behaves exactly like `simple-semantics.c` when we removed the call to `sema_up`. The difference is that `semantics.c` started another thread that *will* call `sema_up` and thereby wake up the main thread.

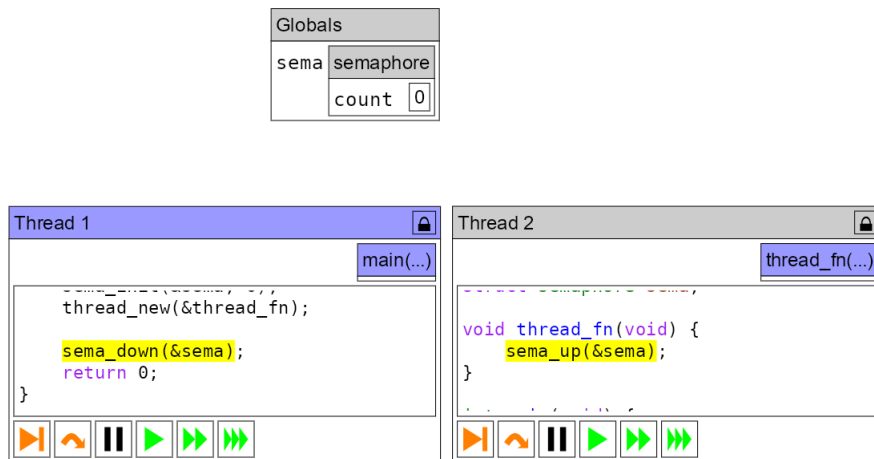


Figure 4.4: The program `semantics.c` just before thread 1 calls `sema_down`.

To see how `semantics.c` behaves, we examine it in Progvis. If we step thread 1 until just before the call to `sema_down` the program will be in the state depicted in Fig. 4.4. Just like in the previous example the semaphore's counter is at zero. This means that just as the previous example, the call to `sema_down` will cause thread 1 to wait.

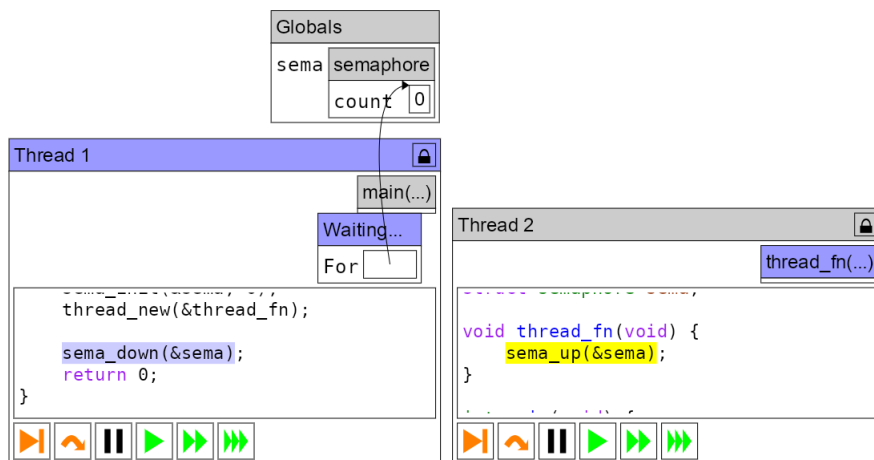


Figure 4.5: The program `semantics.c` when thread 1 is waiting for the semaphore.

Indeed, if we let thread 1 continue its execution we will reach the state in Fig. 4.5. Just as in the previous example the counter inside `sema` was zero, so the semaphore causes the thread to wait until the counter is increased. Since thread 1 is waiting, clicking its *step* button in Progvis does nothing. A similar thing happens if we run `semantics.c` outside of Progvis. When a thread is waiting it is not eligible to be scheduled by the system's scheduler, and it may therefore not continue executing. This is why the *step* button has no effect.

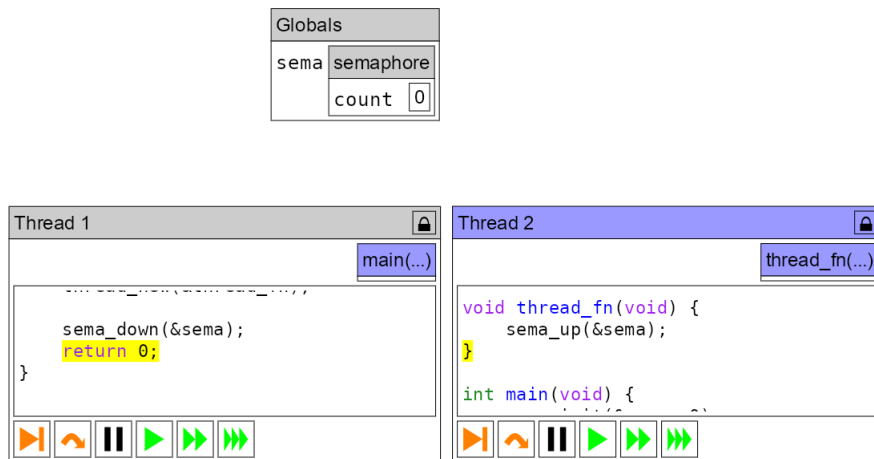


Figure 4.6: The program `semantics.c` after thread 2 has called `sema_up`.

Since thread 1 is waiting, our only option is to step thread 2. Luckily, thread 2 is about to call `sema_up`. If we step thread 2, then two things will happen after each other. Progvis do, however, consider them to be a single step, so they appear to happen at the same time:⁴

1. Thread 2 calls `sema_up` and increases the counter inside `sema` to 1.
2. Since the counter is now positive, thread 1 wakes up and continues executing `sema_down`.⁵ Since the counter is 1, `sema_down` decreases the counter as usual, bringing it back to 0.

After this has happened we get the state in Fig. 4.6. In particular, we note that thread 1 is no longer waiting and that it has returned from `sema_down`. We can also see that the counter in the semaphore is still at zero, even though thread 2 called `sema_up`. As mentioned above, this is because thread 1 executed `sema_down` immediately after thread 2 called `sema_up`.

Practice: The example above showed what happens when thread 1 calls `sema_down` before thread 2 calls `sema_up`. Use Progvis to investigate what happens when thread 1 calls `sema_down` *after* thread 2 calls `sema_up`. What is the final value of the semaphore in this situation?

By examining what happens when thread 1 calls `sema_down` *after* thread 2 calls `sema_up` you should have observed that the counter in `sema` ended up at 0 regardless of the order. This is an important observation, since it means that we don't have to worry about the order in which `sema_down` and `sema_up`

⁴Depending on how semaphores are implemented, they may actually be one step. The mental model presented here is, however, still accurate and makes it easier to think about the behavior in the future.

⁵Technically, thread 1 cannot “notice” anything since it is sleeping. Thread 2 notices that threads are waiting for the semaphore when it executes `sema_up` and wakes one of them up.

was called. If we think about it, this is not too surprising. We initialized the counter to 0, and we called `sema_up` once and `sema_down` once. Since the sum of the changes to the counter is zero regardless of the order, it is natural that the final value of the counter is the same as the value we initialized the counter to. Furthermore, this is true regardless of what we initialized the counter to.

Another important observation here is that since we initialized the semaphore to 0, we know that when `sema_down` returns, some other thread must have called `sema_up` first. The reason we know this is that the counter inside the semaphore is not allowed to become negative. As such, the only way for `sema_down` to successfully decrement the counter is if thread 2 has incremented the counter first. As such, regardless of which interleaving was chosen by the scheduler, we know that the code before `sema_up` in thread 2 (marked as A in Listing 4.2) has to finish executing *before* the code after `sema_down` in thread 1 (marked as B in Listing 4.2). This is how we can use semaphores to wait for other threads to complete their work.

```

void thread_fn(void) {
    // A
    sema_up(&sema);
}

int main(void) {
    sema_init(&sema, 0);
    thread_new(&thread_fn);
    sema_down(&sema);
    // B
    return 0;
}

```

Listing 4.2: Illustration of how semaphores can be used to ensure that A completes before execution of B begins.

Practice: Using the semaphore to wait for thread 2 to call `sema_up` only works if the semaphore is initialized to 0. What happens if you change the line `sema_init(&sema, 0)` to `sema_init(&sema, 1)`? Why does the program behave differently?

4.4 Waiting for Completion

Now that we know how semaphores work, it is time to use them to solve one of the problems from Section 3.5. To make it easier to visualize the program flow in Progvis, we will focus use the version of the problem adapted to Progvis in this chapter (i.e., the ones called `*-progvis.c` in Chapter 3). As such, in this chapter the files called `add-N.c` contain the version adapted for Progvis. Of course, the solution presented here works equally well for both versions of the problem since the only difference is that the number of iterations in the loop have been reduced from 100 000 to 4.

The file `add-1.c` (Listing 4.3) contains our attempted solution to the problem from Section 3.5. Remember that we realized that the code in `thread_fn` needs to run to completion before the main thread calls `printf`. We attempted

```

int result;
bool done;

void thread_fn(void) {
    for (int i = 0; i < 4; i++) {
        result += 2;
    }
    done = true;
}

int main(void) {
    done = false;
    thread_new(&thread_fn);
    while (!done)
        ;
    printf("result=%d\n", result);
    return 0;
}

```

Listing 4.3: The attempted solution from Section 3.5.

to solve this by adding the variable `done`, and wait for it to become `true` by using a `while` loop. However, since `done`⁶ is accessed from both threads concurrently and thereby constitutes a data race. Since data races are undefined, anything may happen when we run the program. As we saw earlier, the program appeared to work fine initially, but started to misbehave once we turned on optimizations.

Practice: Open `add-1.c` in Progviz. Make sure that *Run* \Rightarrow *Report data races* is set to *Full*. Find an interleaving that causes Progviz to report a data race that involves the `done` variable.

To solve this issue we need to use some synchronization primitive to avoid data races. Since we have only introduced semaphores so far, we will use a semaphore to solve the problem. Luckily, the problem here has a similar to the structure we saw in Listing 4.2. Since our goal is to the code in `thread_fn` complete before we run the last part of `main`, we can do the following:

1. Remove `done` and replace it with a semaphore that we call `has_result`.
2. Instead of initializing `done` to `false` at the start of `main`, we initialize the semaphore to zero.
3. Instead of the `while`-loop in `main`, we call `sema_down`.
4. Instead of setting `done` to `true` at the end of `thread_fn`, we call `sema_up`.

These changes result in the code in `add-2.c` (Listing 4.4). The variable `done` no longer causes a data race since it was replaced with the semaphore `has_`

⁶Technically also `result`, but as we will see we don't have to worry about it if we solve the problem with `done`.

```
int result;
struct semaphore has_result;

void thread_fn(void) {
    for (int i = 0; i < 4; i++) {
        result += 2;
    }
    sema_up(&has_result);
}

int main(void) {
    sema_init(&has_result, 0);
    thread_new(&thread_fn);
    sema_down(&has_result);
    printf("result=%d\n", result);
    return 0;
}
```

Listing 4.4: A correct solution using a semaphore.

`result`. Sharing a semaphore (or other synchronization primitives) in this way is safe, as long as the semaphore is initialized *before* it is shared. We initialize the semaphore in the main function, before the second thread is started. This makes it impossible for the second thread to access the semaphore before it is initialized. After that, we only access the semaphore using `sema_down` and `sema_up`, which are designed to be safe to use from multiple threads concurrently.

What is perhaps less obvious is why the global variable `result` does not lead to a data race even though it is used by both threads. Remember that for a data race to occur, threads have to access shared data concurrently, and at least one access has to be a write. Here, `result` is shared, both threads access the variable, and one of them write to the variable (`thread_fn`). However, `thread_fn` and `main` do *not* access `result` concurrently. The `main` function only accesses `result` after calling `sema_down`. Since `sema_down` waits until the second thread calls `sema_up`, we know that the second thread has finished its work that involves `result` when `main` is allowed to continue.

4.4.1 Verification in Progvis

We can verify that the program is correct using Progvis. To be sure that the program is correct, we would need to try all possible interleavings of the two threads. In this case there are not very many interleavings, but it is easy to miss one or more. To help out in situations like this, Progvis contains a *model checker* that automates this. Intuitively, the model checker simply tries all possible interleavings and verifies that none of them cause an error.⁷ To run the model checker, simply open the program in Progvis and select *Run* \Rightarrow *Look for errors*. For `add-2.c`, Progvis will quickly conclude that the program is free from concurrency errors.⁸

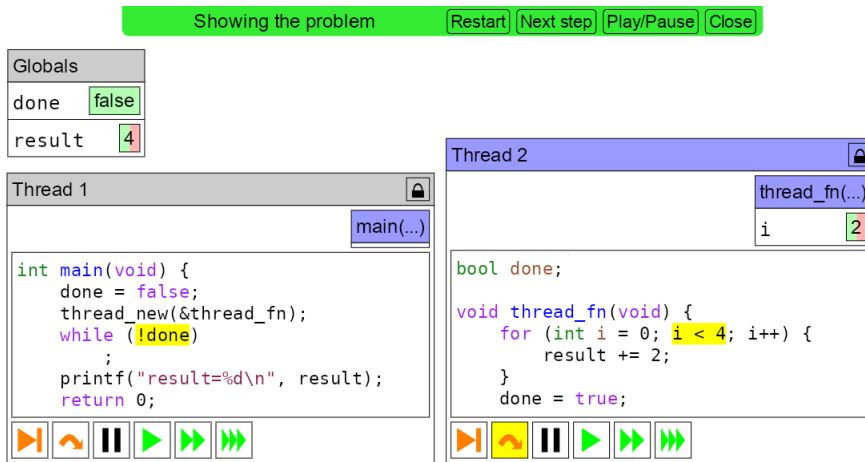


Figure 4.7: Progvis showing an interleaving that leads to a data race.

If the loaded program contains concurrency errors, Progvis will provide an example of an interleaving that causes the issue it found. To see how this works, load the program `add-1.c`, which we know contains an error, and select *Run* \Rightarrow *Look for errors* once again. Progvis will quickly find a *data race* and notify you about this. When you click the *Show* button you will now see a green bar with the title *Showing the problem* at the top of the Progvis window, as shown in Fig. 4.7. This green bar indicates that Progvis is now in control of stepping the program. In this state you are therefore not able to step threads manually (clicking the buttons will have no effect). Instead, Progvis highlights the button that corresponds to the next step in yellow. In Fig. 4.7, the second button in the right thread is highlighted. This means that when you click *Next step* in the green bar, Progvis will click that button for you, thereby letting thread 2 execute for a while, and then highlight the next step in the problematic interleaving. If you don't want to click *Next step* manually, you can also click *Play/Pause* to have Progvis automatically go through the full sequence for you. At any point, you may click *Close* to go back to the normal view where you

⁷It does this in a somewhat intelligent manner to reduce the search space, but the end result matches this intuition.

⁸The message includes "...at least not with this main program." to remind you that it will only find errors in code that is actually executed by your `main` function. We will see the importance of this later in the book.

are in control. You can even click *Close* mid-way through the sequence to see what happens if you make different decisions towards the end of the sequence.

Practice: The file `add-3.c` (see below) contains 3 commented lines marked A, B, and C. Which of these three lines cause a data race when uncommented? First try to reason about the behavior of the program, and then use the model checker in Progviz to verify your answer.

```
int result;
struct semaphore has_result;

void thread_fn(void) {
    for (int i = 0; i < 4; i++) {
        result += 2;
    }
    sema_up(&has_result);
    // result += 1; // (A)
}

int main(void) {
    sema_init(&has_result, 0);
    // printf("result=%d\n", result); // (B)
    thread_new(&thread_fn);
    // printf("result=%d\n", result); // (C)
    sema_down(&has_result);
    printf("result=%d\n", result);
    return 0;
}
```

4.4.2 How to Use Semaphores?

Now that we have seen how to solve a simple problem in practice, it is worth taking a moment to reflect on the solution process. In this case, we were lucky enough that the program we wanted to synchronize was similar enough to our previous experiment. As you might imagine, many situations are not that clear-cut in practice. We want to find a more general approach.

In Section 4.1 (when we introduced the operations), we noted that it is possible to consider the semaphore to keep track of some resource in the program. This view is indeed often useful to help figure out what value to initialize the semaphore to, when to call `sema_up` and when to call `sema_down`.

To illustrate this idea, consider the program `add-2.c` that we just added a semaphore to. The semaphore here is named `has_result` rather than `done`. This change in name was done to highlight the idea that the semaphore counts the number of finished results that have been produced by the other threads. In this case we only have one other thread, so the semaphore will either contain 0, meaning “`result` does not contain anything useful yet” or 1, meaning “`result` contains a meaningful value”.

Once we have found a proper resource to count, the usage of the semaphore follows naturally. Since we don’t have a result at the start of the program, it follows that we need to initialize the semaphore to 0. When `thread_fn` is done, it has *produced* a value inside `result`. As such we want to increase the semaphore

from 0 to 1, hence we call `sema_up`. Similarly, in `main` when we call `printf`, we want to *consume* a result, so we call `sema_down` to decrease the semaphore. Note that we don't technically "destroy" the contents of the variable in this program. The idea of *producing* and *consuming* values is, however, useful regardless as it allows us to think about the program in terms of which thread "owns" some shared data currently. In this case we consider `thread_fn` to "own" the variable `result` initially. This means that `thread_fn` is free to do what it pleases with the variable until it declares that it is done by calling `sema_up`, thereby officially declaring that it has *produced* the final value of `result`. This allows `main` to *consume* the value (i.e., `sema_down` is allowed to return). At this point `main` has successfully *consumed* the value, and thereby "owns" the variable `result`, again meaning that `main` may do what it pleases with the variable. If other threads would be interested in the value, we could let `main` *produce* the value in `result` after printing it (an potentially modifying it). This would let other threads *consume* it and use it for themselves.

As you might imagine, the difficult part of the idea of counting some resource is to *find* a suitable resource to count. This is crucial since only `sema_down` waits when the semaphore's counter would become negative. The key idea is to find some resource that is zero exactly when some thread wishes to wait for it. One way to perform a sanity-check of some choice is to first add the relevant calls to `sema_init`, `sema_up`, and `sema_down` as dictated by the choice of resource. Then, imagine that any threads started by calls to `thread_new` take a very long time to start, and think about what happens when the remaining thread continues to execute. Will it eventually reach a call to `sema_down` where you would expect the thread to wait? Furthermore, does the semaphore passed to `sema_down` contain zero at that time? If so, you have likely picked a good resource to count. We can see this in `add-2.c`. If we imagine that the thread running `thread_fn` is delayed for a long time, we can easily see that the main thread quickly reaches `sema_down` and that the semaphore `has_result` is zero, which means that the thread will wait for `thread_fn` to start. This can either be done by tracing the program in your mind, or by using Progviz and simply choosing to only step one of the threads.

Practice: The program `add-two.c` (see below) contains a program that is similar to the one we have used so far. The difference is that it starts *two* threads and produces *two* result variables (`result1` and `result2`). How can we use *a single* semaphore to synchronize the program? You can solve the problem by adding one or more calls to `sema_init`, `sema_up`, and `sema_down` to the lines that contain a question mark (?). If done correctly, the program should always print `result_a=8` and `result_b=12`. You can use the model checker to verify your solution.

Hint: Consider `result_a` and `result_b` to be a single resource, and count how many of them are filled.

```
int result_a;
int result_b;
struct semaphore has_result;

void thread_fn_a(void) {
    for (int i = 0; i < 4; i++) {
        result_a += 2;
    }
    // ?
}

void thread_fn_b(void) {
    for (int i = 0; i < 4; i++) {
        result_b += 3;
    }
    // ?
}

int main(void) {
    // ?
    thread_new(&thread_fn_a);
    thread_new(&thread_fn_b);
    // ?
    printf("result_a=%d\n", result_a);
    printf("result_b=%d\n", result_b);
    assert(result_a == 8);
    assert(result_b == 12);
    return 0;
}
```

4.5 Multiple Semaphores

In many cases it is not enough to use a single semaphore to synchronize a program. In particular, when we want a thread to wait, we usually want it to *wait for* something particular to happen. If there are multiple things that happens in the program that some thread might want to wait for, then we need one semaphore for each thing we want to wait for to achieve this.

To illustrate this idea, consider the program in `add-multi-1.c` (also in Listing 4.5). It is a version of the program `add-2.c` but with two threads, `thread_fn_a` and `thread_fn_b`, that write to `result_a` and `result_b` respectively. As

```

int result_a;
int result_b;
struct semaphore has_result;

void thread_fn_a(void) {
    for (int i = 0; i < 4; i++) {
        result_a += 2;
    }
    sema_up(&has_result);
}

void thread_fn_b(void) {
    for (int i = 0; i < 4; i++) {
        result_b += 3;
    }
    sema_up(&has_result);
}

int main(void) {
    sema_init(&has_result, 0);
    thread_new(&thread_fn_a);
    thread_new(&thread_fn_b);

    sema_down(&has_result);
    printf("result_a=%d\n", result_a);

    sema_down(&has_result);
    printf("result_b=%d\n", result_b);

    return 0;
}

```

Listing 4.5: The program add-multi-1.c.

before, the code in `main` starts the threads, and prints the results after waiting for the threads to finish. In this case, the main function wishes to print `result_a` first. To do this, the main thread only needs to wait for `thread_fn_a` to complete, and therefore it calls `sema_down` once before printing the result. Afterwards, it calls `sema_down` once more to wait for `thread_fn_b` to complete and then prints `result_b`.

Practice: The program `add-multi-1.c` is not correct. It contains a data race and may print `result_a=0`. Can you see why? Use Progviz to find the error. Start by trying to find the issue yourself by stepping the program manually. You can also use the model checker to find the problem.

As you likely realized above, the fundamental problem is that we are not able to specify *which* thread to wait for since we only have one semaphore. Using the analogy of resources from the previous section, we have decided to use the semaphore to count *how many results are produced*, thereby treating both

results as equivalent.⁹ This means that after we call `sema_down` once in the main function, we know that *one* result is ready, but we don't know which one. As such, if `thread_fn_b` finishes first and signals its semaphore, the main thread will wake up and print whatever is in `result_a`, even though `result_a` was not yet ready.

```

int result_a;
int result_b;
struct semaphore has_result_a;
struct semaphore has_result_b;

void thread_fn_a(void) {
    for (int i = 0; i < 4; i++) {
        result_a += 2;
    }
    sema_up(&has_result_a);
}

void thread_fn_b(void) {
    for (int i = 0; i < 4; i++) {
        result_b += 3;
    }
    sema_up(&has_result_b);
}

int main(void) {
    sema_init(&has_result_a, 0);
    sema_init(&has_result_b, 0);
    thread_new(&thread_fn_a);
    thread_new(&thread_fn_b);

    sema_down(&has_result_a);
    printf("result_a=%d\n", result_a);

    sema_down(&has_result_b);
    printf("result_b=%d\n", result_b);

    return 0;
}

```

Listing 4.6: The program `add-multi-2.c` which fixes the issue in `add-multi-1.c`.

To solve the problem, we need to use two separate semaphores. One that counts whether `result_a` contains a result, and another that counts whether `result_b` contains a result. This can be implemented as in `add-multi-2.c`. The program replaces the single semaphore `has_result` with two separate semaphores, `has_result_a` and `has_result_b`. The function `thread_fn_a` then calls `sema_up` on `has_result_a` when it is done producing `result_a`. Similarly, `thread_fn_b` calls `sema_up` on `has_result_b` when it is done producing `result_b`. This means

⁹As you probably saw in the last section, this is completely fine to do as long as we wait for both results to become ready, or have some other way to determine which result to actually use.

that when the main thread wishes to consume `result_a`, it can call `sema_down` on `has_result_a`. This means that the main thread no longer waits for any of the two threads to become ready, but instead explicitly waits for `thread_fn_a` which produces `result_a` to become ready. This change solves the problem, and the program now works as expected.

Practice: Use Progvis to compare the differences between `add-multi-1.c` and `add-multi-2.c`. Pay special attention to what happens if you let the main thread execute until it starts waiting during its first call to `sema_down`, and you then let the thread running `thread_fn_b` (thread 3) run to completion without touching the thread running `thread_fn_a` (thread 2).

The situation above illustrates why we need to use multiple semaphores to specify what to wait for. A similar (but related, depending on your view) situation is when it is important *which* out of multiple threads that should wake up when `sema_up` is called. This issue is illustrated by the program `multi-sema-1.c` (see Listing 4.7).

```
int result;
struct semaphore has_result;

void thread_a(void) {
    result = 10;
    sema_up(&has_result);
}

void thread_b(void) {
    sema_down(&has_result);
    result += 5;
    sema_up(&has_result);
}

int main(void) {
    sema_init(&has_result, 0);
    thread_new(&thread_a);
    thread_new(&thread_b);

    sema_down(&has_result);
    printf("result=%d\n", result);
    // assert(result == 15);

    // Avoid confusing the model checker:
    sema_up(&has_result);
    return 0;
}
```

Listing 4.7: The program `multi-sema-1.c`.

This program resembles the programs we have seen so far, but with some differences. First and foremost, the loops are replaced with assignments. This simplification is just done to help focus on the synchronization issues, and has

no impact on the synchronization. Perhaps more importantly, the program now attempts to run the code in the three functions in sequence. First, `result = 10`, then `result += 5` and finally `printf(..., result)`.

To achieve this, the programmer considered the semaphore `has_result` to count whether `result` contains a useful value or not. Similarly to before, the programmer then thought that `thread_a` first “owns” the variable `result`. As such, the thread is free to store the value 10 in `result`. It then declares that a value has been *produces* by calling `sema_up`. As such, some other thread may now take over ownership of the variable by *consuming* it using `sema_down`. The intention was that `thread_b` would then *consume* the the variable by calling `sema_down`. Once `sema_down` returns, the thread owns the variable and may update it by increasing it by 5. When it is done, it declares that it has *produced* a new value in the variable and calls `sema_up`. This means that the main thread is finally able to *consume* the value and print the final result.

Practice: As you might have realize from the wording above, the program `multi-sema-1.c` does not work as intended. The intended behavior is for the program to always print `result=15`. However, the program sometimes prints `result=10`. In contrast to most other examples we have seen so far, this program is actually free from data races and thereby well-defined according to the C standard. It misbehaves in spite of this.

Use Progvis to examine the program and find the problem. Since the model checker does not know what the expected output of the program is, it will not automatically find the error. You can use an `assert` statement (available as a comment in the program) to let the model checker know that you expect `result` to be 15. That way, the model checker will be able to find the error.

From above, you have likely realized that the problem in `multi-sema-1.c` is that when `thread_a` calls `sema_up`, *one of* the two threads (`thread_b` and `main`) are allowed to continue. The programmer intended that `thread_b` should always wake, but as we saw above this is not always the case. It is equally valid for the implementation to let the main thread wake up immediately, which means that `thread_b` never gets a chance to execute.¹⁰ This is a key insight, which is worth highlighting:

- When one thread calls `sema_up` on a semaphore that two or more threads are waiting for, the implementation chooses any one of the threads to wake up. This choice is *arbitrary*.

Since the implementation may pick an *arbitrary* thread out of the ones that are waiting, it means that it is equally valid for either `thread_b` or `main` to wake up when `thread_a` calls `sema_up`. This arbitrary choice is what causes the non-deterministic behavior in `multi-sema-1.c`. In this case it leads to the program behaving incorrectly, since the program did not take this behavior into account.

Before we continue, it is worth noting that Progvis simulates a slightly stricter model in order to be easier to understand. In Progvis, semaphores

¹⁰This is the reason for the extra `sema_up` at the end of `main`. Without it, `thread_b` never wakes up, which the model checker in Progvis reports as a deadlock.

(and other synchronization primitives) wake threads in the order they started waiting, using a FIFO queue. This is *one* of many valid choices for this arbitrary behavior. Semaphores in your system are likely not this strict, even though they will still be fair in the long run (i.e., if we do the same thing many times after each other, it will pick either thread roughly 50% of the time). However, in spite of the semaphore being completely fair, we see the same problem in Progviz. The reason for this is that we generally cannot control in what order threads start waiting for the semaphore. This uncertainty therefore gives rise to the same type of non-determinism that is already a part of the semaphore, and the reason why it is usually not worth the performance impact of implementing a FIFO behavior in practice.

```
int result;
struct semaphore thread_a_done;
struct semaphore thread_b_done;

void thread_a(void) {
    result = 10;
    sema_up(&thread_a_done);
}

void thread_b(void) {
    sema_down(&thread_a_done);
    result += 5;
    sema_up(&thread_b_done);
}

int main(void) {
    sema_init(&thread_a_done, 0);
    sema_init(&thread_b_done, 0);
    thread_new(&thread_a);
    thread_new(&thread_b);

    sema_down(&thread_b_done);
    printf("result=%d\n", result);
    assert(result == 15);
    return 0;
}
```

Listing 4.8: The program multi-sema-2.c which solves the problems from multi-sema-1.c.

To be able to control which thread wakes up when we call `sema_up` we must once again use more than one semaphore. That is, instead of treating whether `result` contains a value or not as a *single* resource that we count with a semaphore, we think of it as two different resources. The first one is: “has `thread_a` produced a value in `result`?” and the second one is “has `thread_b` updated the value in `result`?”. Again, if we replace the single semaphore `has_result` with two different semaphores to reflect this new view of the problem (we call them `thread_a_done` and `thread_b_done`), we get the program in multi-sema-2.c (Listing 4.8).

It is particularly interesting to note that it is now clear that `thread_a` produces a result that only `thread_b` consumes, since those are the only two functions that use `thread_a_done`. Similarly, the link between `thread_b` and `main` is clear, since those two functions are the only places where `thread_b_done` is used. In essence, we have applied the pattern of using semaphores to wait for something twice.

It is worth noting that some care needs to be taken in cases like this. Here we have multiple semaphores that together make sure that `result` is never accessed by more than one thread at the same time. In this case, it is quite easy to see that this is the case since the semaphores cause the threads to execute in a fixed order. This may, however, not always be the case.

To make the separation of a single resource into two separate resources clearer, we could have also split the `result` variable into two separate variables: `result_a` that contains the output from `thread_a`, and `thread_b` that contains the output from `thread_b` (see `multi-sema-3.c` and Listing 4.9). This division makes it clearer that we are working with two different resources, and thereby what the purpose of the two semaphores are. It is, however, not always trivial to do this kind of restructuring.

```
int result_a;
int result_b;
struct semaphore thread_a_done;
struct semaphore thread_b_done;

void thread_a(void) {
    result_a = 10;
    sema_up(&thread_a_done);
}

void thread_b(void) {
    sema_down(&thread_a_done);
    result_b = result_a + 5;
    sema_up(&thread_b_done);
}
```

Listing 4.9: Another possible solution of the problem in `multi-sema-1.c`. The full code is available in `multi-sema-3.c`.

4.6 Mutual Exclusion

As mentioned in the beginning of this chapter, concurrency problems can be grouped into two large groups. We have now seen how semaphores can be used to solve problems in the first group, waiting for something. What remains is the second group, to ensure *mutual exclusion*. This means that we need to protect shared data that may be modified concurrently in order to avoid *data races*.

It turns out that semaphores are powerful enough for this task too. In fact, it is possible to implement all other synchronization primitives using only semaphores. Thereby it is possible to solve all synchronization problems using

only semaphores if we wish to. However, as we shall see in the next section, semaphores are not always the *most convenient* way to solve all problems. Ensuring *mutual exclusion* is one such area. While it is possible to use semaphores for this task, it is often better to use *locks*. Since locks are intended to ensure mutual exclusion they communicate the intent of the programmer clearer, and they are able to detect some common usage bugs related to how locks are used. Since locks are more suitable for this task, this chapter only introduces the problem and the semaphore-based solution briefly. The topic will be discussed in greater depth in the next chapter.

To illustrate the problem, remember the last program in Section 3.5 where we attempted to add 2 and 5 to a shared variable concurrently. If we solve the waiting problem using semaphores, we get the program in `mutex-1.c` (see Listing 4.10).

```
int result;
struct semaphore has_result;

void thread_fn(void) {
    for (int i = 0; i < 4; i++) {
        result += 2;
    }
    sema_up(&has_result);
}

int main(void) {
    sema_init(&has_result, 0);
    thread_new(&thread_fn);
    for (int i = 0; i < 4; i++) {
        result += 5;
    }
    sema_down(&has_result);
    printf("result=%d\n", result);
    return 0;
}
```

Listing 4.10: The program `mutex-1.c`.

This program is similar to `add-2.c` (Listing 4.4). The only difference is that we have a loop that adds 5 to `result` in `main`, before the call to `sema_down`. Since the loop is *before* `sema_down`, it may execute concurrently with the loop in `thread_fn`, and both threads may therefore modify `result` concurrently. If you use the model checker in Progviz, it will quickly point out the issue.

One way to solve this issue would be to move the `sema_down` call to before the loop in `main`. This would make sure that we execute the loop in `thread_fn` first, and the loop in `main` afterwards, thereby avoiding the data race. However, imagine a situation where the numbers 2 and 5 need to be computed through some expensive computations. This may look like in `mutex-2.c` (Listing 4.11) which simulates expensive computations with a call to `timer_msleep`, which just sleeps for the specified number of milliseconds.¹¹

¹¹Note that Progviz ignores this call, since you are in control of scheduling anyway.

```

int result;
struct semaphore has_result;

int expensive_computations(int x) {
    // Simulate time-consuming computations.
    timer_msleep(1000);
    return x;
}

void thread_fn(void) {
    for (int i = 0; i < 4; i++) {
        int to_add = expensive_computations(2);
        result += to_add;
    }
    sema_up(&has_result);
}

int main(void) {
    sema_init(&has_result, 0);
    thread_new(&thread_fn);
    for (int i = 0; i < 4; i++) {
        int to_add = expensive_computations(5);
        result += to_add;
    }
    sema_down(&has_result);
    printf("result=%d\n", result);
    return 0;
}

```

Listing 4.11: The program `mutex-2.c`, which simulates time-consuming computations.

In this situation, we can imagine that the point of using an additional thread was to speed up the program. In this case, the program uses two threads to call `expensive_computations` concurrently. This means that the system may use 2 different threads to halve the execution time of the program. Therefore, the solution above (moving the call to `sema_down` to before the loop) is not ideal, since it would make `thread_fn` do all of its work first, then let the main thread do all of its work, thereby eliminating the speedup.

Practice: Compile and run `mutex-2.c` on your system and measure its execution time. As before, you can compile the program using:

```
make mutex-2
```

To measure the execution time, you can run the program using:

```
time ./mutex-2
```

Measure the execution time of the program as shown in Listing 4.11. Then move the call to `sema_down` to before the loop inside `main` and see how this changes the execution time.

As you likely have seen from above, if you run the program as originally written it will take 4 seconds to run, but if we move the semaphore it will instead take 8 seconds to run.

Because of this, we have a different situation compared to before. Here, we actually want to run the two loops concurrently. However, we need to make sure that the two threads don't touch `result` concurrently. As previously mentioned, we can solve this using semaphores. We just need to view the problem in a different way compared to before. Here, the goal is to make sure that at most one thread accesses `result` at the same time. As such, if we can find a resource related to that property and use a semaphore to count it, we can solve the problem.

In this case, we can choose to count "how many additional threads may access `result`?". So we start by creating a semaphore named `access_result` to count this resource. At the start of the program no thread is accessing `result`, which means that we can allow one additional thread to access it. Therefore we initialize the new semaphore to `1` at the start of `main`.

Since we now count threads that access `result`, we need to look for places in the code that access `result` and update the semaphore accordingly. Before accessing `result`, the code needs to *consume* one instance of the result (i.e., decrease the number of additional threads that may access `result`). As such, we need to add a call to `sema_down` before the two lines `result += to_add`. Similarly, when we are done using `result` we need to *produce* an instance of the resource, since we are done accessing `result`. Therefore, we add a call to `sema_up` after the two lines `result += to_add`. Similarly, we need to add calls to `sema_down` and `sema_up` in the same way since it accesses `result`. These changes results in the code in `mutex-3.c` (Listing 4.12), which solves the issues with shared data.

Practice: Verify that the program is free from data races using the model checker in Progvis. Also verify that the program still executes the two loops in parallel by running the program outside of Progvis and verifying that it finishes in around 4 seconds rather than 8.

As mentioned previously, there are many nuances to the placement of `sema_down` and `sema_up` when working with shared data. We will examine them in further detail in the next chapter, but using locks instead of semaphores.

```
int result;
struct semaphore has_result;
struct semaphore access_result;

int expensive_computations(int x) {
    // Simulate time-consuming computations.
    timer_msleep(1000);
    return x;
}

void thread_fn(void) {
    for (int i = 0; i < 4; i++) {
        int to_add = expensive_computations(2);
        sema_down(&access_result);
        result += to_add;
        sema_up(&access_result);
    }
    sema_up(&has_result);
}

int main(void) {
    sema_init(&has_result, 0);
    sema_init(&access_result, 1);
    thread_new(&thread_fn);
    for (int i = 0; i < 4; i++) {
        int to_add = expensive_computations(5);
        sema_down(&access_result);
        result += to_add;
        sema_up(&access_result);
    }
    sema_down(&has_result);
    sema_down(&access_result);
    printf("result=%d\n", result);
    sema_up(&access_result);
    return 0;
}
```

Listing 4.12: The program `mutex-3.c` which solves the issues in `mutex-1.c` and `mutex-2.c`.

4.7 Exercises

1. The file `ex1.c` contains a program that starts another thread. Both threads access the global variable `shared`. Use semaphores to ensure that the two lines `shared += 2` and `shared *= 10` executes in the correct order so that it ends up containing 20 at the end of the program. You can verify your solution using Progviz.
2. The file `ex2.c` contains a program that starts two other thread. As indicated in comments, both threads need to wait for `main` to set `shared`. Use a semaphore to ensure that both threads wait properly. You can verify your solution using Progviz.
3. The file `ex3.c` contains a program that starts two other threads. The idea is that `thread1` computes a value that `thread2` needs, so `thread2` needs to wait for it. Similarly, `thread2` produces a value that `main` needs, so `main` needs to wait for the value to be finished. Use semaphores to ensure that the threads wait for each other properly. You can verify your solution using Progviz.
4. In the program `mutex-3.c` we protected the final call to `printf` with calls to `sema_down` and `sema_up`. This is actually not necessary. Use the model checker in Progviz to verify that they are indeed not needed. Why is this the case?
5. What happens if you move the call to `sema_down` inside the two loops one line earlier in the program `mutex-3.c` (i.e., so that it is before the call to `expensive_computations`). Does the program still behave correctly? Does this affect the program's performance?

Locks

In the previous chapter we saw how semaphores can be used to solve synchronization errors. Even though semaphores are enough to solve all concurrency issues,¹ they are not always the most convenient option. One example is the common problem of achieving *mutual exclusion* for some variables. As we saw in the previous chapter, semaphores are more than capable of doing this, but due to the versatility of semaphores it is easy to make mistakes and the code quickly gets difficult to read.

Since this problem is common, there is a separate synchronization primitive whose sole purpose is to protect variables through mutual exclusion. It is called a *lock* since we can consider it to lock some data so that only one thread may access it concurrently. Since locks are used to achieve mutual exclusion, they are sometimes called *mutexes*.

Since locks are specialized for mutual exclusion, they are able to provide help in the form of better error checking. However, as we shall see, this also means that they are less powerful than semaphores. It is thus not possible to use locks to wait for something to happen. In Chapter 8 we will add this capability to locks by introducing *condition variables* that can be used alongside locks to wait for something to happen.

5.1 Semantics

As in the previous section we start by introducing the available operations and their semantics. As before, we focus on the usage here. The full definitions are available in Chapter A for the interested reader. Again, the `lock` is an opaque data structure, and it is therefore only possible to modify it through the functions below.

A lock itself can be thought of as the lock on the door to a room that contains some data that need to be protected. A thread may then try to *acquire* the lock by entering the room and locking the door from the inside. This means that other threads who also try to acquire the lock needs to wait for their turn. A thread who has previously acquired a lock may then *release* the lock by opening the door and leaving the room. The lock ensures that at most one thread is ever inside the room. We say that the thread that is inside

¹Unless we work with low-level programming and/or interrupts.

the room is *holding* the lock. One thing that is important to note is that locks are not special in the sense that they automatically protect some variables in the program. We must manually remember to acquire and release the relevant locks for the analogy described above to work.

```
struct lock l;
```

The line above defines a variable named `l` that contains a lock. Note: In contrast to C, Progviz uses the same namespace for types and variables (like C++ does). This means that while it is possible to define a variable named `lock`, it will shadow the type `lock`, which makes it difficult to create other locks later on in the program.

```
lock_init(&l);
```

As with semaphores, each lock variable needs to be initialized exactly once before it is used. The first and only parameter is a pointer to the lock that should be initialized.

```
lock_acquire(&l);
```

Attempt to *acquire* the lock. If the lock is not held by another thread, the calling thread will be allowed to continue and will hold the lock from this point onward. If the lock is already held by another thread, `lock_acquire` will make the calling thread sleep until the thread that is holding the lock releases it. At that point the current thread will finish acquiring the lock.

```
lock_release(&l);
```

Called by a thread that is currently holding the lock to *release* the lock. This may cause another thread that is currently waiting for the lock to wake up and acquire the lock.

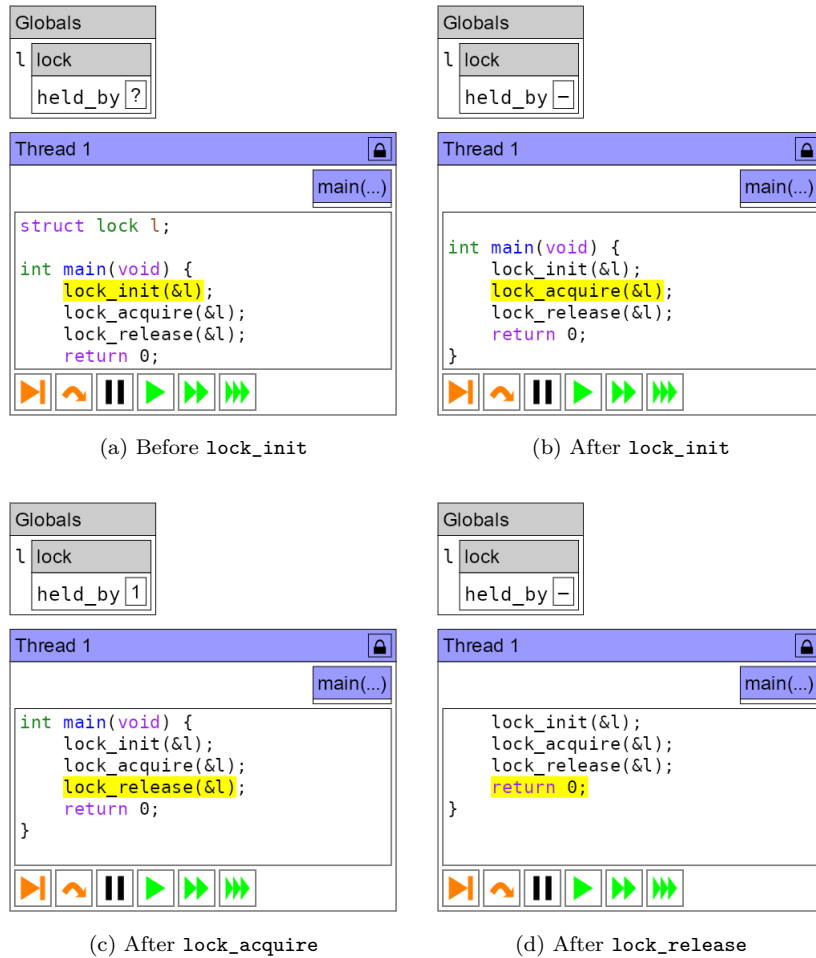
To see how locks are visualized by Progviz, we will use the sequential program `simple-semantics.c`. As we can see from Listing 5.1, the program simply initializes a lock, acquires it and releases it.

```
struct lock l;

int main(void) {
    lock_init(&l);
    lock_acquire(&l);
    lock_release(&l);
    return 0;
}
```

Listing 5.1: A program that illustrates the semantics of the program.

Figure 5.1 shows how Progviz illustrates the state of the lock throughout the program. In the first picture (Fig. 5.1a), we can see that since a lock is a complex data type, it is drawn as a box labeled *lock* to indicate the type of the data. The contains a single line named `held_by` that shows which thread is currently holding the lock. Since the lock is not yet initialized in Fig. 5.1a, the contents are shown as a question mark (?).

Figure 5.1: Progvis' visualization of the program `simple-semantics.c`.

In the next picture (Fig. 5.1b) we have let the program execute `lock_init`. As such, the lock is initialized and ready to be used. However, no thread has yet acquired the lock, so the contents of `held_by` is shown as a dash (-) to reflect this. In the next step (Fig. 5.1c) the main thread has executed `lock_acquire` and thereby acquired the lock. We can see this by observing that the value of `held_by` is now 1. Finally, after the thread executes `lock_release` (in Fig. 5.1c) the contents of `held_by` is a dash again, since the lock is no longer held by any thread.

5.2 Relation to Semaphores

Since semaphores are powerful enough to solve all synchronization problems, it is possible to implement a lock using a semaphore. While this rarely needs to be done in practice, it is useful to see the relation between the two to understand their semantics in more detail.

Table 5.1: List of how a lock can be implemented using a semaphore.

Lock operation	Semaphore operation
<code>struct lock l;</code>	<code>struct semaphore l;</code>
<code>lock_init(&l);</code>	<code>sema_init(&l, 1);</code>
<code>lock_acquire(&l);</code>	<code>sema_down(&l);</code>
<code>lock_release(&l);</code>	<code>sema_up(&l);</code>

Table 5.1 shows how each operation for a lock can be replaced by an equivalent operation for a semaphore. This implementation is correct for all programs that use locks correctly. However, the equivalence in Table 5.1 omits the additional error checking performed by locks.

The error checking is illustrated by the program `lock-as-sema.c` (Listing 5.2). This program is similar to the program we used in Section 4.3 to illustrate the semantics of a semaphore. The difference is that the semaphore operations have been replaced with equivalent lock operations according to Table 5.1, as indicated by the comments to the right of the relevant lines. Looking at the corresponding semaphore operations, the only difference from the code used in Section 4.3 is that the initialization is split into two steps. Since there is no operation that corresponds to `sema_init(&l, 0)`, we instead call `sema_init(&l, 1)` to initialize the semaphore to 1, and then call `sema_down` to decrease the counter to zero.

```

struct lock l;           // struct semaphore l;

void thread_fn(void) {
    lock_release(&l);    // sema_up(&l);
}

int main(void) {
    lock_init(&l);       // sema_init(&l, 1);
    lock_acquire(&l);    // sema_down(&l);
    thread_new(&thread_fn);
    lock_acquire(&l);    // sema_down(&l);
    return 0;
}

```

Listing 5.2: Trying to use a lock as a semaphore to wait for `thread_fn`.

Even though the code in Listing 5.2 would work properly if we replaced the lock with a semaphore as indicated by the comments, it does *not* work as it is currently written. This is because the lock assumes that it is protecting some data. The implementations in the library provided with this book and in Progviz both verify this, but other implementations may not give an explicit error and silently behave incorrectly instead.

Illustrates two problems that lock implementations assume, and are often able to verify. Load the program in Progviz and step the threads as described below to see the errors yourself.

1. A lock assumes that the lock is *released* by the thread that currently *holds* the lock (i.e., the thread that previously *acquired* the lock). This helps the programmer to detect cases where they have forgotten to call `lock_acquire` or multiple calls to `lock_release`.

You can see this error if you step the main thread until it has executed `thread_new`, and then step the newly started thread so that it executes `lock_release`. In Progviz, this causes the thread to crash and Progviz informs you of the problem. The C library behaves similarly. You can cause this problem to occur by adding a call to `timer_msleep(100)` after `thread_new`.

If a semaphore would be used instead of a lock, this would just increment the counter one additional time, meaning that the semaphore would no longer ensure mutual exclusion.

2. Locks also detect cases where one thread attempts to *acquire* a lock that is already held. The locks used in this book do not allow this to happen and report it as an error. As in the previous point, an implementation may also assume that this does not happen and do something unexpected instead.²

A third option is to implement what is sometimes referred to as *recursive locks* or *recursive mutexes*. In this case, acquiring a lock that is already only causes the lock to require an additional *release* before it is actually released. This is useful in cases where two functions, say `f` and `g` both acquire the same lock, but `f` calls `g` while holding the lock.

You can see this error if you step the main thread until the end without stepping the other thread. In the C library, you can cause this to happen by adding a call to `timer_msleep` before `lock_release` in `thread_fn`.

If a semaphore were used instead of a lock in this situation, the extra call to `sema_down` would cause the thread to wait forever, causing a deadlock.

5.3 Using Locks for Mutual Exclusion

Now that we have seen how a lock works and how it relates to a semaphore, we can exemplify this knowledge by replacing the semaphore `access_result` in Listing 4.12 with a lock. If we just use the information in Table 5.1 to replace `sema_init`, `sema_down`, and `sema_up` with `lock_init`, `lock_acquire`, and `lock_release` we get the program in `mutex.c` (Listing 5.3). Note that we also replaced the semaphore variable `access_result` with a lock named `result_lock`.

Examining the code in Listing 5.3, we can now clearly see how locks are designed to protect shared data. Each time the program uses the variable `result`, the program has made sure to *acquire* the associated lock beforehand and to *release* the lock afterwards. This makes it possible to think about locks as if they are marking a region in the source code where the current thread has exclusive access to some variable or resource.³ While this is often a useful way to think about locks, some care needs to be observed. For example, if the call to `lock_acquire` is inside an `if` statement, the lock is not always acquired

²Often, this causes the thread to wait forever, but anything is allowed to happen.

³This idea is used for *scoped locks* in languages like C++, for example.

```

int result;
struct semaphore has_result;
struct lock result_lock;

int expensive_computations(int x) {
    // Simulate time-consuming computations.
    timer_msleep(1000);
    return x;
}

void thread_fn(void) {
    for (int i = 0; i < 4; i++) {
        int to_add = expensive_computations(2);
        lock_acquire(&result_lock);
        result += to_add;
        lock_release(&result_lock);
    }
    sema_up(&has_result);
}

int main(void) {
    sema_init(&has_result, 0);
    lock_init(&result_lock);
    thread_new(&thread_fn);
    for (int i = 0; i < 4; i++) {
        int to_add = expensive_computations(5);
        lock_acquire(&result_lock);
        result += to_add;
        lock_release(&result_lock);
    }
    sema_down(&has_result);
    lock_acquire(&result_lock);
    printf("result=%d\n", result);
    lock_release(&result_lock);
    return 0;
}

```

Listing 5.3: Replacing the semaphore in Listing 4.12 with a lock.

and the idea of a syntactical region where it is safe to access some variable falls apart.

In the code, the connection between the `result` variable and the lock is also made clear by naming the lock `result_lock`. Note, however, that there is no formal connection between the `result` variable and the lock `result_lock`. As such, creating a lock with the appropriate name does not automatically protect variable. Since the language does not know which variables we want to protect, we need to protect them explicitly by calling `lock_acquire` and `lock_release`. Similarly, there is no way for the language to verify that we have not forgotten to protect some variable. This is why it is important to be meticulous about documenting which variable(s) the lock is intended to protect (e.g., through naming), so that it is easy to verify that the lock is actually acquired when

necessary.

5.4 Benefits of Error Checking

We have now seen how both semaphores and locks can be used to achieve mutual exclusion and thereby protect shared data from concurrent modification. Even though locks can be implemented using semaphores (as we saw in Table 5.1), they are designed to protect shared data. The benefit of this is that it makes the intent of the programmer clearer. If we see a lock in the code, we know that it protects some shared data. This is not necessarily true if we see a semaphore, since the semaphore might also be used to wait for something to happen in the program. The focus on protecting shared data also means that the lock can check for errors in using the lock. We have already seen some benefits of this in Section 5.2, but in a contrived setting. The remainder of this section will show how the error checking can help catching more realistic issues.

To illustrate this, we will use the `decrease` function in the program `decrease.c` as an example. The function is shown in Listing 5.4. It decreases the global variable `counter` with the value of the parameter `with`, but ensures that the value of `counter` does not become negative. For example, calling `decrease(14)` does nothing since the counter would become negative. Calling `decrease(2)` would subtract 2 from the counter, resulting in it containing 8 after the call.

```
int counter = 10;

void decrease(int with) {
    if (counter < with)
        return;
    counter -= with;
}
```

Listing 5.4: Implementation of the function `decrease` from the program `decrease.c`.

Now assume that we wish to call the `decrease` function from multiple threads concurrently. For this to work properly, we need to protect the `counter` variable somehow. Since the goal is to protect data (i.e., create *mutual exclusion*), we use a lock. However, imagine that we were a bit careless when adding the lock and ended up with the code in `decrease-error-1.c` (Listing 5.5). As we can see, we defined the variable `counter_lock` to protect the `counter` variable, acquire the lock at the start of `decrease` and release it at the end of `decrease`. However, we forgot to account for the fact that `decrease` may return early if it determines that `with` is larger than `counter`, and forgot to release the lock in that case.

```
int counter = 10;
struct lock counter_lock;

void decrease(int with) {
    lock_acquire(&counter_lock);
    if (counter < with)
        return;
    counter -= with;
    lock_release(&counter_lock);
}
```

Listing 5.5: Incorrect application of a lock from `decrease-error-1.c`.

Practice: Assume that a thread calls `decrease(14)` followed by `decrease(2)` using the implementation of `decrease` in Listing 5.5. What happens when `decrease(2)` is called after the call to `decrease(14)` failed to release the lock?

The `main` function in `decrease-error-1.c` calls `decrease(14)` followed by `decrease(2)`. As such, you can see what happens by running the program in Progviz.

What would happen if we use a semaphore instead of a lock to protect `counter`?

As you have noticed from above, since the implementation of `decrease` in Listing 5.5 never releases the lock when `decrease(14)` is called, the lock is already held by the current thread when `decrease(2)` is called later. The lock implementation is able to detect this, since it is always invalid for a single thread to acquire the same lock more than once.⁴ If we would have used a semaphore instead, the call to `sema_down` that replaces `lock_acquire` would instead have caused a deadlock that would have been much harder to debug.

⁴Unless *recursive locks/mutexes* are used.

The program `decrease-error-2.c` contains another error that is perhaps more obvious but not uncommon. Here, the programmer has forgotten to acquire the lock at the start of the `decrease` function, and they thereby fail to protect the `counter` variable. While it is easy to spot the issue in this particular case, the same type of error would also arise if the programmer had accidentally acquired some other lock at the start of the function.

```
int counter = 10;
struct lock counter_lock;

void decrease(int with) {
    if (counter < with)
        return;
    counter -= with;
    lock_release(&counter_lock);
}
```

Listing 5.6: Another incorrect lock application of a lock from `decrease-error-2.c`.

Practice: Assume that a thread calls `decrease(2)` using the implementation of `decrease` in Listing 5.6. What happens?

The `main` function in `decrease-error-2.c` calls `decrease(2)`, so you can see what happens using `Progvis`.

What would happen if we use a semaphore instead of a lock to protect `counter`?

By running the code above, you will see that the call to `lock_release` fails since the lock was not acquired previously. If we would have used a semaphore to protect `counter`, the situation would actually have been worse than in `decrease-error-1.c`. The mistake in `decrease-error-1.c` at least led to a deadlock, which is quite noticeable. On the other hand, the mistake in `decrease-error-2.c` would not only lead to failure to achieve mutual exclusion in the `decrease` function. Since the final call to `sema_up` (which replaced `lock_release`) would increment the semaphore without previously decrementing it, it means that *all other parts* of the code that use the semaphore *correctly* would now *fail to achieve mutual exclusion* since they change the semaphore between 2 and 1 instead of between 1 and 0. As such, if we would have use a semaphore instead of a lock to synchronize `decrease-error-2.c`, the program would most likely have *appeared to work* since failure to synchronize some data is often not immediately visible. However, since the mistake breaks mutual exclusion the program is likely to exhibit strange behavior at some point in the future, which will not be fun to debug. As such, even though the error checking that can be done by locks is far from perfect, the cases it *is* able to detect can save quite a lot of debugging work down the line.

As we have seen in the two previous examples, it is easy to make mistakes when working with locks (or semaphores). Synchronization of shared data is usually one of the most tricky parts to get right, especially since mistakes in

synchronization are not always directly visible in form of incorrect behavior or a crash. Because of this, it is often useful to structure the code so that the synchronization becomes trivial to verify just by reading the code, even if it complicates the program logic to some extent. For example, if we restructure the `decrease` implementation to remove the `return` statement in the middle of the function, we get the code in Listing 5.7 (`decrease-good.c`). Here, it is trivial to verify that we always acquire and release the lock, since the `lock_acquire` and `lock_release` is at the same level of indentation (i.e., one is not in an `if` statement or a loop), and no `return` statements are between them.

```
int counter = 10;
struct lock counter_lock;

void decrease(int with) {
    lock_acquire(&counter_lock);
    if (counter >= with)
        counter -= with;
    lock_release(&counter_lock);
}
```

Listing 5.7: A good way to structure the implementation of the `decrease` function to make it trivial to verify that the synchronization is correct.

In this particular case, one can argue that it makes it easier to read the function as well. This is not always the case, as this kind of rewrite would sometimes require us to add additional variables. At the end of the day, there is a trade-off between how easy it is to follow the program logic and how easy it is to verify that the synchronization is correct. It is usually a good idea to keep the synchronization logic much simpler than the program logic, since it is *much* easier to write unit-tests to verify that the program logic is correct compared to verifying that the synchronization is correct.

5.5 Exercises

1. The program `replace.c` contains a variable `data` that is shared between two threads. The program currently uses two semaphores, `a` and `b`, to make sure that the shared variable is accessed properly. One of these can be replaced with a lock. Find out which one, and replace it with a lock. You can verify your solution using Progvis.

Try to replace the other semaphore with a lock as well, and use Progvis to see why it is not possible to use a lock to replace that semaphore.

2. The program `bug-hunt.c` contains a function `update_counter` that modifies the variable `counter` according to the value of the parameter `delta`. The variable `counter` is protected using a semaphore that is used as a lock (`counter_lock`). The programmer who wrote the `update_counter` function added an “optimization” that checks if `delta` is zero and avoids synchronization in that case.

The program is, however, not working correctly. While the program *appears* to work correctly, the model checker in Progvis (*Run* \Rightarrow *Look for errors...*) reports that the program behaves incorrectly. Use Progvis to understand *why* the program is incorrect (even if you see *what* the error is). Then replace the semaphore `counter_lock` with a proper lock and see what error Progvis reports when you run the program (either the model checker or normally). Did the semaphore or the lock make it easier to find the error? What was the error?

Critical Sections

In Chapters 4 and 5 we have seen how both semaphores and locks can be used to avoid data races by enforcing *mutual exclusion*. That is, we make sure that at most one thread is allowed to execute code that uses some shared data and thereby avoid data races. The piece of code that at most one thread is allowed to execute is what we will call a *critical section*.

It turns out that it is usually not enough to just find all accesses to shared variables and synchronize them individually. In most cases we need to synchronize larger pieces of the code than just a single line. This is where the notion of a *critical section* is useful. It allows us to speak about sections in the code that need to be executed together as a unit, and that must not be interrupted by other code that accesses the same data.

6.1 Why Critical Sections?

To illustrate why it is useful to think in terms of critical sections we will use three versions of the program `increment` (i.e., `increment-1.c`, `increment-2.c`, and `increment-3.c`), which is similar to the program used to illustrate the problems with concurrent execution in Chapter 3. The central parts of the program `increment-1.c` is shown in Listing 6.1.

As can be seen in Listing 6.1, the program is centered around the global variable `result` that is protected with the lock `result_lock`. Two functions, `thread_a` and `thread_b`, both increment the variable 4 times in a loop by 2 and 5 respectively. The full program also contains a `main` function that initializes the lock, starts the threads, waits for them to terminate using a semaphore, and prints the value of `result`. The `main` function is not shown in the figure as it is not central to the discussion about critical sections.

```

const int times = 4;
int result = 0;
struct lock result_lock;

void thread_a(void) {
    for (int i = 0; i < times; i++) {
        lock_acquire(&result_lock);
        result += 2;
        lock_release(&result_lock);
    }
    sema_up(&done);
}

void thread_b(void) {
    for (int i = 0; i < times; i++) {
        lock_acquire(&result_lock);
        result += 5;
        lock_release(&result_lock);
    }
    sema_up(&done);
}

```

Listing 6.1: The important parts of the program `increment-1.c`.

The version `increment-1.c` currently behaves correctly. As we would expect, the variable `result` contains $(2+5) \cdot 4 = 28$ at the end of the program. However, remember that even though the statement `result += 2` looks like one operation, it is actually executed as (at least) three operations: load the value of `result` from memory, increment it by 2, and store the value back into `result`. We can illustrate this by explicitly rewriting `thread_a` as shown in Listing 6.2 (also in `increment-2.c`).

```

void thread_a(void) {
    for (int i = 0; i < times; i++) {
        lock_acquire(&result_lock);
        int tmp = result;
        tmp += 2;
        result = tmp;
        lock_release(&result_lock);
    }
    sema_up(&done);
}

```

Listing 6.2: Explicitly implementing the separate steps of `result += 2`.

This version of the code also works correctly. After all, the only change we made was to make the steps that the hardware needs to do anyway explicit in the source code. It does, however, highlight that only the first and the third steps actually access the shared variable `result`. Based on this observation, it looks like we only need to hold the lock around the lines that actually access `result`, as shown in Listing 6.3 (also in `increment-3.c`).

```

void thread_a(void) {
    for (int i = 0; i < times; i++) {
        lock_acquire(&result_lock);
        int tmp = result;
        lock_release(&result_lock);
        tmp += 2;
        lock_acquire(&result_lock);
        result = tmp;
        lock_release(&result_lock);
    }
    sema_up(&done);
}

```

Listing 6.3: Adjusting the locks to minimize the time where mutual exclusion is enforced.

The implementation in Listing 6.3 is interesting. Since the implementation makes sure to hold `result_lock` whenever `result` is accessed, the lock ensures that the two threads never access it concurrently, and therefore we have eliminated all possibilities of *data races* in the code. The program is therefore well-defined according to the C standard. However, the program is still not correct in the sense that it behaves the way we expect it to behave.

Practice: Load the program `increment-3.c` in Progvis, and use *Run* \Rightarrow *Look for errors...* to ask Progvis to find any errors. Does Progvis find any errors? Why/why not?

The end of the `main` function in `increment-3.c` contains a commented-out line that contains: `assert(result == times * (2 + 5));` Uncomment the line and select *Run* \Rightarrow *Look for errors...* again. What happens now?

As you will have noticed, Progvis will *not* find any errors in the code. This is because, as we noted earlier, the program neither crashes nor contains any data races. It is therefore a correct program according to the C standard. However, we as programmers have additional expectations about the behavior of the program. In particular, we expect that none of the additions should be “forgotten”, which is why we consider the program to be incorrect. If we let Progvis know about our expectations, for example by adding an `assert` statement at the end of the `main` function, it will find the error for us since a failed `assert` statement causes the program to crash.

This example illustrates an important point. It is usually not enough to simply surround each access to shared data with `lock_acquire` and `lock_release`. Doing so does indeed eliminate data races and thereby avoids undefined behavior. However, it is usually not enough to make the program behave as we expect it to. In most cases, the program updates shared data in multiple steps, and it is important that other threads are neither able to observe nor interfere with the shared data before all steps are completed.

This is exactly the problem in Listing 6.3. The code illustrates that the `+=` operator is really executed as (at least) three steps: (1) read from memory, (2) increment the value, and (3) write to memory. If another thread is able to

observe and/or interfere with the process before it is done, the result from step (3) in one thread might be overwritten by step (3) of another another thread that read an old version of the shared data in step (1). The solution, as we have seen, is to use locks to ensure that once one thread starts step (1), all other threads have to wait until the thread is done with step (3) before accessing the shared data. This ensures that at most one thread is working on any of steps (1)–(3), and thereby that no thread is able to observe or interfere during the steps. As such, we would think of these steps as being a single critical section.

6.2 Critical Sections and Conditionals

The previous section illustrated that shared data is often updated in multiple steps, even though it might not seem to be the case initially. As such, it is important to properly determine the critical section where shared data is updated so that we can protect it with locks. It is not uncommon for updates to shared data to include conditionals (e.g., `if` statements or `for` loops) as well. In this section we will therefore return to the `decrease` program from the previous chapter to illustrate how conditionals interact with critical sections. We start with the program `decrease-1.c`, which is the same as the program we synchronized at the end of the last chapter.

As can be seen in Listing 6.4, the program contains a global variable `counter` that is protected by the `counter_lock`. The function `decrease` decreases `counter` by a specified amount after checking that the counter is large enough to not become negative when decreased (not too dissimilar from `sema_down`). To ensure that `decrease` behaves correctly when called concurrently, the function acquires `counter_lock` at the start of the function and releases it at the end of the function.

```
int counter = 10;
struct lock counter_lock;

void decrease(int with) {
    lock_acquire(&counter_lock);
    if (counter >= with)
        counter -= with;
    lock_release(&counter_lock);
}
```

Listing 6.4: The `counter` variable and `decrease` function from `decrease-1.c`.

The program also contains code that starts two threads that both call `decrease(6)`. Since `counter` starts at 10 and we call `decrease(6)` two times, one of the calls (whichever happens to acquire the lock first) decreases the counter to 4 while the other does nothing. As such, `counter` contains 4 at the end of the program.

The crucial question that remains is, how do we arrive at the solution in Listing 6.4? The first observation is, of course, that the variable `counter` is shared and all accesses to it needs to be protected. The question that remains is how we arrive at the proper size of the critical section that we need to protect with locks. That is, do the two lines that access `counter` need to be in a

single large critical section, or is it enough to have two different critical sections? The key observation here is that the line `counter -= with` assumes that `counter >= with` (i.e., the condition in the `if` statement) is still true. As such, there is a dependency between the `if` statement and decrementing `counter`.

```
void decrease(int with) {
    lock_acquire(&counter_lock);
    bool ok = counter >= with;
    if (ok) counter -= with;
    lock_release(&counter_lock);
}
```

Listing 6.5: Rewritten `decrease` function as found in `decrease-2.c`.

We can clarify this dependency by rewriting the `decrease` function as in Listing 6.5 (`decrease-2.c`). Here, we store the result of the condition in the boolean variable `ok`, and use it on the next line. The `if`-statement on the next line is written on a single line to illustrate that we think of it as a single unit rather than altering the control flow of the program. That is, we think of the line as “decrease `counter` if `ok` is true”¹ rather than “if `ok` is false, skip the body of the `if` statement”. Since the variable `ok` is an “input” to the line, the dependency between the two lines is now quite apparent.

```
void decrease(int with) {
    lock_acquire(&counter_lock);
    bool ok = counter >= with;
    lock_release(&counter_lock);

    lock_acquire(&counter_lock);
    if (ok) counter -= with;
    lock_release(&counter_lock);
}
```

Listing 6.6: Rewritten `decrease` function as found in `decrease-3.c`.

This way of expressing `decrease` also makes it easier to see what happens if we separate the two accesses to `counter` into separate critical sections. To do this, we can simply release the lock and re-acquire it between the lines, as is done in Listing 6.6 (`decrease-3.c`).

Practice: Use the `Run ⇒ Look for errors...` option in Progviz to verify that the program `decrease-3.c` is incorrect, and why.

As you will have seen above, dividing the critical section into two causes the program to misbehave in a way that is similar to the behavior we saw in the previous chapter before we added synchronization. Even though they fail in

¹This is actually not far-fetched. Many CPUs have an instruction called *conditional move* that can be used to implement this line without altering the control flow. If you enable optimizations, it is likely that your compiler uses such an instruction if your CPU has one.

a similar way, there is a very important difference. The non-synchronized version is undefined according to the C language, which means that anything may happen.² The code in Listing 6.6 on the other hand is well-defined since it avoids data races.

The issue in Listing 6.6 is that the line in the first critical section (`ok = counter >= width`) observes the value of `counter` and makes a decision based on the observation. In this case, the decision is if `counter` is large enough so that it can be decreased without becoming negative. The line in the next critical section (`if (ok) counter -= with`) then uses the value in `ok` to act on the decision. Note that the program does *not* verify that `counter` is still large enough, but blindly trusts the value in `ok` and thereby assumes that all is well. For the program in Listing 6.6, this is *not* true since it releases the lock between the two statements. Once the lock is released, other threads that operate on `counter` (e.g., another thread that calls `decrease()`) are able to acquire the thread and modify `counter`. This is why the value of `ok` may be stale when the lock is re-acquired.

This illustrates that we once again have code that updates shared data in multiple steps. In this case, the first step is that we inspect shared data to make a decision, and then act on the decision to possibly update a shared variable (which we have previously seen is multiple steps). Since it is important that the data we used when making a decision is the same as when we act on the decision, the two need to be a part of the *same* critical section. Otherwise, other threads may change the data in a way that invalidates the decision.

Finally, it is worth noting that while the rewrite in Listing 6.5 makes the dependency between the decision and the update explicit through the `ok` variable, it does *not* affect the behavior of the program. It is possible to split the critical section into two parts if the code is written as in Listing 6.4 as well, but it is less clear. It would look like Listing 6.7.

```
void decrease(int with) {
    lock_acquire(&counter_lock);
    if (counter >= with) {
        lock_release(&counter_lock);

        lock_acquire(&counter_lock);
        counter -= with;
    }
    lock_release(&counter_lock);
}
```

Listing 6.7: Split critical section without introducing additional variables. Note that the lock is no longer acquired and released at the same level of indentation, which makes it harder to determine where critical sections start and end.

²In particular, even if the program works for one version of a compiler, it may stop working the day you receive a small update to your compiler, or even when you change some seemingly unrelated code.

6.3 Shared Data in Multiple Functions

A critical section is not necessarily tied to a single function. If multiple functions in the same program access the same shared data, they need to use the same lock to protect the shared data. As such, the critical sections in the functions need to use the same lock since they operate on the same shared data.

```

int counter = 10;
struct lock counter_lock;

void decrease(int with) {
    lock_acquire(&counter_lock);
    if (counter >= with)
        counter -= with;
    lock_release(&counter_lock);
}

void increase(int with) {
    counter += with;
}

```

Listing 6.8: Extension of the `decrease` program to also include an `increase` function.

To illustrate this idea, we add an `increase` function to the `decrease` program we used in the previous section. It simply increases `counter` with a given value as shown in Listing 6.8. The full program `updown-1.c` also includes a `main` function that starts two threads that each call `decrease(6)` as before. It additionally calls `increase(1)` from the main thread. As such, we expect the `counter` to contain 5 at the end of the program as indicated by the `assert` at the end of the `main` program.

Practice: The program `updown-1.c` is not correct. In particular it contains a data race, even though it uses `counter_lock` to protect the shared variable `counter`. Use Progvis to find the data race.

As you have likely already seen, the issue with the program `updown-1.c` (Listing 6.8) is that the `increase` function does *not* use `counter_lock` to protect `counter`. As such, it is possible for one thread to run `increase(1)` concurrently with another thread that runs `decrease(6)`. Even though the thread running `decrease(6)` acquires the lock, the thread running `increase(1)` is allowed to continue and modify `counter` as it wishes since it does not attempt to acquire `counter_lock`.

One way to view this problem is that the `increase` function violates an assumption in the code. In this case, we assume that all threads that modify `counter` have first successfully acquired the `counter_lock` lock. This assumption is *not* validated by the programming language, which is why we will not get any errors or warnings when we compile `updown-1.c`. The assumption is, however, important to uphold. As we saw in the previous section, the `decrease` function

utilizes the assumption to ensure that no other threads modify `counter` after it has decided to decrease `counter`, but before it has actually had time to do so.³ Additionally, we use the convention to ensure that we avoid data races when accessing `counter`. Since we fail to do so, the program currently contains a data race and is therefore undefined.

To address the issue, we need to modify the `increase` function to follow the convention of holding `counter_lock` whenever `counter` is accessed. By doing that, we properly protect `counter` with `counter_lock` throughout the program, thereby avoiding data races. In this case, we can simply modify the function as in Listing 6.8 (`updown-2.c`).

```
void increase(int with) {
    lock_acquire(&counter_lock);
    counter += with;
    lock_acquire(&counter_lock);
}
```

Listing 6.9: Updated version of the `increase` function to properly protect access to the `counter` variable.

Practice: The program `updown-3.c` (below) contains an alternative solution that uses two locks, one for `decrease` and one for `increase`. This does *not* work as intended, even though it may look like all code that uses `counter` is protected with locks. Use Progviz to find out why the program is incorrect.

```
int counter = 10;
struct lock decrease_lock;
struct lock increase_lock;

void decrease(int with) {
    lock_acquire(&decrease_lock);
    if (counter >= with)
        counter -= with;
    lock_release(&decrease_lock);
}

void increase(int with) {
    lock_acquire(&increase_lock);
    counter += with;
    lock_acquire(&increase_lock);
}
```

³In this particular example, increasing `counter` will never invalidate the decision made by `decrease`. However, we still have the issue that neither `-=` nor `+=` are executed as a single step, which may cause the program to behave incorrectly.

6.4 Synchronization Granularity

Up until now we have only looked at critical sections that involve a single shared variable. In real programs, critical sections are likely to involve multiple variables that together represent some state that needs to be kept consistent. As we shall see, there are different ways to synchronize such programs, each with different trade-offs. We will also see that the extent of critical sections may change based on requirements from other parts of the code. The end-goal of this section is therefore to show that it is important to consider the program as a whole when determining how it needs to be synchronized. In other words, the best way to synchronize a program often depends on how the program is structured at a larger scope than individual lines of code or even individual functions. Therefore it is important to look at the big picture before diving into details.

To illustrate the intricacies involved with synchronization of larger programs we will use different versions of the program `bank`. The first version, `bank-1.c`, is not synchronized at all and acts as the starting point for our exploration. The central parts of `bank-1.c` is shown in Listing 6.10.

```
struct account {
    int balance;
};

int accounts_count;
struct account *accounts;

bool transfer(int amount, int from, int to) {
    struct account *f = &accounts[from];
    struct account *t = &accounts[to];

    if (f->balance >= amount) {
        f->balance -= amount;
        t->balance += amount;
        return true;
    } else {
        return false;
    }
}
```

Listing 6.10: Initial, unsynchronized version of the `bank` program.

As indicated by the name, the program manages a simple fictional bank. The bank consists of a number of accounts, each represented by an instance of `struct account`. For the purposes of this example, we are only concerned about the current balance in each account (i.e., the amount of money currently in the account). All accounts are stored in the global array `accounts`. The `accounts` variable is defined as a pointer since the actual array is allocated dynamically by the `create_accounts` function (not in the figure). The variable `accounts_count` stores the number of elements in the `accounts` array. Since the array is allocated dynamically, it is possible to change the size of the array at a later point. For the time being we do, however, assume that the array is created

once and then never resized.

As shown in Listing 6.10, the program also contains a function named `transfer` that transfers money between two accounts within the bank. The function first checks that account `from` has a sufficient balance for the transfer.⁴ If the balance is large enough, it then subtracts `amount` money from the balance of account `from`, and adds the corresponding balance to account `to`.

In addition to the code in Listing 6.10, the program also contains a `main` function that creates two accounts with a balance of 10. The `main` function then starts two threads that both transfer 8 units of currency from account 0 to account 1. Finally, `main` waits for both threads to complete and asserts that the balance of account 0 is nonnegative.

Practice: The program `bank-1.c` is currently not correct. First and foremost, it contains data races since access to shared data is not protected. Additionally, it is possible for the balance of account 0 to become negative.

Find the problem in the code that causes these issues (either by reasoning about the code, or by using Progviz). Then, suggest one or more critical sections in the `transfer` function that needs to be protected in order to solve the issues.

As you most likely concluded from above, the issue is similar to the issue in the `decrease` program. The `transfer` function contains an if-statement that first verifies a condition (that the balance of account 0 is large enough) and then acts on it (decreasing the balance of account 0). However, since we have not protected access to `f->balance`, other threads might have invalidated the condition that we have verified. In the case of the program above, both threads may first observe that the balance of account 0 is larger than 8, and decide to continue transferring funds. As such, both threads will decrease the balance of account 0 with 8, so that it now contains the value -6, which is not as intended.

```
bool transfer(int amount, int from, int to) {
    struct account *f = &accounts[from];
    struct account *t = &accounts[to];

    bool ok = f->balance >= amount;
    if (ok)
        f->balance -= amount;
    if (ok)
        t->balance += amount;
    return ok;
}
```

Listing 6.11: Updated version of the `transfer` function to make the dependency explicit. Available as `bank-2.c`.

As before, we can rewrite the logic in the `transfer` function to make the fact that adjusting the balances of both accounts depend on the check of `f->balance`.

⁴The bank does not provide services like lending money.

Initially, checking `ok` twice might seem superfluous. This does, however, become relevant when we think more carefully about the extends of the critical sections in the function.

Practice: At this point we are ready to protect the shared data in the `accounts` array with a lock called `accounts_lock`. Below are three options of possible lock placements, `bank-3-a.c`, `bank-3-b.c`, and `bank-3-c.c`. Reason about the code or use Progviz to determine which of the options most accurately protects the critical section in the `transfer` function.

bank-3-a.c

```
bool transfer(int amount, int from, int to) {
    lock_acquire(&accounts_lock);
    struct account *f = &accounts[from];
    struct account *t = &accounts[to];
    lock_release(&accounts_lock);

    bool ok = f->balance >= amount;
    if (ok)
        f->balance -= amount;
    if (ok)
        t->balance += amount;
    return ok;
}
```

bank-3-b.c

```
bool transfer(int amount, int from, int to) {
    lock_acquire(&accounts_lock);
    struct account *f = &accounts[from];
    struct account *t = &accounts[to];

    bool ok = f->balance >= amount;
    if (ok)
        f->balance -= amount;
    if (ok)
        t->balance += amount;
    lock_release(&accounts_lock);
    return ok;
}
```

bank-3-c.c

```
bool transfer(int amount, int from, int to) {
    struct account *f = &accounts[from];
    struct account *t = &accounts[to];

    lock_acquire(&accounts_lock);
    bool ok = f->balance >= amount;
    if (ok)
        f->balance -= amount;
    if (ok)
        t->balance += amount;
    lock_release(&accounts_lock);
    return ok;
}
```

As you have likely realized, the program in `bank-3-a.c` is incorrect. Even though the data protected by the lock `accounts_lock` is in the `accounts` array, it is not enough to just protect the lines in the code that mention `accounts`. In this case, the lines protected in `bank-3-a.c` compute pointers to the data in `accounts`, and the code that actually modify shared data appears later on in the function. As we discussed above, the relevant part is the code that inspects and modifies the `balance` member of the individual accounts.

As such, we can conclude that `bank-3-b.c` works better. This is true, `bank-3-b.c` does indeed behave correctly. Since we acquire the lock before we touch the `accounts` variable or modify any other accounts, and release it after we have finished all modifications to the `balance` of all accounts, it is not possible for one thread that executes `transfer` to interrupt another thread that also executes `transfer`. As such, the issues we found in `bank-1.c` and `bank-2.c` can not happen.

The remaining question is, what about `bank-3-c.c`? The difference between `bank-3-b.c` and `bank-3-c.c` is that the latter does *not* include the lines below in the critical section that is protected by the lock:

```
struct account *f = &accounts[from];
struct account *t = &accounts[to];
```

Perhaps surprisingly, these lines do *not* need to be included in the critical section. To understand why, we need to think about what the lines above represents and what `account_lock` is actually protecting.

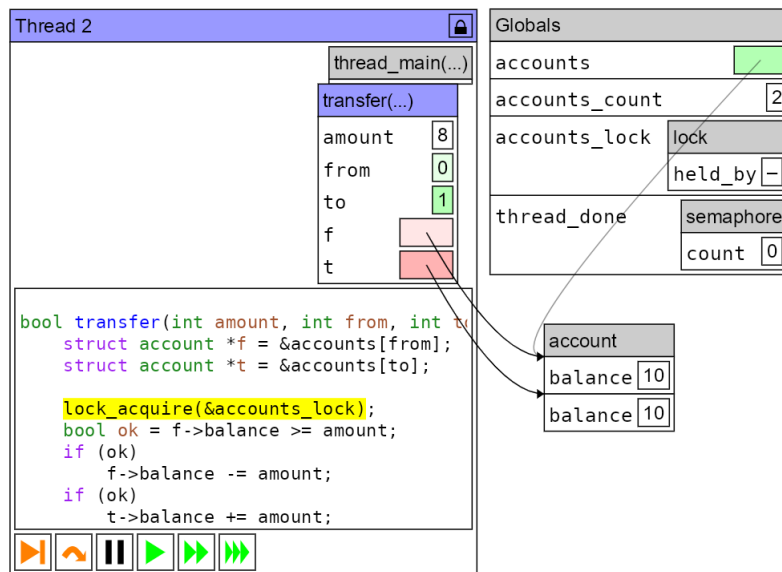


Figure 6.1: Progviz illustration of how the lines not protected in `bank-3-c.c` affects memory.

As a starting point, we can see how Progviz illustrates the memory accesses performed by the two lines. We can do this by opening the program `bank-3-c.c` and stepping the program until thread 2 is started and reaches the point in Fig. 6.1. As usual, Progviz colors memory locations the program has read

from in green, and locations the program has written to in red. The strong colors represent locations accessed by the last statement, and more faded ones represent locations accessed since the previous synchronization operation (e.g., starting threads, acquiring or releasing locks, etc.). In particular, note that `accounts` is colored green, both `t` and `f` are colored red, but the two accounts' `balance` is *not* colored at all. What this means is that the two lines *did not* access the balances⁵ of *any* of the accounts. Since the lock `account_lock` is used to protect the *contents* of the array, this is our first hint that `bank-3-c.c` is indeed correct.

To understand *why* the lines do not access the contents of the array, we need to consider what the lines above actually mean. Remember that arrays are both represented as and handled through pointers (see Section 2.6). As such, the syntax used above is just syntactic sugar for the following:

```

struct account *f = accounts + from;
struct account *t = accounts + to;
```

That is, each line reads the pointer stored in the variable `accounts` and increments it with the value that is stored inside `from` and `to` multiplied with the size of `struct account`. Remember that the underlying representation of a pointer is just an integer value, the special semantics of pointers is something that C adds to help us remember what type of data the pointer refers to. From the rewrite above, it hopefully becomes clear why these lines do not access the contents of the `accounts` array. It is, however, important to remember that the code still reads from the `accounts` variable itself (i.e., the pointer, not the contents of the array). As such, the reason we do not have to include these lines in the critical section is that we assume that `accounts` is initialized once and then never changes.

It is worth noting that this phenomenon is not unique to arrays. The same thing happens when we use the address of operator (`&`) to compute the address of other things as well. For example, if we wished to compute the pointer to the `balance` variable in the `from` account, we could write: `int *f_balance = &f->balance;`. Just as when we computed the address to the `from` account, this line only needs to read the variable `f`. It does not touch the `balance` in the account that `f` points to, since we requested the address of the variable rather than the integer stored there. As you can see, correctly reasoning about what parts of a function that needs to be included in a critical section requires a good understanding of the semantics of the programming language we use. This is one reason why Progviz emphasizes an accurate and detailed representation of the memory and how the code interacts with the memory!

The conclusion is that the program `bank-3-c.c` is correct since the `account_lock` only protects the *contents* of all accounts, and the lines that compute the address to the `from` and `to` accounts do not access any data of any account. Furthermore, the program `bank-3-c.c` is slightly better than `bank-3-b.c`, since it does not disallow two threads from computing the address to their respective accounts concurrently. Even though the difference is negligible in this case, it is good practice to find the minimal critical section and only protect that using a lock. When learning concurrency and synchronization, this mindset challenges

⁵Or any other part for that matter.

us to refine our understanding of the program we synchronize and learn more nuances.

Furthermore, in “real” multithreaded programs, the main bottleneck is often the critical sections. Since only one thread at a time is allowed to execute code in a critical section, the performance of critical sections will *not* scale with the number of available CPU cores. As such, there are tangible benefits to structuring programs so that critical sections can be minimized or avoided altogether.

6.4.1 Improved Parallelism

Our current implementation of the bank program (e.g., `bank-3-c.c`) uses a single lock to protect all accounts, and thereby only allows *a single* thread to transfer money between accounts. This is overly restrictive, since there would be no issues allowing one thread to transfer money from account 0 to account 1 while another thread transfers money from account 1 to account 2.

Practice: The file `bank-4.c` contains a copy of `bank-3-c.c`, but with a different main program. This program starts two threads as described above. One thread transfers money from account 0 to account 1, while the other thread transfers money from account 2 to account 3. Open the program in Progviz and observe the behavior of the program:

1. The way the program is currently written, note that the two threads need to wait for each other since both acquire `accounts_lock` inside the `transfer` function.
2. What happens if we remove the lock? From before, we know that this causes issues when two or more threads operate on the *same* accounts concurrently. However, the main program here always operates on *different* accounts. Use `Run ⇒ Look for errors...` to see if the removal of the lock causes any issues in this particular case.

As you have likely noticed from above, using Progviz we can see the problem in `bank-4.c`. In Fig. 6.2, we can see that the thread labeled Thread 2 has acquired `accounts_lock` and has started to transfer money from account 0 to account 1. We can also see that the thread labeled Thread 3 also tried to acquire `accounts_lock`, and therefore has to wait until Thread 2 is done. This happens even though Thread 3 is trying to transfer money between different accounts (from 2 to 3, as shown by the arrows).

The reason for this is that we have been too coarse when determining our critical sections. In Section 6.4, we treated the `accounts` array as one large unit. This was initially intuitive since the data we needed to protect was stored in that variable. However, as we saw above, this gives us an implementation that seems to impose stricter restrictions than the program actually needs to fulfill its requirements. As we shall see, the synchronization will depend on our requirements of the system as a whole. Therefore it is a good idea to start by clarifying our requirements of the system. So far, we only need to focus on the `transfer` function (`create_accounts` is only called once to initialize everything). Our requirements are that `transfer` should...

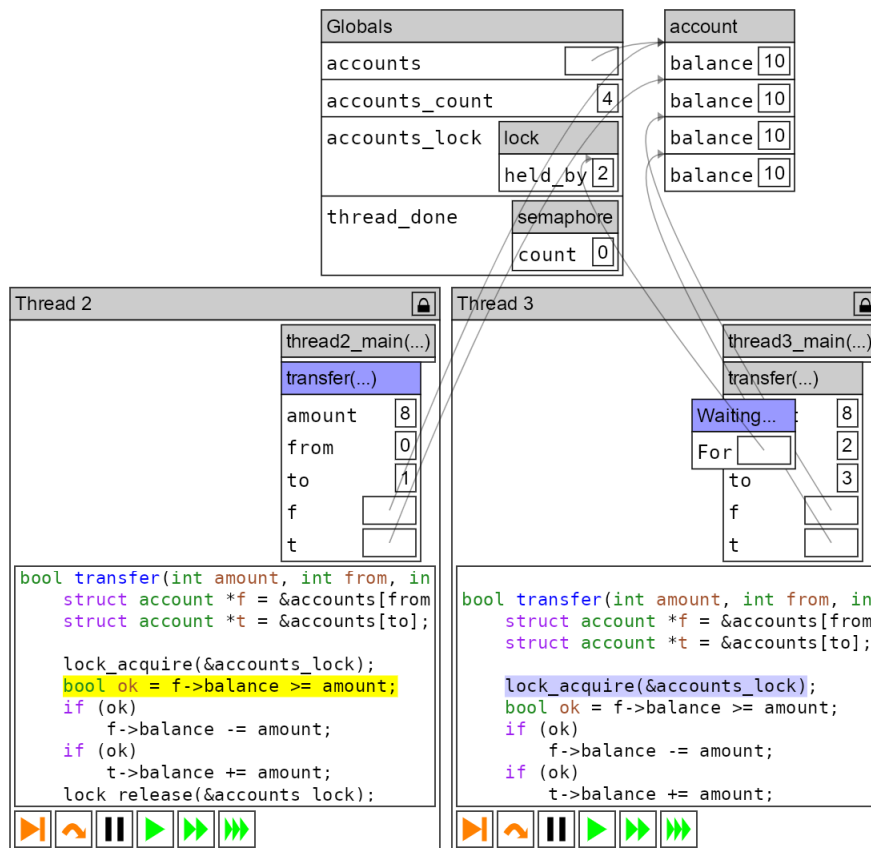


Figure 6.2: Even though the two threads are trying to operate on different accounts, the implementation in `bank-4.c` requires the threads to wait for each other.

1. ...ensure that the balance of the `from` account remains non-negative by refusing to transfer money if that would be the case.
2. ...ensure that a transfer from some account should not be overwritten by another transfer from the same account.
3. ...ensure that a transfer to some account should not be overwritten by another transfer to the same account.

From our analysis before we noted that requirements 1 and 2 above are linked, since the check for a sufficient balance needs to be performed in the same critical section as the removal of the funds. Furthermore, we can observe that requirements 1 and 2 are related to the `from` account, while requirement 3 is related to the `to` account. Since it is always acceptable to *add* money to an account in our system (there is no maximum balance, for example), this requirement can be fulfilled separately from the other two.

Since none of the requirements speak about more than one account at a time, we can conclude that there is no need for us to guarantee a consistent

state for the *entire array* of account. Rather, it is enough for us to treat *individual accounts* separately, and ensure consistency for individual accounts in isolation. Since the program `bank-4.c` uses a single lock to protect the entire account, we ensure that only one thread accesses the entire array at any one time, and therefore we can guarantee a consistent state for the entire array at the cost of not being able to perform multiple transfers in parallel. However, we have concluded that we do not need that property, and we would instead like to trade some consistency with increased parallelism.

To achieve our goal, we need to re-consider our lock placement. A good starting point is to declare the lock together with the data that should be protected. In this case, our goal is to protect the `balance` of individual accounts. Therefore, we remove `accounts_lock` and instead add a lock inside `struct account` that we call `balance_lock` to remind ourselves that it protects `balance`, as shown below:

```
struct account {
    int balance;
    struct lock balance_lock;
};

int accounts_count;
struct account *accounts;
```

Now that we have determined at what level we wish to synchronize the data (i.e., one lock for each account rather than one lock for *all* accounts), we need to re-visit the implementation of `transfer` to find the critical sections that need to be protected. To do this, we will start over from the implementation in `bank-2.c` below:

```
1 bool transfer(int amount, int from, int to) {
2     struct account *f = &accounts[from];
3     struct account *t = &accounts[to];
4
5     bool ok = f->balance >= amount;
6     if (ok)
7         f->balance -= amount;
8     if (ok)
9         t->balance += amount;
10    return ok;
11 }
```

We previously concluded that lines 2 and 3 do not need to be a part of any critical section since we are just computing the address of the accounts that will be used in the remainder of the function. Furthermore, lines 5–7 need to be a part of the same critical section, since line 5 checks that the balance of the `from` account is sufficient and line 7 acts on that decision. These lines address requirements 1 and 2, and are all related to the balance in the `from` account. As such, this critical section needs to be protected by the `balance_lock` in the `from` account.

The only other part of the function that accesses shared data is line 9. Line 9 needs to be protected by a lock to fulfill requirement 3. However, even though line 9 is only executed if `ok` is `true`, it is not important that line 9 is a part of the *same* critical section as the modifications to the `from` account. This

is because while we only wish to add money to the `to` account if we removed money from the `from` account, adding money to the `to` account can always be done and does not require that the state of either account remains in the state it was previously. As such, line 9 is a separate critical section that needs to be protected by the `to` account's `balance_lock`.

Now that we have determined the critical sections and what locks need to be used to protect them, we can simply add calls to `lock_acquire` and `lock_release` to protect the critical section. This ends up looking like the code in `bank-5.c`, also depicted in Listing 6.12.

```
bool transfer(int amount, int from, int to) {
    struct account *f = &accounts[from];
    struct account *t = &accounts[to];

    lock_acquire(&f->balance_lock);
    bool ok = f->balance >= amount;
    if (ok)
        f->balance -= amount;
    lock_release(&f->balance_lock);

    if (ok) {
        lock_acquire(&t->balance_lock);
        t->balance += amount;
        lock_release(&t->balance_lock);
    }
    return ok;
}
```

Listing 6.12: The implementation of `transfer` in `bank-5.c`.

Practice: Use Progviz to verify that the proposed synchronization in `bank-5.c` is indeed correct. The main program in `bank-5.c` is a combination of the two previous ones. It creates three threads in addition to the main thread. Thread 2 transfers money from account 0 to account 1, Thread 3 transfers money from account 0 to account 2, and Thread 4 transfers money from account 2 to account 3. In particular, verify that:

- Transfers between different pairs of accounts can occur without having to wait for each other (i.e., Thread 2 never has to wait for Thread 4, but Thread 3 may cause both to wait).
- Transfers do not cause concurrency errors (e.g., by using *Run* \Rightarrow *Look for errors...*).

From above (and Listing 6.12) we can make an additional important observation. Since both the `from` and `to` parts of the code uses `balance_lock` (for different accounts), it is possible that the critical section for the `from` account interferes with the critical section for the `to` account in another thread. For example, when Thread 4 acquires the lock before reading the balance of the

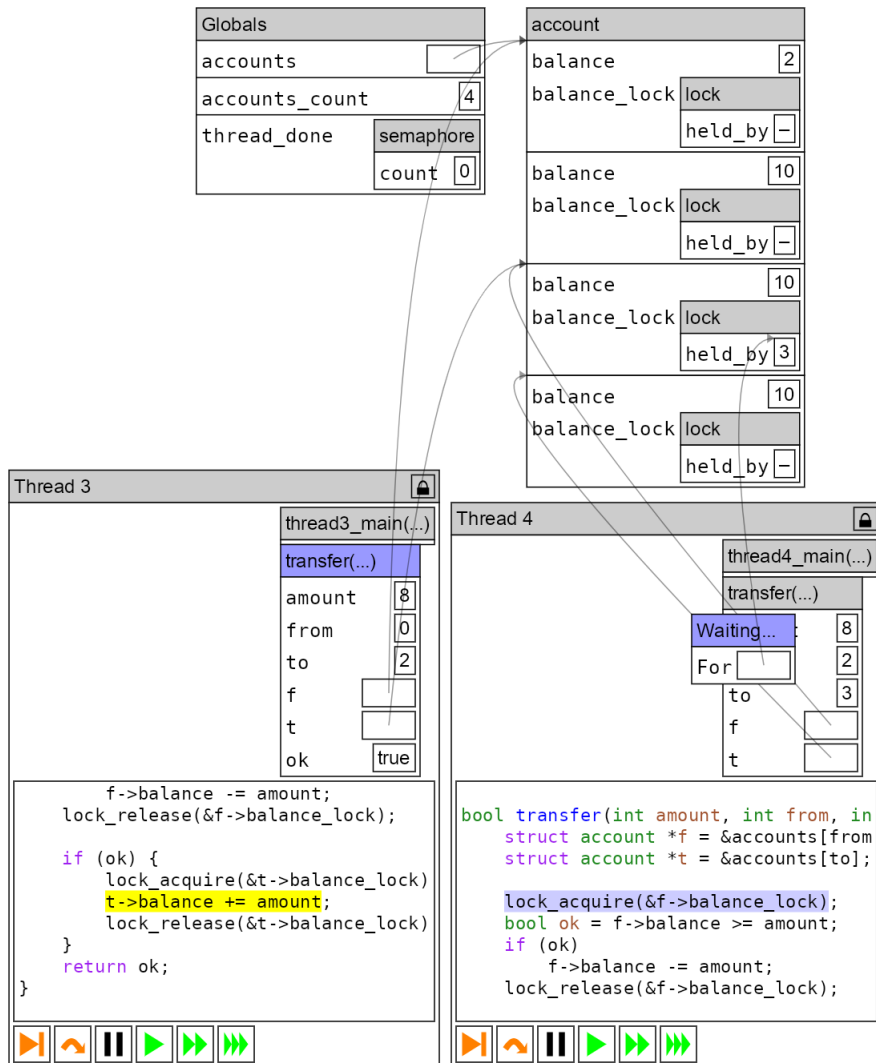


Figure 6.3: Using the same lock for different critical sections means that a thread executing one part of the code blocks another thread that wishes to execute another part of the code.

`from` account, it may need to wait for Thread 3 to update its `to` account as depicted in Fig. 6.3.

Practice: We can avoid one critical section blocking another critical section by using two locks for each account, as is done in `bank-5-incorrect.c` (depicted below). As the name of the program implies this is, however, not correct. Why? You can use *Run* \Rightarrow *Look for errors...* in Progviz to find an example of the issue.

```

struct account {
    int balance;
    struct lock balance_from_lock;
    struct lock balance_to_lock;
};

int accounts_count;
struct account *accounts;

bool transfer(int amount, int from, int to) {
    struct account *f = &accounts[from];
    struct account *t = &accounts[to];

    lock_acquire(&f->balance_from_lock);
    bool ok = f->balance >= amount;
    if (ok)
        f->balance -= amount;
    lock_release(&f->balance_from_lock);

    if (ok) {
        lock_acquire(&t->balance_to_lock);
        t->balance += amount;
        lock_release(&t->balance_to_lock);
    }
    return ok;
}

```

As you have likely noted from above, the behavior illustrated in Fig. 6.3 is exactly what we want. Since both critical sections may access `balance` in the same account, we want to protect the shared data to avoid data races. Otherwise it is possible for one thread to add funds to an account at the same time as another thread removes funds. Just as we have seen when two or more threads execute `+=` concurrently, it is possible that one of the operations will be overwritten by another thread and we thereby either create or destroy money from the bank.

This example illustrates a good rule of thumb when synchronizing code: it is usually a good idea to place the synchronization close to the variable or variables that need to be protected. As we have seen from this example, this means that we sometimes need to think about at what granularity we operate on data. For example, we found it overly restrictive to treat the entire array as a big unit, and ended up protecting accounts individually. This is, however, a rule of thumb. Sometimes it is desirable to lump data together and synchronize them as a bigger unit. One way to find such cases is to start by synchronizing individual variables and then pay attention to which locks always need to be acquired together. Locks that always need to be acquired

together can then be merged into a single lock that protects multiple variables without any performance loss.

6.4.2 Extending the Functionality of the Bank

As we noted in the previous section, the extent of the critical sections we need to protect in the program depends on the requirements of the program *as a whole*. As such, if the requirements for the program changes, we might need to revisit parts of the code we have synchronized earlier.

To illustrate this, the program `bank-sum-1.c` adds a new function `accounts_total` to the bank program. The function computes the total balance of all accounts in the bank. The initial version of the implementation is not synchronized at all, as shown in Listing 6.13.

```
int accounts_total(void) {
    int total = 0;
    for (int i = 0; i < accounts_count; i++) {
        total += accounts[i].balance;
    }
    return total;
}
```

Listing 6.13: The initial version of the function `accounts_total`.

Since the function accesses the `balance` of all accounts in the bank, we need to protect the access with the `balance_lock` as before. Since the function only reads each account's `balance` once in sequence, it appears to be enough to treat the body of the loop as a critical section that should be protected by the current account's `balance_lock`. As such, we add `lock_acquire` and `lock_release` to the function as shown in Listing 6.14 and `bank-sum-2.c`.

```
int accounts_total(void) {
    int total = 0;
    for (int i = 0; i < accounts_count; i++) {
        lock_acquire(&accounts[i].balance_lock);
        total += accounts[i].balance;
        lock_release(&accounts[i].balance_lock);
    }
    return total;
}
```

Listing 6.14: First attempt at synchronizing the `accounts_total` function in `bank-sum-2.c`.

This approach does indeed make sure that no data races are possible. Since we have made sure to hold `balance_lock` every time we access the `balance` of an account, the lock ensures that two threads are not able to access `balance` of the same account at the same time. This is good, since it makes our program well-defined according to the C standard. However, as we have seen before, all well-defined programs are not necessarily deterministic, nor do they always behave as we want them to. This is the case for `bank-sum-2.c`.

Practice: The main program in `bank-sum-2.c` illustrates a situation where the behavior of the program differs from our expectations. Since we only allow transferring money within the bank, we would expect `accounts_total` to always return the same value once the bank is initialized.

To verify this, the main program in `bank-sum-2.c` initializes a bank with 2 accounts with a balance of 10. It then creates two threads that transfer 5 money from account 0 to account 1. While the threads running, the main thread calls `accounts_total` and verifies that the total is 20.

Use Progvis to find a situation where `accounts_total` does *not* return 20. You can use *Run* \Rightarrow *Look for errors...* to do this. As the program is written, `accounts_total` may return either 10, 15, 20, 25, or 30. Can you find interleavings that cause this to happen? (Hint, with a slight modification to the `assert` statement, you can use *Run* \Rightarrow *Look for errors...* to find the interleavings for you!)

As you have likely noticed from above, there are two main problems with the program `bank-sum-2`. They are illustrated in Figs. 6.4 and 6.5. Note, that they only illustrate the key insights in the interleavings that cause `accounts_total` to return 10 and 30. The interleavings where `accounts_total` returns 15 and 25 are simply less extreme versions of the interleavings shown here.

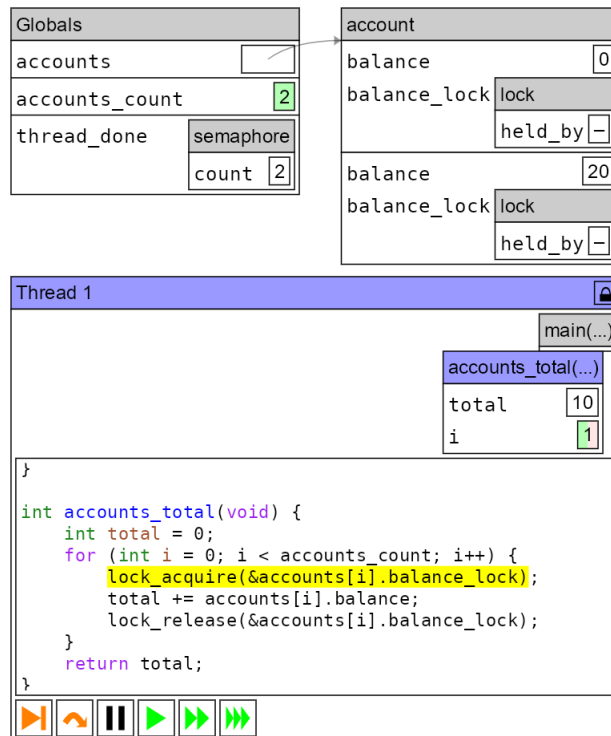


Figure 6.4: The first main issue in `bank-sum-2.c`.

The first issue is illustrated in Fig. 6.4. Here, Thread 1 executed the first iteration of the loop inside `accounts_total` before Thread 2 and 3 transferred

any money from account 0. As such, `total` was increased to 10 and Thread 1 continued to its second iteration. Before Thread 1 acquired the lock for account 1, the two other threads finished their transactions and terminated. This is the state shown in the figure. As we can see, `total` in Thread 1 is 10, but the 10 money from account 0 has already been transferred to account 1 (i.e., the balance of account 0 is 0 and the balance of account 1 is 20). As such, when Thread 1 continues, `total` will be increased to 30 and then returned.

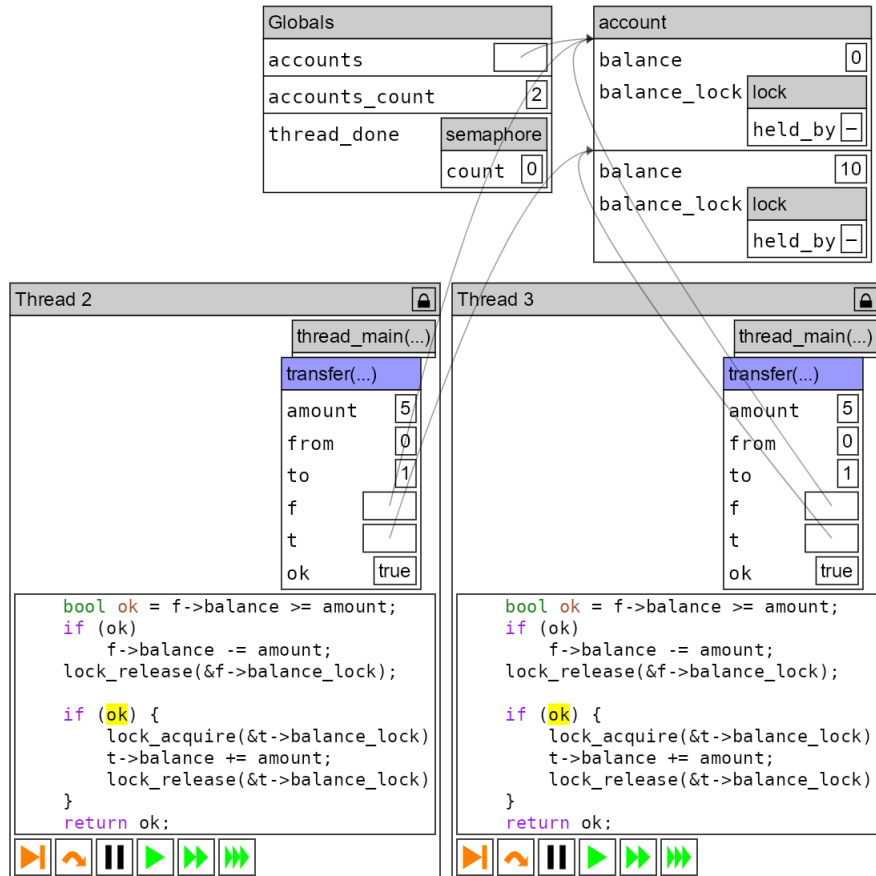


Figure 6.5: The second main issue in `bank-sum-2.c`.

The second issue is illustrated in Fig. 6.4. Here, Thread 2 and 3 have both started to transfer money. As we can see in the figure, both threads are partway through the transfer. They have both deducted 5 from the balance of account 0, but they have not yet increased the balance of account 1. If Thread 1 (not in the figure) were to execute `accounts_total` in its entirety at this point it would return 10 since the total balance of the two accounts is $0 + 10$ at this point in time.

This shows, again, that even though the program is free from data races, `accounts_total` may produce surprising results when used in concurrent programs. If we would have used `transfer` and `accounts_total` in a sequential program (i.e., a program with only one thread), then `accounts_total` would

always return the same value since a transfer would never occur concurrently with the summation.

At this point we need to take a step back and define our expectations of the `accounts_total` function in terms of requirements on the program. Most importantly, we need to determine whether we need `accounts_total` to provide the same guarantees as it would have in a sequential program, or if inaccurate results are acceptable. This very much depends on the context in which the program or the data structure is being used. In this case, we decide that it *is* important that `accounts_total` always gives an accurate result. Both in order to avoid surprising the developers that will use the functions down the line, but also since we are working with money, and losing money due to race conditions feels like a very bad idea. As such, we add a requirement to our previous list of requirements. As such, our requirements now look like this:

1. The `transfer` function should ensure that the balance of the `from` account remains non-negative by refusing to transfer money if that would be the case.
2. The `transfer` function should ensure that a transfer from some account should not be overwritten by another transfer from the same account.
3. The `transfer` function should ensure that a transfer to some account should not be overwritten by another transfer to the same account.
4. The `accounts_total` function should behave as if no transfers are in progress while it is being called. That is, when money is being transferred, it should always include money “in flight” as well as counting newly transferred money twice or not at all.

The new requirement (number 4) implies that `accounts_total` needs stronger guarantees regarding the consistency of multiple accounts. As we saw, the first problem arose because there are places in the code where `accounts_total` holds *no* lock. This means that other threads are able to interrupt it and shuffle money around in the bank (we saw that money “appeared” by transferring from low to high account numbers, but money can also “disappear” by transferring from high to low account numbers). Similarly, the second problem arose because `bank_transfer` releases the first lock before it acquires the second lock. This means that `accounts_total` is free to observe a state where money has “disappeared” since some thread is in the middle of a transfer.

Since we are trying to behave *as if* the program is sequential, the easiest way to achieve that is of course to go back to using a single lock for all accounts, as was the case in `bank-2-c.c`. This makes it easy to ensure that `accounts_total` does not execute during a transfer simply by acquiring the global lock. This solution is implemented in the program `bank-sum-slow.c` and depicted in Listing 6.15.

While this is a valid solution to the problem, we are back to the situation where we do not allow transfers between separate pairs of accounts concurrently. As such, we try to fix the issues in `bank-sum-2.c` instead. We start by addressing the first issue we found, that `accounts_total` can be interrupted by other threads that call `transfer` and thereby accidentally count the same

```

bool transfer(int amount, int from, int to) {
    struct account *f = &accounts[from];
    struct account *t = &accounts[to];

    lock_acquire(&accounts_lock);
    bool ok = f->balance >= amount;
    if (ok)
        f->balance -= amount;
    if (ok)
        t->balance += amount;
    lock_release(&accounts_lock);
    return ok;
}

int accounts_total(void) {
    int total = 0;
    lock_acquire(&accounts_lock);
    for (int i = 0; i < accounts_count; i++) {
        total += accounts[i].balance;
    }
    lock_release(&accounts_lock);
    return total;
}

```

Listing 6.15: Solution to the problems with `accounts_total` using a single lock in `bank-sum-slow.c`.

money more than once. In other words, `accounts_total` needs to see a consistent state of the balance in *all* accounts. Therefore, we can solve the problem by having it acquire the `balance_lock` for all accounts, and then release each lock after it has read the `balance`. This is implemented in `bank-sum-3.c` and depicted in Listing 6.16.

```

int accounts_total(void) {
    int total = 0;

    for (int i = 0; i < accounts_count; i++)
        lock_acquire(&accounts[i].balance_lock);

    for (int i = 0; i < accounts_count; i++) {
        total += accounts[i].balance;
        lock_release(&accounts[i].balance_lock);
    }

    return total;
}

```

Listing 6.16: Solving the first problem by acquiring *all* accounts' locks in `bank-sum-3.c`.

Practice: Use Progviz to verify that the program `bank-sum-3.c` solves the first problem. This means that it is no longer possible for `accounts_total` to return a value that is too large (i.e., only 10, 15, and 20 are possible). As before, you can use *Run* \Rightarrow *Look for errors...* to help you by modifying the assert statement.

As we have seen, the code in `bank-sum-3.c` does not solve the second problem, namely that `transfer` releases `balance_lock` for the `from` account before acquiring the `balance_lock` for the `to` account. As such, `accounts_total` is allowed to sum all accounts after `transfer` has removed money from the `from` account, but before it has returned money to the `to` account.

To avoid this, we simply need change `transfer` so that it acquires both locks at the same time. What happened here is that since `accounts_total` needs to be able to have a consistent view of all accounts, it changes the critical sections in `transfer`. Before we introduced `accounts_total` we concluded that `transfer` could treat the `from` and `to` accounts as separate entities. That was correct at that time. However, since `accounts_total` needs to be able to see the state of all accounts at they are either before or after a `transfer`, it means that `transfer` needs to hold the locks from when it starts accessing the first account until it is done with the second account. Essentially, the additional requirements introduced by `accounts_total` made the critical section in `transfer` larger. This again highlights that we need to consider the system *as a whole*, and we can not only focus on individual functions in isolation! Luckily, if we utilize abstraction to modularize our programs, it is usually sufficient to consider individual modules at a time. This is covered in more detail in Chapter 9.

Based on these observations, we can now solve the second problem by moving the calls to `lock_acquire` and `lock_release` as depicted in Listing 6.17. The final program is available as `bank-sum-4.c`.

```
bool transfer(int amount, int from, int to) {
    struct account *f = &accounts[from];
    struct account *t = &accounts[to];

    lock_acquire(&f->balance_lock);
    lock_acquire(&t->balance_lock);

    bool ok = f->balance >= amount;
    if (ok)
        f->balance -= amount;
    if (ok)
        t->balance += amount;

    lock_release(&t->balance_lock);
    lock_release(&f->balance_lock);
    return ok;
}
```

Listing 6.17: Solving the second problem by moving `lock_acquire` and `lock_release` as done in `bank-sum-4.c`.

At this point it is useful to take a step back and compare `bank-sum-4.c` to `bank-sum-slow.c`. The program `bank-sum-slow.c` does indeed look quite a bit simpler than `bank-sum-4.c`. However, `bank-sum-4.c` allows transfers between different pairs of accounts to occur concurrently. Note that neither of the two solutions allow two threads to run `accounts_total` concurrently: `bank-sum-4.c` acquires the locks from all accounts, while `bank-sum-slow.c` acquires the global lock. As such, which of the two is the best depends on how the code is used. If the program using the code calls `transfer` from multiple threads very frequently, but only calls `accounts_total` rarely, then the extra complexity in `bank-sum-4.c` is probably worthwhile. However, if `accounts_total` is called about as frequently as `transfer`, then the additional complexity in `bank-sum-4.c` is probably not worth it. Both regarding to the additional memory and CPU overhead involved in managing a larger set of locks, but also with regards to the complexity of the code. This is another example of why it is relevant to consider the system as a whole when working with concurrency, even though this particular aspect does not directly affect the correctness of the program.

Finally, it is worth noting that `bank-sum-4.c` *still* contains problems. We will look closer at these in the next chapter.

6.5 Working with Critical Sections

This section has been dedicated to the concept of critical sections. Determining what code needs to be a part of a critical section, and which lock should protect the critical section is not easy. However, as a part of illustrating the reasoning behind critical sections, we have also seen a strategy for finding and protecting the critical sections throughout the chapter. While it is not extremely formal and requires a good amount of creative thinking, it can be summarized as follows:

1. Start by formulating your expectations of the program or module that you need to synchronize as requirements.
2. Find variables that are shared between threads. For all variables that are modified outside of the initialization code, create a lock that protects the variable and add calls to `lock_acquire` and `lock_release` to protect accesses to the variable.
3. Use the requirements from step (1) and try to find valid uses of your program or module that behave incorrectly. This is probably the hardest step, and requires the ability to scrutinize one's own work, even though it is believed to be correct.
4. Based on the outcomes of step (3), either refine the synchronization (e.g., enlarge critical sections, merge different locks into one, etc.) or refine the requirements.
5. Repeat steps (3) and (4) until no further problems are found.

Of course, once you get more familiar with concurrent programming you can likely skip writing the first version of the program from step (2), and start with code that already addresses the obvious issues (e.g., starting not synchronizing

the if-statement and the decrement of the balance as two separate steps initially). However, the main takeaway from the steps above still remains. That is, write a solution, *critically analyze your solution*, and refine it accordingly. As mentioned above, trying to find problems in a solution you believe is correct is probably the hardest step of them all. Hopefully, this book and the help of the tools in Progvis will help you get in the right mindset at least!

6.6 Exercises

1. The program `seats1.c` contains a data structure `struct seat` that keeps tracks of whether or not a particular seat in a train is reserved. The function `reserve_seat` can be used to reserve the next free seat. The aim is that it should be possible to call `reserve_seat` from multiple threads, for example as done by `main` and `thread_main`.

Use the locks already defined in `struct seat` to make `reserve_seat` behave correctly. You can use Progvis to verify that your solution works correctly.

2. The program `seats2.c` contains an extension to the program in the previous problem. Instead of `reserve_seat`, the program now has a function `reserve_range` that reserves a range of seats (e.g., seats 0–2 to reserve seat 0 and seat 1). The aim is that the function shall first verify if all seats are currently free, and in that case all seats should be reserved. Importantly, this should work if `reserve_range` is executed concurrently, which is not the case currently.

Use the locks already defined in `struct seat` to make `reserve_range` behave correctly. You can use Progvis to verify that your solution works correctly.

3. The program `pack.c` contains data structures to keep track of the inventory of a small mail-order business. `struct product` stores all products that are available, and the variable `boxes_available` stores the number of boxes that are available to ship products. The function `buy` is called when a customer wishes to buy a product. It checks that both the product itself and a box are in stock, and reduces the availability accordingly. It also contains a function `additional_boxes`, that computes how many additional boxes are needed to ship all available products.

The program is written to support multiple threads accessing the data. The function `buy` behaves correctly in isolation (i.e., when only `buy` is called). However, when `buy` is used together with `additional_boxes`, the latter returns an incorrect number in some cases, as shown by the main program.

Revise the synchronization of `buy` so that `additional_boxes` works as intended. You can use Progvis to verify your solution.

Deadlocks

In the previous chapter we analyzed and synchronized a program that manages accounts in a fictional bank. At the end of the chapter, the program had two functions `transfer` and `accounts_total` that manipulated accounts with seemingly proper synchronization. Indeed, the final version of the program was free from data races and behaved according to the synchronization. However, it has the potential to *deadlock*.

Practice: The program `bank-1.c` contains the code of the final version of `bank-sum` from the previous chapter, but with a different main program. In `bank-1.c`, the main program spawns one thread that transfers money from account 1 to account 0, while the main thread itself transfers money from account 0 to account 1.

Use Progvis to find the problem in the program. Progvis will tell you when the program is in a deadlock. You can use *Run ⇒ Look for errors...* to help you. Think about *why* the program ended up in a deadlock, and why the program is not able to continue.

If you open `bank-1.c` in Progvis and select *Run ⇒ Look for errors...*, it will quickly state that it found a *deadlock* in the code. In visualization that shows the problem ends up in the state shown in Fig. 7.1. First and foremost we can see that both threads are waiting for something. Progvis shows this both by highlighting the relevant line in the source code blue, and by adding the box labeled *Waiting...* in the call stack. If we follow the arrow in the *Waiting...* box from Thread 1, we can see that it is waiting for `balance_lock` in account number 1 to be released, which is held by Thread 2. If we follow the arrow in the *Waiting...* box in Thread 2, we can see that it is waiting for `balance_lock` in account 0 to be released. However, that lock is currently held by Thread 1, which as we saw before is waiting for Thread 2 to release its lock.

This is a typical illustration of a deadlock: two or more threads wait for each other in a cycle. Since they are all waiting, none of them are able to continue executing code and thereby release any of the locks. As such, the typical symptoms of a deadlock is that some part of the program simply stops. Since deadlocks do not cause programs to crash, they are not always easy to find and debug. Similarly to other parts of concurrent programming, it

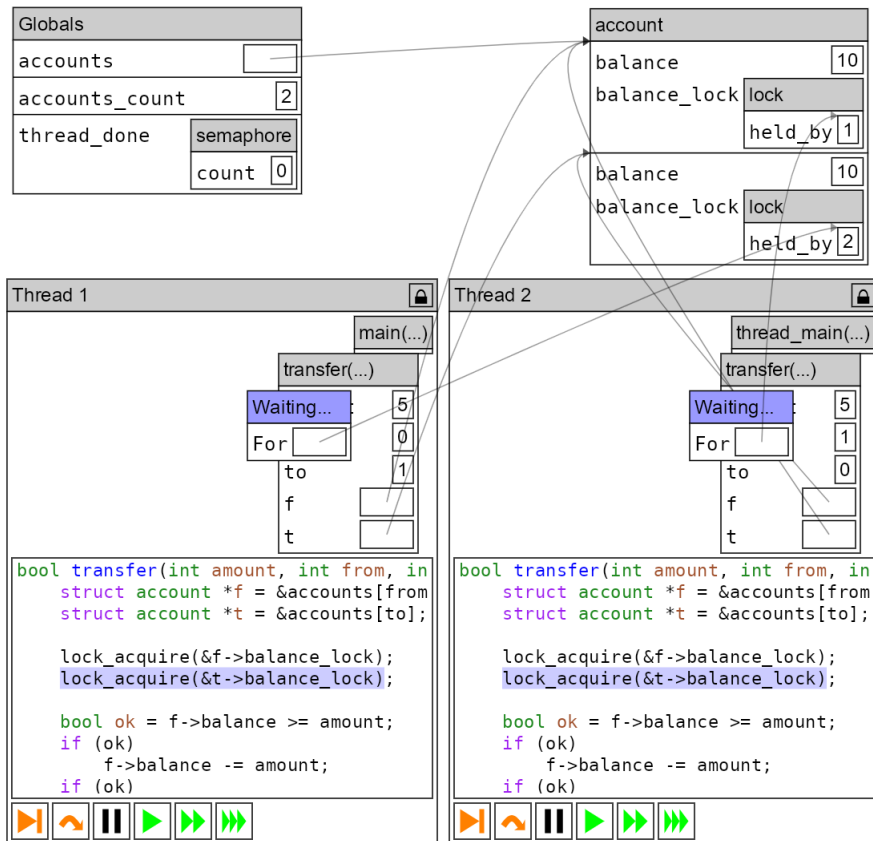


Figure 7.1: Deadlock found in `bank-1.c` by Progis.

is therefore highly beneficial to address the issue up-front, when designing programs. Luckily, there are several approaches to ensure that deadlocks are not possible.

Before we dive deeper into deadlocks, it is perhaps relevant to mention how Progis detects and reports deadlocks. Whenever program execution is paused (i.e., when Progis shows the current state of the program), Progis checks if all threads are waiting for something. If they are, Progis reports that the program is in a deadlock. This means that even if a deadlock has occurred in some part of the program, it will not be reported until all other threads have terminated. This is not a problem when using *Run ⇒ Look for errors...*, since Progis will find the issue eventually, but it is worth noting when stepping the program manually. Another implication of this behavior is that Progis will report a single thread calling `sema_down` on a semaphore initialized to zero as a deadlock. While this is of course an error, it is not typically considered a deadlock in the literature.

7.1 What is a Deadlock?

For a deadlock to occur, the following 4 conditions all need to be fulfilled. They are sometimes called *the 4 conditions of deadlock*:

1. Mutual exclusion

There are resources (e.g., variables or other data) in the program that threads need exclusive access in order to access. In other words, there are critical sections in the program that we protect using locks to ensure that only one thread access the resources at a time.

2. No preemption

Once a thread has gained exclusive access to some resource, only the thread itself can voluntarily relinquish the exclusive access. For example, once a thread has successfully acquired a lock, the thread will continue to hold the lock until the thread itself releases the lock by calling `lock_release`. Other threads can not force the thread to release the lock prematurely.

3. Hold and wait

A thread that has gained exclusive access to some resource requests exclusive access to another resource and needs to wait. For example, the thread T currently holds lock A and calls `lock_acquire` to acquire lock B. However, another thread holds lock B, so T needs to wait. As such, T holds lock A and waits for lock B.

4. Circular wait

Two or more threads fulfill *hold and wait* so that they form a cycle. For example as we saw in the example above: we have two threads, 1 and 2. Thread 1 holds lock A and waits for lock B. Thread 2 waits for lock B and holds lock A.

The first three conditions (1–3) are usually referred to as *necessary conditions* for deadlock, while the last one (4) is referred to as the *sufficient condition* for deadlock. The necessary conditions stipulate properties that are required in order for circular wait to arise (i.e., the sufficient condition), but does not necessarily mean that deadlock *will* arise. The sufficient condition is, as the name implies, sufficient for a deadlock to arise. If a circular wait has occurred, then a deadlock is happening, and the necessary conditions have to be fulfilled.

7.1.1 Implications of the Four Conditions

To better understand the implications of the four conditions, we will first take a closer look at each of the four conditions of deadlock. Both at the relevance of each of them, but also their implications on programs.

As mentioned above, *mutual exclusion*, essentially means that the program uses some form of synchronization to ensure mutual exclusion when accessing some resource. This is most commonly done using locks, and we will therefore use locks as an example in the remainder of this chapter. However, it is important to remember that locks are simply one of multiple possible means to

achieve mutual exclusion. Deadlock may just as well happen if semaphores or other synchronization primitives are used. The key insight is that regardless of *how* mutual exclusion is enforced, there are times when threads need to wait for some other thread to finish using the shared resource. The waiting is the problematic part, not the locks themselves.

The name *no preemption* refers to the fact that it is not possible to *preempt* a thread *from its exclusive access to a resource*. It is worth noting that the word *preemption* is also used to describe schedulers. A *preemptive* scheduler is able to force the currently running thread to stop executing on the CPU and give other threads the opportunity to execute (as discussed in Chapter 3). In contrast, a *non-preemptive* or *cooperative* scheduler is only allowed to switch between whenever the running thread explicitly allows it.

The same word is used for condition 2 above, because it describes the same situation. However, the key difference is that condition 2 speaks about whether threads can be forced to give up exclusive access to a resource or not. As described above, *no preemption* means that the only way a thread can give up exclusive access to a resource is to give it up voluntarily. In terms of locks, this means that if thread T acquires a lock by calling `lock_acquire`, the thread itself must call `lock_release` to release the lock. There is no other way a second thread can forcefully take over the lock from T without T calling `lock_release`. As you can imagine, writing correct code in a system where this is possible would be borderline impossible. The whole point of a lock is to protect a critical section so that other threads are unable to interfere with shared data. If other threads would be able to take over locks, this is no longer true and most benefits of mutual exclusion disappear.¹

One important observation from above is that without the ability to preempt locks from threads, the only way a thread can release a lock is by calling `lock_release` itself. That is, the thread needs to execute code in order to release a lock. As such, if a thread is waiting for something else, it is unable to release any locks it is currently holding.²

The last of the necessary conditions, *hold and wait*, combines the key insights of the two previous conditions. That is, *mutual exclusion* means that threads may need to wait for something, and *no preemption* means that threads that wait are unable to release resources. If we combine them, we get *hold and wait*, which is the observation that if a thread, T, currently holds a lock, A, and attempts to acquire another lock, B, the T will be unable to release A until it acquires B. One way to think of this is that T creates a link between A and B. In this case, A can not be acquired until the thread that holds B releases it.

The final condition, *circular wait*, essentially uses *hold and wait* to create a cycle of threads waiting for each other. We just saw that we could use hold and wait to link threads together. If we have two threads, T and U, and assume that T holds A while U holds B. Then if T attempts to acquire B, we get a situation similar to in the previous paragraph. Lock A can not be released until

¹One possibility is to use *transactional memory*, but the hardware support required for that is scarce at the time of writing.

²For example, this means that even if you implement a mechanism that allows threads to ask the thread currently holding a lock to release its lock, we still have *no preemption*. Since the other threads *ask* rather than *take* the lock, the thread holding the lock still needs to realize that it has been asked to release the lock, perhaps finish a part of what it is doing, and then release the lock.

thread T successfully acquires lock B, which requires that thread U releases it. However, if instead of releasing lock B, thread U attempts to acquire lock A we get a problem. At this time, we have essentially created two links: lock A will not be released until lock B is released, and lock B will not be released until lock A is released. Since we have created these links in a circle, we can conclude that none of the locks will be released (due to *no preemption*), and therefore that neither thread T nor thread U will ever acquire their second lock.

7.1.2 Resource Allocation Graphs

As you have most likely realized from the textual descriptions above, it is difficult to get an overview of even relatively small systems without a good representation of the state of the system. Resource allocation graphs is representation that is useful for representing threads in relation to mutually exclusive resources.




Figure 7.2: The initial state of a system that contains two threads (T and U) and two shared resources (A and B).

We will use the scenario above to demonstrate how resource allocation graphs can be used to illustrate the state of a concurrent system. As a starting point, a program contains two threads, T and U, and two shared resources that require mutually exclusive access. This initial state is depicted in Fig. 7.2. The threads are represented as circles that contain the name of the thread. Each resource is represented as a dot. The dots are in turn grouped into squares with a label. All dots within the same box are treated as equivalent and interchangeable, and as such only the boxes are labeled. In this example, we imagine that the resources A and B correspond to locks that each protect a shared variable. Therefore, each box contains a single dot (we can think of the dot as corresponding to “the ability to access the shared variable”).

If a box contain more than one dot, all dots are treated as equivalent. This represent cases when a certain thread will need one instance of a resource but does not care which one it gets, since they are equivalent. As an example, imagine a thread that needs to store temporary data on a disk connected to the computer. If the system has multiple disks, the program does not care *which* one it gets to use, as long as it gets exclusive access to that disk. We could also model RAM usage in the same way by letting each dot represent 1 MiB of RAM for example. When working with examples in the scope of this book, these situations will be rare, since none of the standard synchronization

primitives support this use-case out of the box (i.e., there is no version of the lock that acquires one out of many resources, but it is possible to implement one if we need one as we will see in Section 8.4).




Figure 7.3: Thread T has acquired one instance of resource A and thread U has acquired one instance of resource B.

The next step in our scenario is that thread T acquires one instance of resource A, and thread U acquires one instance of resource B. As shown in Fig. 7.3, we illustrate this by drawing an arrow from the dot in the resource to the thread that has access to it. One way to think of the representation is that the arrow indicates that a particular dot is moved to the thread itself. In this case, that the threads have the ability to access the variables protected by locks A and B.




Figure 7.4: Thread T has acquired one instance of resource A and thread U has acquired one instance of resource B.

After T has successfully acquired an instance of resource A (e.g., acquired the lock), it attempts to acquire an instance of resource B. As we can see from Fig. 7.3, this is not possible, since the only instance of resource B is currently held by thread U. As such, thread T needs to wait. Again, we represent this as an arrow, but this time from the thread to the box that represent the resource(s) the thread is waiting for, as shown in Fig. 7.4. The reason we draw the arrow to the box rather than one of the dots in the box is to make it clear that the thread is waiting for *one* of the possible multiple equivalent resources that are in the box.

At this point we can start to see the benefits of the graphical notation. By following the arrows, we can see that resource A is acquired by thread T, that thread T is waiting for an instance of resource B, that is in turn acquired by thread U. As such, resource A is currently indirectly blocked by thread U, as it can not be released until thread U releases B.




Figure 7.5: Thread T has acquired one instance of resource A and thread U has acquired one instance of resource B.

Finally, thread U attempts to acquire an instance of resource A. Again, no instance of resource A is available, so thread U needs to wait. As before, we indicate this by adding a line from thread U to the box labeled A. This results in the representation in Fig. 7.5. In the figures we can see that the lines form a cycle. This means that resource A is held by thread T, which is waiting for resource B, which is held by thread U, which in turn waits for resource A. As such, the two threads are indirectly waiting for themselves, which is clearly visible in the resource allocation graph.




Figure 7.6: An example of a resource allocation graph with multiple instances of resource B.

It is worth noting that some care needs to be taken when working with resources that have multiple instances. For example, resource B in Fig. 7.6 has two instances. One instance is acquired by thread U and the other by thread V. In this case, even though the graph *does* contain a cycle, the threads are actually not currently in a deadlock. Since one instance of B is held by thread V, and thread V is currently not waiting for anything, V can release its instance of resource B, which means that thread T can acquire it and resolve

the deadlock, which results in the situation in Fig. 7.7. As such, for a deadlock to have occurred in graphs where there are multiple instances of some resource the graph needs to contain cycles for *all* instances.




Figure 7.7: The situation in Fig. 7.6 after thread V released its instance of resource B.

It is worth noting that the situation in Fig. 7.6 may eventually lead to a deadlock. For example, if thread V would have attempted to acquire an instance of resource A rather than releasing its instance of resource B then both instances of resource B would have been part of a cycle, and the system would have been in a deadlock.

7.2 Preventing Deadlocks

The 4 conditions for deadlock (Section 7.1) can also be used to prevent deadlocks from occurring. The idea is simple, we know that all conditions need to be fulfilled for a deadlock to occur. This means that we can avoid deadlocks by ensuring that at least one of the conditions can never be fulfilled.

As such, we will discuss how we can break each of the conditions below. Then we will see how we can apply them to prevent deadlocks from occurring in the `bank-1` program.

1. Mutual exclusion

Since deadlocks can not occur without mutual exclusion, we can avoid deadlocks by avoiding mutual exclusion. As we have seen earlier in this book, we typically need mutual exclusion *somewhere* in a concurrent program in order to avoid data races, so it is often not possible to avoid mutual exclusion altogether.

2. No preemption

Another option to avoid deadlocks would be to allow threads to take over locks that other threads have acquired. This is often not a viable approach either, since this more or less eliminates the benefits of locks altogether.³

³At least without mechanisms like *transactional memory*.

3. Hold and wait

Sometimes it is possible to structure the program so that it never holds a lock while waiting for something else. One way to achieve this is to make sure to never hold more than one lock at a time. Doing this means that it is impossible to build “links” between two resources, and thereby circular wait can not occur.

4. Circular wait

Even when the other conditions are fulfilled, it is still possible to avoid a circular wait. That is, implementing some way to make sure that cycles in the resource allocation graph can not form. A simple way to achieve this is to ensure that all threads acquire their locks according to some total order.

It is worth examining the idea for how circular wait can be broken in more detail. Consider a system with 5 resources and 5 threads as depicted in Fig. 7.8. As can be seen in the figure, the resources are labeled R1–R5 and the threads are labeled T1–T5. The system is also in a deadlock since the edges in the graph form a cycle that involves all threads and all resources.




Figure 7.8: Generalization of a circular wait in a resource allocation graph.

One important observation from the graph is that for all threads except T5, thread T_n has acquired resource R_n and waits for resource R_{n+1} . That is, thread T1 has acquired R1 and then attempted to acquire R2, but had to wait since R2 was already acquired by another thread. In particular, this means that each of the threads hold resources that have lower numbers than the resource they are waiting for (i.e., the arrows go to the right in the figure). This is, however, *not* true for T5, which holds R5 and waits for R1 (i.e., the arrows go to the left in the figure). As such, T5 needs to have acquired R5 before it attempted to acquire R1.

The key observation here is that we need threads like T5 that “break the rule” in order to form a cycle. Without T5, all edges from resources in the figure would go from left to right, and thereby we can not form a cycle. We need at least one edge from a resource that goes right to left to be able to close the cycle. If T5 would follow the same pattern as the other threads, that is it acquires the lower numbered resource before the higher numbered resource, either T1 or T5 would have to wait for R1 before attempting to acquire its second resource. For example, if T1 acquired R1 before T5, the resource allocation graph would

look like Fig. 7.9. Importantly, the rule that multiple resources need to be acquired according in the same order prevents cycles from forming since the highest numbered resource will always be acquired last.




Figure 7.9: Generalization of a circular wait in a resource allocation graph.

Practice: In the example above, threads acquire resources with numbers adjacent to each other (i.e., R1 and R2, R2 and R3, etc.). Of course the idea works regardless of exactly which resources need to be acquired, and even the number of resources that should be acquired. To illustrate this, consider the scenario below. First, draw one resource allocation graph where threads acquire the resources in the order they are listed below. Then, draw another resource allocation graph where threads instead acquire resources in a fixed order (e.g., increasing or decreasing). Remember that the thread stops execution once it needs to wait for a resource that is acquired by another thread!

T1: R2, R4

T2: R1, R3, R4

T3: R5, R2

T1 (after T3 is done): R5

If you have done the diagrams correctly, you will see that the system enters a deadlock if you follow the order above, but the deadlock is avoided when locks are acquired in a specific order.

Finally, it is worth noting that the numbering of the threads does not matter at all. It just makes it easier to see what happens in the example. The numbering of the locks is not important. What matters is that all locks that are used together are always acquired in the same order. For example, taking the locks above in the order 4, 2, 3, 1, 5 will still avoid deadlock. This shows that the important part is that there is *a* consistent order, not *which* order is used.

7.3 Deadlock Avoidance in Programs

With knowledge of both what a deadlock is and how deadlocks can be avoided, it is time to apply the knowledge to the `bank` program. As we saw in Section 7.2, we have two options. We either break *hold and wait* or *circular wait* (as we have seen, we need *mutual exclusion* since there is shared data). Breaking *hold and wait* is possible by replacing the per-account `balance_lock` with a single lock for all accounts, which gives us the program `bank-sum-slow.c` covered in the previous chapter (see Listing 6.15). However, as mentioned in the previous chapter this means that it is not possible for multiple threads to transfer money between different pairs of accounts concurrently. As mentioned previously, this might be a reasonable solution depending on how `transfer` and `accounts_total` are used. After all, the implementation of `bank-sum-slow.c` is simpler than that of `bank-1.c`, and since it does not have *hold and wait*, it can never cause a deadlock.

As we saw above, breaking *hold and wait* sometimes requires compromising on the available parallelism.⁴ As such, we will instead try to break *circular wait* and hope that it leads us to a solution that avoids deadlock while allowing multiple transactions to happen concurrently. The easiest way to achieve that is to use the idea outlined at the end of Section 7.2, namely to make sure that threads acquire locks according to some total ordering. One way to do this is to give each lock a number and make sure to acquire locks in increasing numerical order (at least conceptually, we do not necessarily need to write it down other than “always acquire lock X before lock Y”). In the case of `bank-1.c`, this is quite easy. The only locks we have in the program are the `account_locks` for each account. Since the accounts are already numbered, we can simply use the account number to order the locks.

Practice: Apply the idea outlined above to solve the deadlock we discovered in `bank-1.c` at the start of the chapter. The file `bank-1-larger-test.c` contains the same code as `bank-1.c`, but has a main program that also calls `accounts_total` so that you can verify that `transfer` and `accounts_total` work properly together. You can use `Run ⇒ Look for errors...` to verify your solution.

As you have likely realized above, the idea of utilizing account numbers to ensure threads acquire the `account_locks` in the same order can be applied quite easily. One way is the approach implemented in `bank-2.c`. As can be seen in Listing 7.1, we simply check which account number is larger before acquiring the locks, and acquire the lock for the account with the lowest number first. Note that even though it may seem natural to release the locks in the opposite order we acquired them in, the order does not matter at all since `lock_release` never needs to wait.

Even though we have focused on `transfer` so far in the chapter, it is important to note that it is important that *all* parts of the program adhere to the

⁴This depends on the program as a whole. For example, the program `bank-3.c` from the previous chapter does not have *hold and wait* even though it uses one lock for each account. The same is true for `bank-sum-2.c`, even though that program has other odd behaviors that are probably undesirable.

```

bool transfer(int amount, int from, int to) {
    struct account *f = &accounts[from];
    struct account *t = &accounts[to];

    if (from < to) {
        lock_acquire(&f->balance_lock);
        lock_acquire(&t->balance_lock);
    } else {
        lock_acquire(&t->balance_lock);
        lock_acquire(&f->balance_lock);
    }

    bool ok = f->balance >= amount;
    if (ok)
        f->balance -= amount;
    if (ok)
        t->balance += amount;

    lock_release(&t->balance_lock);
    lock_release(&f->balance_lock);
    return ok;
}

```


Listing 7.1: Modifications to `bank-1.c` to avoid deadlocks.

order in which locks should be acquired. In this case, we also need to make sure that `accounts_total` acquires locks in the same order as `transfer`. By accident, `bank-2.c` happens to be correct simply because it was natural to acquire the lock for the account with the lowest number first in both cases.


Practice: Change the condition in the if statement from `from < to` into `from > to` (available as `bank-2-error.c`), and use *Run* \Rightarrow *Look for errors...* to see what happens when `transfer` acquires locks in a different order compared to `accounts_total`.

7.4 Exercises


1. Below is a resource allocation graph that shows four threads and three resources. Which threads (if any) are a part of a deadlock?




2. Below is a resource allocation graph that shows four threads and three resources. Which threads (if any) are a part of a deadlock?



3. Below is a resource allocation graph that shows four threads and three resources. Which threads (if any) are a part of a deadlock?



4. Below is a resource allocation graph that shows four threads and three resources. Which threads (if any) are a part of a deadlock?



5. The file `increase.c` contains a simple program that attempts to keep the values of two global variables in sync with each other. The variable `high` should always contain a higher value than `low`. Furthermore, `low` should not be more than 2 lower than `high`.

The program contains two functions to increment each variable: `inc_high` increments `high` and `inc_low` increments `low`. Both functions also check the other variable to see if that also needs to be incremented to maintain the rules mentioned above.

- Look at the implementation of `inc_low` and `inc_high`. Is it possible for a deadlock to occur if they are called from different threads?
 - Modify the implementation of `inc_low` and `inc_high` to eliminate any deadlocks you found in a). You can verify your solution in Progvis. If you do, uncomment the call to `verify` in `main`. You might need to modify how `verify` acquires its locks to match with your solution.
6. The file `books.c` contains a program that stores the names of books in a bookshelf. The function `move` moves a book to a new location. The implementation is currently synchronized using a lock for each location in the bookshelf.
- It is possible for a deadlock to occur if `move` is called in the right way. You will, however, notice that Progvis currently reports that all is well. This means that `move` can not end up in a deadlock as it is currently used.
Modify the code in `main` to find a situation where `move` may end up in a deadlock. Progvis will tell you when you have found it.
 - Modify `move` to make it impossible for a deadlock to occur. The scenario you found in a) should be enough to test your solution.

Condition Variables

Up to this point we have introduced semaphores and locks. As we have seen, semaphores are powerful enough to solve any concurrency issue while locks are specialized for mutual exclusion. This makes locks easier to use to protect shared data, but also limits their capabilities. In particular, it is *not* possible to use locks to wait for some event to occur (i.e., we can *not* implement a semaphore using only one or more locks).

This chapter introduces *condition variables*, which can be used in conjunction with locks to make it possible to waiting for events. As such, condition variables can be viewed as an extension of the capabilities of locks to make them as powerful as semaphores (i.e., we can implement a semaphore using a lock and a condition variable). Even though we do not gain any additional power in terms of problems that are solvable by introducing condition variables, they work differently from semaphores. Therefore it is easier to solve some synchronization problems with locks and condition variables than with semaphores, and vice versa.

8.1 Semantics

As in Chapter 5, we start by introducing the available operations and their semantics. Again, we focus on how the operations are used. The full definitions are available in Chapter A for the curious reader.

A condition variable can be thought of as containing a set of zero or more threads that are currently waiting. Just as semaphores and locks, a condition variable is an opaque data structure that is only possible to modify by calling the operations below.

It is also useful to consider the condition variable to be associated with a lock that is used to protect some variables that we wish to wait until they have a particular value. To mirror the semantics in Pintos¹ we consider the condition variable itself to *not* be threadsafe, and the lock to protect the condition variable as well. While this strict view is not always necessary (many implementations of condition variables *are* threadsafe), it is useful since it avoids sporadic wakeups, and some implementations (e.g., pthreads) still require that

¹Which is the source of the nomenclature for the synchronization primitives used in this book.

each condition variable is used with one lock. As such, thinking of the condition variable as not being threadsafe and thereby needing protection from a lock makes it easier to reason about the condition variable and the behavior of the program.

```
struct condition c;
```

The line above defines a variable named `c` that contains a condition variable. As we can see, the type is named `struct condition`.

```
struct lock l;
```

To use the condition variable, we also need a lock. Each condition variable should be associated with exactly one lock (i.e., each condition variable should be protected by one lock, but using one lock to protect multiple condition variables is fine). For this reason it is useful to think of the condition variable as a shared variable that needs to be protected by the associated lock, even if this is not strictly necessary in most implementations.

```
cond_init(&c);
```

Each condition variable needs to be initialized before it is used. As with locks, the initialization function does not require any additional parameters apart from a pointer to the condition variable to initialize. Of course we need to initialize the lock as well, but we omit the initialization since this chapter focuses on condition variables.

```
cond_wait(&c, &l);
```

Put the current thread to sleep and store the thread in the condition variable.² The function assumes that the current thread holds the lock (`l`). It makes sure to release the lock just before the thread is put to sleep, and to re-acquire the lock once the thread wakes up again. As such, the calling thread needs to hold the lock before calling `cond_wait` and will still hold the lock afterwards. However, it is important to remember that the lock is released during the call to `cond_wait`, since other threads may have acquired the lock and changed the variables protected by the lock.

```
cond_signal(&c, &l);
```

Wakes *one* of the threads that is currently sleeping and stored inside the condition variable `c`. Requires that the lock `l` is held by the current thread. If no threads are stored in the condition variable, nothing happens. If more than one thread is sleeping, an arbitrary thread is picked to wake up.

```
cond_broadcast(&c, &l);
```

Wakes *all* of the threads that is currently sleeping and stored the condition variable. Requires that the lock `l` is held by the current thread.

It is worth noting that many implementations of condition variables do not require that the lock (`l`) is passed to `cond_signal` and `cond_broadcast`.

²The condition variable typically stores some form of reference to the thread (e.g., a thread identifier) rather than the thread itself, but conceptually we can consider it to store the thread.

They also do not require that the lock associated with the condition variable is held by the current thread when `cond_signal` and `cond_broadcast` is called. However, as mentioned above, holding the lock while calling `cond_signal` and `cond_broadcast` is still good practice as it eliminates certain forms of sporadic wakeups. Furthermore, in most cases the thread that calls `cond_signal` and `cond_broadcast` needs to hold the lock just before the call anyway, so the call to `lock_release` can simply be placed after `cond_signal` and `cond_broadcast` rather than acquiring the lock again.

8.1.1 Using `cond_wait` and `cond_signal`

Now that we have seen the operations provided by a condition variable, we start by observing them in practice in Progviz. For this, we will use the program `semantics-signal.c` which is shown in Listing 8.1. The program declares two global variables, a condition named `cond` and a lock named `cond_lock`. Note that we name the lock `cond_lock` to remind ourselves that we use the lock to protect the variable `cond`, just as we have done with other variables previously. After initializing the lock and the condition variable, the `main` function starts a new thread that executes `thread_fn`, acquires the lock, calls `cond_signal`, and then exits. The second thread simply acquires the lock, calls `cond_wait`, and releases the lock.

```
struct condition cond;
struct lock cond_lock;

void thread_fn(void) {
    lock_acquire(&cond_lock);
    cond_wait(&cond, &cond_lock);
    lock_release(&cond_lock);
}

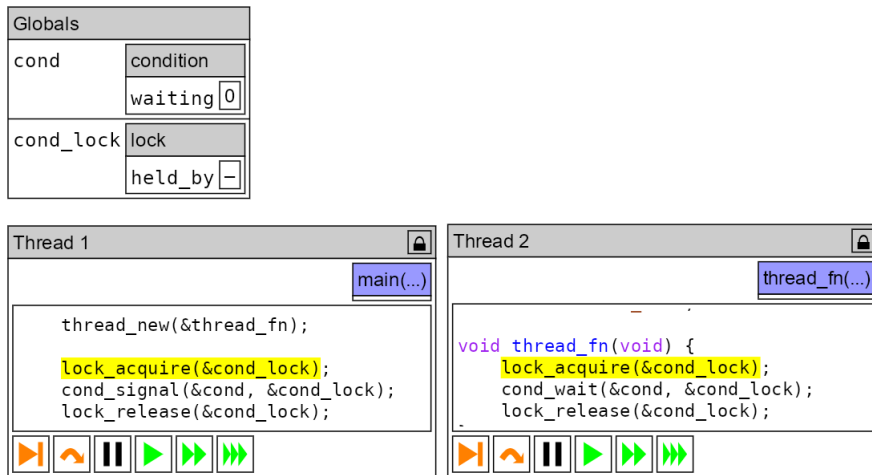
int main(void) {
    cond_init(&cond);
    lock_init(&cond_lock);

    thread_new(&thread_fn);

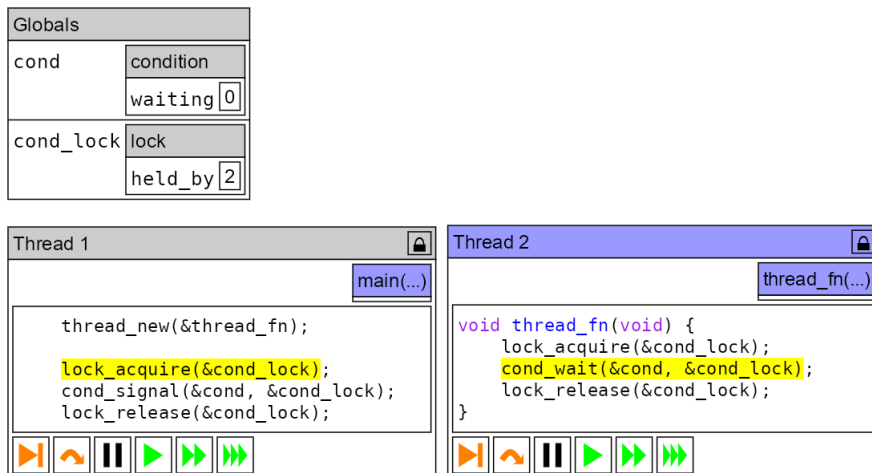
    lock_acquire(&cond_lock);
    cond_signal(&cond, &cond_lock);
    lock_release(&cond_lock);

    return 0;
}
```

Listing 8.1: The program `semantics-signal.c` used to illustrate the semantics of condition variables.

Figure 8.1: The program `semantics-signal.c` after initialization.

If you open the program in Progvis and step Thread 1 to the end of the initialization, you will see a figure similar to the one in Fig. 8.1. Note that Progvis represents a condition variable as a box that contains a value named *waiting* whose value is currently zero. This is how Progvis represents the number of threads that are currently waiting on the condition variable.

Figure 8.2: The program `semantics-signal.c` after Thread 2 has acquired the lock.

If we step Thread 2 once, the program will reach the state shown in Fig. 8.2, where Thread 2 has just acquired the lock `cond_lock`. As we have seen before, Progvis represents this by displaying the number 2 next to the text *held_by* in the box for the lock.

If we step Thread 2 once more, it will call `cond_wait`. As the name implies, this puts the thread to sleep. Progvis illustrates this as in Fig. 8.3. First and foremost, we can see that Thread 2 is sleeping, both since the highlighted line

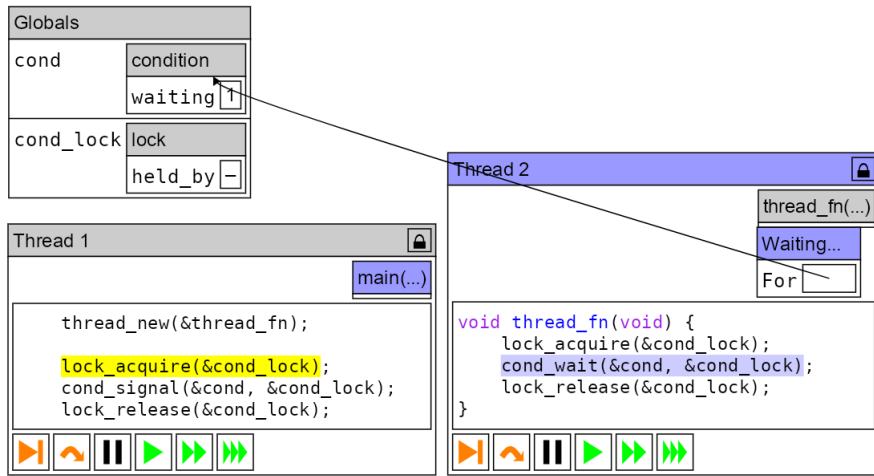


Figure 8.3: The program `semantics-signal.c` after Thread 2 has called `cond_wait`.

is blue, and due to the box labeled *Waiting...* on the thread's call stack that points to the condition variable. We can also see that the condition variable in the *Globals* box says that one thread is waiting. Another important detail that is easy to miss is that Thread 2 *no longer* holds the lock. We can see this by observing that *held_by* is a dash, which means that the lock is not held by any thread. This is fortunate, since Thread 1 needs to acquire the lock to call `cond_signal`.

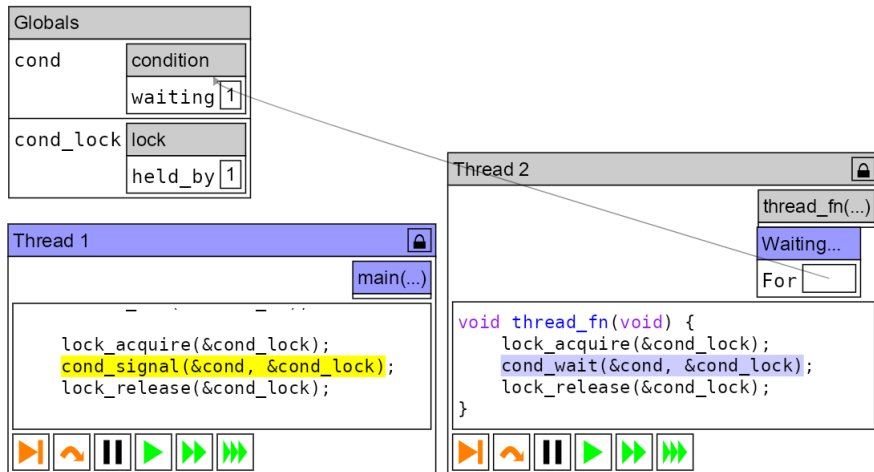


Figure 8.4: The program `semantics-signal.c` after stepping Thread 1.

At this point, our only option is to step Thread 1. If we do that, we will end up in the state depicted in Fig. 8.4, where Thread 1 has acquired the lock.

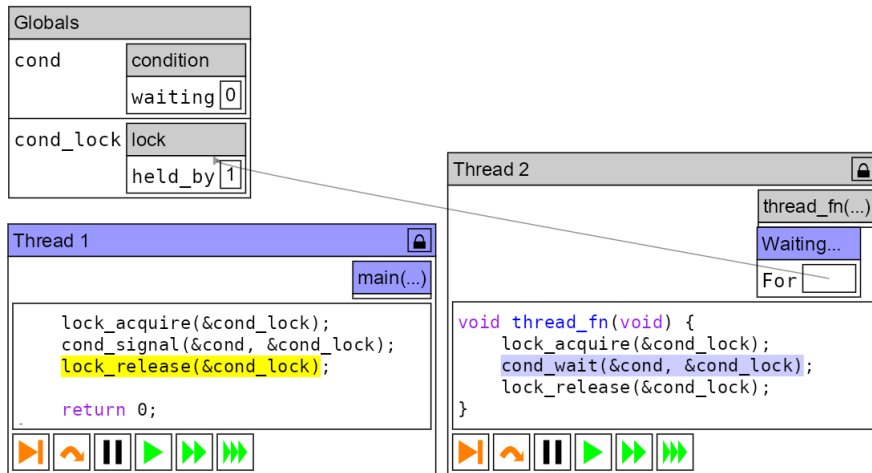
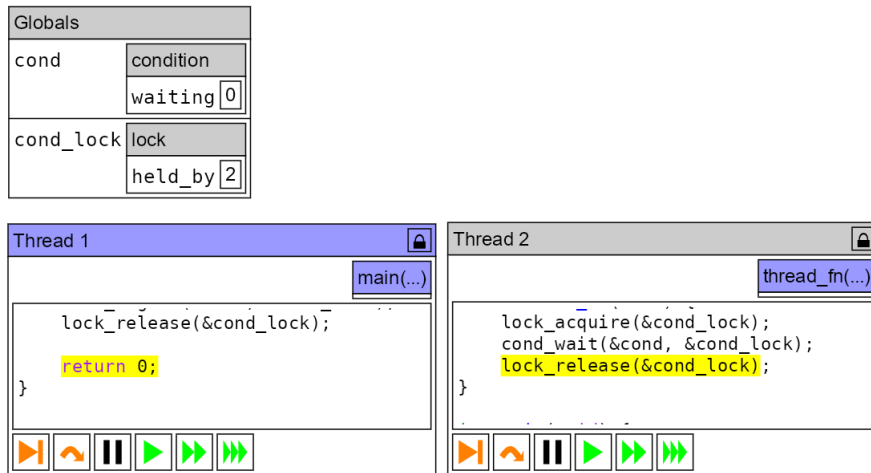


Figure 8.5: The program `semantics-signal.c` after stepping Thread 1 yet again, just after it has returned from the call to `cond_signal`.

If we step Thread 1 once more, it will call `cond_signal`. This instructs the condition variable to wake *one* of the threads that are currently waiting. In this case, Thread 2 is the only option, so the condition variable wakes Thread 2. After this, `cond_signal` returns and the program ends up in the state in Fig. 8.5. Interestingly, Thread 2 is still waiting for something. If we follow the arrow from the `Waiting...` box in Thread 2's call stack, we will see that the thread is indeed no longer waiting for the condition variable, but for the lock! Remember that `cond_wait` does three things: (1) it releases the lock, (2) puts the thread to sleep, and (3) once the thread wakes up it re-acquires the lock. What happened here³ is that the thread was woken up and then immediately attempted to re-acquire the lock, which is still held by the thread that called `cond_signal` (Thread 1 in this case).

If we step Thread 1 a final time, it will call `lock_release` and thereby allow Thread 2 to continue. As shown in Fig. 8.6, Thread 2 now holds the lock. This might initially seem strange. However, if we consider `lock_acquire` and `lock_release` as markers for the start and end of a critical section, it is sensible to let `cond_wait` denote a temporary break in the critical section. That way, the intuition that `lock_acquire` and `lock_release` denotes the start and end of a critical section remains, and we do not have to remember if `cond_wait` is replaces the start or the end of the critical section.

³This is indeed a common case, especially in our implementation where it is necessary to hold the lock when calling `cond_signal`. Because of this, many implementations optimize this case by not actually waking the thread just for it to immediately sleep again since the lock is still held by the current thread. Instead, they simply move the thread from the waiting queue for the condition variable to the waiting queue for the lock.

Figure 8.6: The program `semantics-signal.c` after Thread 1 releases the lock.

Practice: The example above only illustrated what happens if Thread 2 calls `cond_wait` *before* Thread 1 calls `cond_signal`. Use Progvis to investigate what happens if the threads execute the other way around!

As you have likely noticed above, if Thread 1 calls `cond_signal` before Thread 2 has called `cond_wait`, no threads are waiting for the condition when `cond_signal` is called, and thereby no threads are woken up. This is not considered an error, so Thread 1 simply continues executing. However, once Thread 2 calls `cond_wait`, no thread will wake it up and it will therefore sleep forever. Progvis will report this as a deadlock once Thread 1 terminates. However, as discussed in Chapter 7, this situation is not what is usually meant by the term *deadlock*, even if the symptoms are similar to a cycle of threads waiting for each other.

8.1.2 The Difference Between `cond_signal` and `cond_broadcast`

The example above only used `cond_signal`. Before we move on to examine how programs are used we first take a closer look at the difference between `cond_signal` and `cond_broadcast`. For this, we use the program `semantics-broadcast.c`, which is the same as `semantics-signal.c` except that the main function is replaced with the implementation in Listing 8.2.

```

int main(void) {
    cond_init(&cond);
    lock_init(&cond_lock);

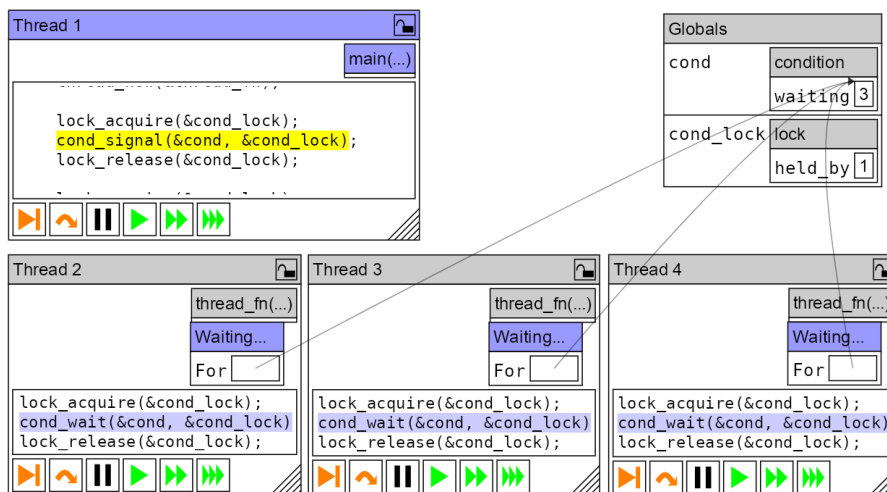
    thread_new(&thread_fn);
    thread_new(&thread_fn);
    thread_new(&thread_fn);

    lock_acquire(&cond_lock);
    cond_signal(&cond, &cond_lock);
    lock_release(&cond_lock);

    lock_acquire(&cond_lock);
    cond_broadcast(&cond, &cond_lock);
    lock_release(&cond_lock);

    return 0;
}

```

Listing 8.2: The main function of `semantics-broadcast.c`.Figure 8.7: The program `semantics-broadcast.c` once threads 2–4 are waiting for the condition and Thread 1 is about to call `cond_signal`.

We will start from the state in Fig. 8.7, where threads 2–4 are waiting on the condition inside `cond_wait`, and Thread 1 is about to call `cond_signal`. To reach the state, start by stepping Thread 1 until it has called `thread_new` three times. Then step threads 2–4 until they all wait inside `cond_wait`. Finally, step thread 1 until it reaches the line `cond_signal` in Fig. 8.7.

At this point, three threads are waiting on the condition. If we step Thread 1 it will call `cond_signal` and we will be able to see how it behaves. As noted before, `cond_signal` wakes *one* of the threads that wait for the semaphore (as with other synchronization primitives, it is not well-defined which one, so it is not possible to rely on first-in first-out). In this case, Thread 2 is woken, and

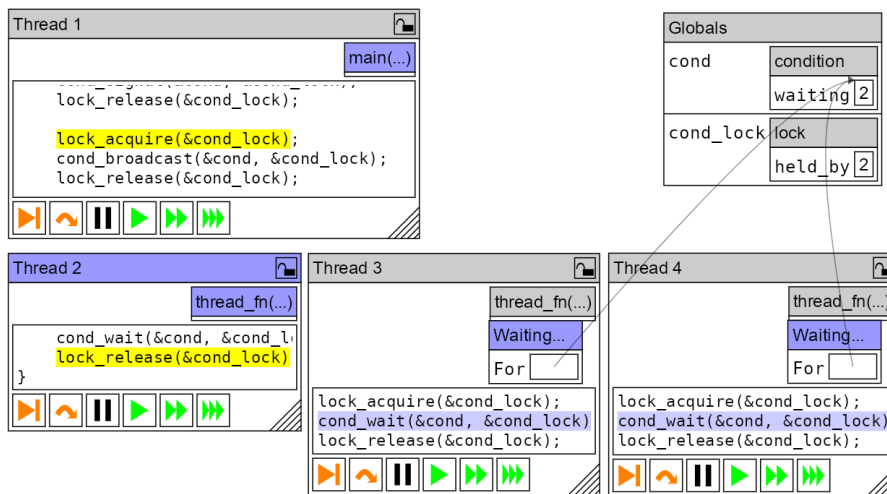


Figure 8.8: The program `semantics-broadcast.c` after `cond_signal` wakes one of the threads.

as before Thread 2 will immediately start waiting for the lock that is currently held by Thread 1. Once we step Thread 1 once more, it releases the lock and thereby allows Thread 2 to acquire the lock and continue, as shown in Fig. 8.8.

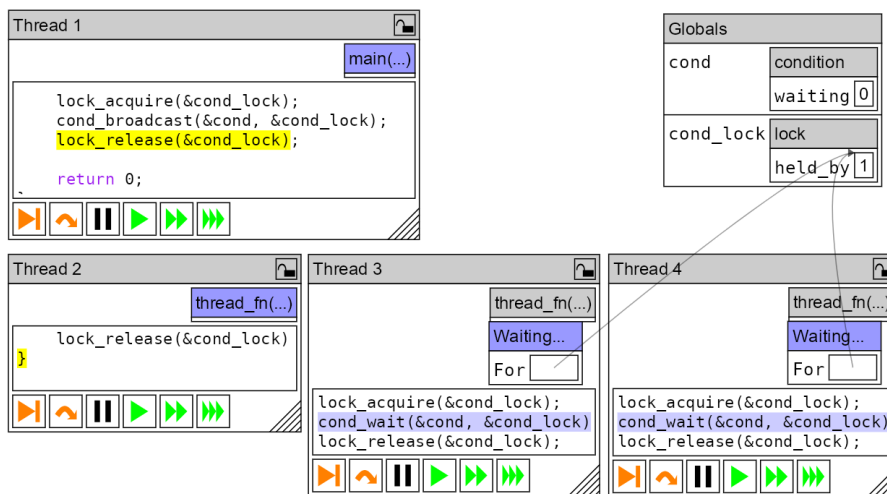


Figure 8.9: The program `semantics-broadcast.c` just after `semantics-broadcast.c` after `cond_broadcast` has been called.

If we continue stepping Thread 1 until it has called `cond_broadcast` (which requires stepping Thread 2 until it releases the lock), we see the difference between `cond_signal` and `cond_broadcast`. While `cond_signal` only wakes *one* of possible multiple threads that wait on the condition variable, `cond_broadcast` wakes *all* of them. We can see this in Fig. 8.9, since both Thread 3 and Thread 4 are now waiting for the lock rather than the condition. Just as before, this is because `cond_wait` re-acquires the lock when threads wake up, before it returns

to the caller.

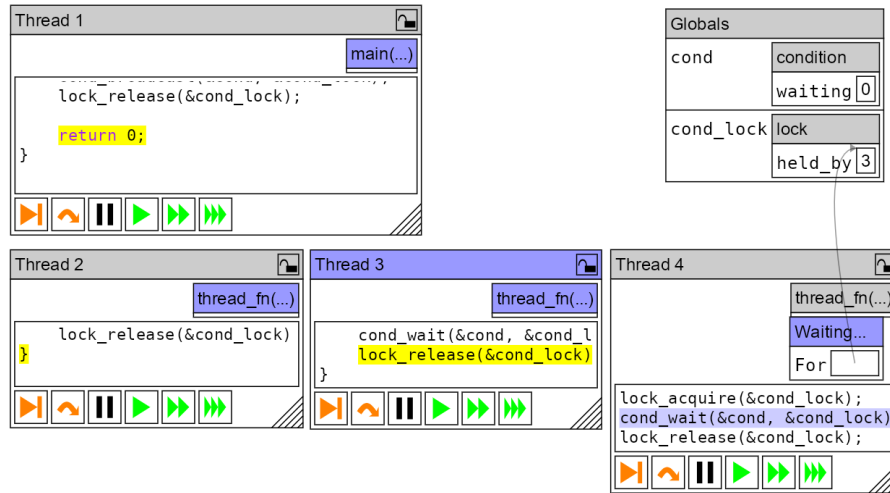


Figure 8.10: The program `semantics-broadcast.c` after Thread 1 releases the lock, which allows Thread 3 to acquire the lock.

If we let Thread 1 execute another step, it will release the lock. This allows *one* of the two threads to acquire the lock and continue execution. In Fig. 8.9, this happens to be Thread 3 (again, just as with other synchronization primitives, the order is typically not well-defined so we can not rely on some particular ordering even though Progvis follows first-in first out).

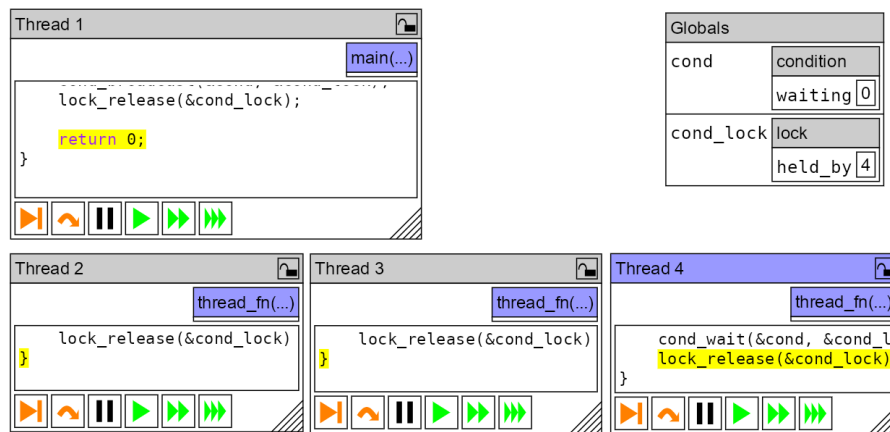


Figure 8.11: The program `semantics-broadcast.c` after Thread 3 releases its lock.

At this point, it might seem like there is no real difference between `cond_signal` and `cond_broadcast` since the lock only allows one thread to continue anyway. The key difference is, however, that Thread 4 is currently waiting for the lock rather than the condition variable (in contrast to Fig. 8.8). This

means that once Thread 3 releases the lock (which it does if we step it once), Thread 4 will be allowed to acquire the lock and thereby continue. As such, `cond_broadcast` allows all threads to continue, but the lock makes them start one by one in some order.

We could of course achieve the same behavior by calling `cond_signal` two times in Thread 1. However, since it is not possible to ask a condition variable if any threads are currently waiting on it, we must either provide such a mechanism ourselves, or call `cond_signal` enough times, which is likely at least a few too many. Neither of these solutions are very convenient, so `cond_broadcast` is a good alternative here, especially since `cond_broadcast` is a fairly common operation.

8.2 Condition Variables and Semaphores

The two programs we used above had the same issue. If `cond_signal` or `cond_broadcast` was executed before `cond_wait`, some threads would sleep forever. This was not a problem we encountered when using semaphores! Why is this an issue with condition variables and not semaphores?

<pre> int result; struct semaphore sema; void thread_fn(void) { result = 10; sema_up(&sema); } int main(void) { sema_init(&sema, 0); thread_new(&thread_fn); sema_down(&sema); printf("%d\n", result); return 0; } </pre>	<pre> int result; struct condition cond; struct lock l; void thread_fn(void) { result = 10; lock_acquire(&l); cond_signal(&cond, &l); lock_release(&l); } int main(void) { cond_init(&cond); lock_init(&l); thread_new(&thread_fn); lock_acquire(&l); cond_wait(&cond, &l); lock_release(&l); printf("%d\n", result); return 0; } </pre>
--	--

Figure 8.12: Comparison between using a semaphore and a condition variable to wait until a thread is finished. The left program is `compare-sema.c` and the right is `compare-cond.c`.

To illustrate the difference, consider the two programs in Fig. 8.12. The program on the left uses a semaphore to wait for the thread running `thread_`

`fn` to set `result` to 10. As you can probably easily conclude by now, the program works as expected. The program on the right attempts to use `cond_wait` and `cond_signal` to do the same thing. As you can immediately see, the approach that uses a condition variable involves more code. More problematic, the program does not always work as expected. Remember that `cond_signal` only wakes the threads that are waiting on the condition variable when `cond_signal` is called. As such, if the thread running `main` reaches `cond_wait` *after* the other thread calls `cond_signal`, the main thread will never wake up.

Practice: Verify the problem in `compare-cond.c` using Progviz. Using *Run* \Rightarrow *Look for errors...* you can see a visualization of the error. As noted above, `compare-sema.c` does *not* have the same issue, even though the thread running `thread_fn` may still call `sema_up` before the thread running `main` calls `sema_down`. What difference between the two makes `compare-sema.c` work correctly while `compare-cond.c` misbehaves sometimes?

As you have likely noticed from above, the crucial difference between semaphores and condition variables is that condition variables simply do nothing if `cond_signal` is called and there is no thread to wake. As such, if a thread later calls `cond_wake`, it will still have to wait. In contrast, the semaphore remembers that `sema_up` was called by incrementing its internal counter. That way, if a thread later calls `sema_down`, it does not have to wait since the counter is at 1. One way to think of this difference is that a semaphore maintains more state than a condition variable (the semaphore also needs to keep track of the number of threads waiting, it is just not shown in Progviz).

To compensate for the fact that condition variables lack the state that semaphores have, we need to maintain that state ourselves. Since this state will be represented as some form of shared data, we need a lock to protect that data. This is the reason why a condition variable needs to be paired with a lock — we need to maintain some external state to use the condition variable correctly anyway. Maintaining this state manually does of course mean it takes a bit more effort to use condition variables. The great benefit, however, is that we as users of the condition variable get to choose how this additional state is represented and how it is updated. If we use a semaphore, the state has to be a counter that we can only increment and decrement. If we instead use a condition variable, the state can be whatever we want and we can update it however we want! This is what makes condition variables much easier to use than semaphores for certain problems.

8.3 Waiting Until a Condition is True

Now that we know that we have to maintain additional state to use condition variables correctly, the next step is to examine how we maintain the state properly. For once, there is actually a fairly straight-forward method that dictates how to maintain this state. One way to describe this method is to convert a non-synchronized program that utilizes while-loops to wait into a properly synchronized program that uses locks and condition variables. We will use the programs `wait-for-two-x.c` to illustrate this process by example, and then summarize it into four concrete steps at the end of the section.

```
int result_a;
int result_b;
int threads_running;

void thread_fn_a(void) {
    result_a = 2;
    threads_running--;
}

void thread_fn_b(void) {
    result_b = 3;
    threads_running--;
}

int main(void) {
    threads_running = 2;

    thread_new(&thread_fn_a);
    thread_new(&thread_fn_b);

    while (threads_running > 0)
        ;

    assert(result_a == 2);
    assert(result_b == 3);
    return 0;
}
```

Listing 8.3: The program `wait-for-two-1.c`.

Listing 8.3 shows the program `wait-for-two-1.c`, which is the initial version of the program. The `main` function launches two threads that run `thread_fn_a` and `thread_fn_b` respectively, and then waits for both of them to complete using a `while` loop. Each of the threads set `result_a` and `result_b` respectively, and tell the main thread that they are done by incrementing `threads_running`. As you can see, the program is currently not synchronized, and since data is shared without properly protecting it, it contains data races. At this point, you are probably confident enough to synchronize the program using semaphores. However, note that doing this requires some creativity, since it is not possible to simply convert `threads_running` to a semaphore and replace `threads_running--` with `sema_down`. As we shall see, using locks and condition variables does not

require altering the program logic at all.

The first step is to identify shared data and the associated critical sections. By examining `wait-for-two-1.c`, we find that the variables `result_a`, `result_b`, and `threads_running` are shared between threads. However, assuming that the logic `main` uses to wait for the two threads to complete works correctly (which is not the case currently, but we will fix it soon), only `threads_running` will actually be accessed concurrently since it ensures that both threads have written to `result_a` and `result_b` by the time the main thread reads from them.

The critical sections for `threads_running` are straight-forward to find. The program contains a total of three critical sections. Two of them are around the two occurrences of the line `threads_running--`. The third and final one surrounds the `while` loop that waits until `threads_running` is zero.

```
int result_a;
int result_b;
int threads_running;
struct lock threads_running_lock;

void thread_fn_a(void) {
    result_a = 2;
    lock_acquire(&threads_running_lock);
    threads_running--;
    lock_release(&threads_running_lock);
}

void thread_fn_b(void) {
    result_b = 3;
    lock_acquire(&threads_running_lock);
    threads_running--;
    lock_release(&threads_running_lock);
}

int main(void) {
    lock_init(&threads_running_lock);
    threads_running = 2;

    thread_new(&thread_fn_a);
    thread_new(&thread_fn_b);

    lock_acquire(&threads_running_lock);
    while (threads_running > 0)
        ;
    lock_release(&threads_running_lock);

    assert(result_a == 2);
    assert(result_b == 3);
    return 0;
}
```

Listing 8.4: The program `wait-for-two-2.c`, where the critical sections are protected with locks.

We then protect the critical sections using locks. Since we only have one variable to protect, we only need one lock. As before, we name the lock `threads_running_lock` to remind ourselves that it protects the variable `threads_running`. Then we simply insert calls to `lock_acquire` and `lock_release` around the three critical sections we found. This results in the code in `wait-for-two-2.c` showed in Listing 8.4.

Interestingly, this version of the program might seem worse than the original version. In particular, if the main thread reaches the `while` loop before `threads_running` is zero, the program will get stuck in the loop. The original version at least had the merit that it *appeared* to work properly even in this case. However, as we saw in Chapter 3 the original program contains data races and may thereby behave as poorly as `wait-for-two-2.c`. The synchronization simply makes this issue readily visible by making the program well-defined.

You can see this issue in Progvis, either by manually stepping the main thread to the `while` loop without touching the other threads. Note that Progvis will not report the error explicitly, but you will note that the main thread gets stuck in the while loop. Progvis does, however, detect and report the error if you select *Run* \Rightarrow *Look for errors...*. In this case, Progvis identifies the issue as a *livelock*, which is when one or more threads do work but do not make any progress in program execution. When Progvis shows you the cycle, you can click *Close* in the green bar at the top of the window to take control of program execution if you wish to experiment.

The reason for this behavior is, as you might already have realized, that the main thread holds the lock across the entire `while` loop. This means that neither of the two other threads may acquire `threads_running_lock` and thereby they can not modify the `threads_running` variable. This means that the main thread has prevented `threads_running` from changing, even though it is waiting for it to change!

```
lock_acquire(&threads_running_lock);
while (threads_running > 0)
    cond_wait(&threads_running_cond,
             &threads_running_lock);
lock_release(&threads_running_lock);
```

Listing 8.5: Adding `cond_wait` to the loop to avoid livelock, available as `wait-for-two-3.c`.

We can solve this issue using a condition variable. We define a new condition variable next to the lock and name it `threads_running_cond` to remind ourselves that we should signal the condition variable when `threads_running` changes. Once we have a condition variable, we can insert a call to `cond_wait` into the body of the loop that waits for `threads_running` to become 0. This makes the loop look like in Listing 8.5 (also in `wait-for-two-3.c`).

This solves the livelock issue we saw before since `cond_wait` releases the lock while the thread is waiting. One way to look at this is to consider `cond_wait` as a pause in the critical section that is delimited by `lock_acquire` and `lock_release`. Since the lock is released, other threads are allowed to acquire the lock and modify `threads_running`. However, we are still not done. If we use *Run* \Rightarrow *Look for errors...* in Progvis at this time, it will report a deadlock.

What it found was that if the main thread reaches the while loop before both threads are done, it will call `cond_wait`. Since no other thread calls `cond_signal` or `cond_broadcast`, it will never wake up again.

This illustrates that we have forgotten to signal the condition variable when the program modifies `threads_running`. In this case, this happens at the end of the functions executed by the two threads. To solve the issue we can thereby add a call to `cond_signal` or `cond_broadcast` after the line `threads_running--` in both functions. Since `cond_wait` is called in a loop (which is good practice regardless, as we shall see), calling `cond_broadcast` is always a safe option. However, doing so might be wasteful as it wakes *all* threads. If we know that we only need to wake up *one* thread, and it does not matter which one (i.e., we do not care which thread the implementation selects), we can use `cond_signal` instead. In this case, the main thread is the only thread that will wait on the condition variable, so `cond_signal` is sufficient. This gives us the code in `wait-for-two-4.c`, which works correctly.

```
void thread_fn_a(void) {
    result_a = 2;
    lock_acquire(&threads_running_lock);
    threads_running--;
    cond_signal(&threads_running_cond,
               &threads_running_lock);
    lock_release(&threads_running_lock);
}
```

Listing 8.6: Adding `cond_signal` to wake the main thread. Available as `wait-for-two-4.c`.

Note that we want to call `cond_signal` or `cond_broadcast` right after we have modified one or more variables that some thread uses the condition to wait for. Since the variables are shared, they need to be protected by the same lock that protects the condition variable. Because of this, the thread that calls `cond_signal` or `cond_broadcast` will always hold the lock just before the call anyway. As such, it is a good idea to keep holding the lock during the call to `cond_signal` and `cond_broadcast` (and required in the library provided with this book), since it means that no other thread will invalidate our assumptions before signaling the condition.

Practice: In `wait-for-two-4.c` we have retained the `while` loop around the call to `cond_wait`. This might seem unnecessary since it means that the main thread will check the condition again after waking up. We can avoid this by replacing the `while` loop with an `if` statement as shown below and in `wait-for-two-error.c`. This will, however, cause the program to fail. Why?

```
lock_acquire(&threads_running_lock);
if (threads_running > 0)
    cond_wait(&threads_running_cond,
             &threads_running_lock);
lock_release(&threads_running_lock);
```

You can use *Run \Rightarrow Look for errors...* to find an interleaving that illustrates the error. Is there another way in which we can solve the issue in `wait-for-two-error.c`?

As you have seen from above, the issue is that the two threads always call `cond_signal` when they have decremented `threads_running`. This means that the first one will decrease `threads_running` to 1 and then call `cond_signal`. If the main thread has reached `cond_wait` at this time, this causes it to wake up even though `threads_running` is not yet zero. This is sometimes called a *spurious wakeup*, since the thread was woken even though the condition was not yet fulfilled. If we use an `if` statement instead of a `while` loop, the main thread will simply continue its execution when this happens. If you try the same interleaving in `wait-for-two-4.c` in Progviz,⁴ you will see that the main thread wakes up, but rather than continuing it checks the condition again, realizes that it is not yet time to continue, and goes back to sleep by calling `cond_wait` again.

In this case, we could avoid the issue altogether by only calling `cond_signal` when `threads_running` reaches zero. However, this means that we need to duplicate the condition that the main thread is waiting for, which makes it more difficult to change the condition in the future. As such, this is only possible to do if a condition variable is used to wait for *one* thing. Furthermore, not calling `cond_wait` in a loop also means that it is not always safe to call `cond_broadcast` rather than `cond_signal` to wake threads, which is not always possible,⁵ but also breaks the expectation that it is safe to replace `cond_signal` with `cond_broadcast` without affecting the correctness of the program. Furthermore, evaluating a condition is typically very cheap in comparison to waking a thread anyway, so the cost of the loop is often negligible. For all of these reasons, it is good practice to always call `cond_wait` in a loop. In fact, condition variables in C++ have a member function that checks the condition (provided as a lambda function) in a loop for you, since the pattern is so common.

For the curious reader, the file `wait-for-two-duplicated.c` contains an implementation that avoids waking the main thread more than once, and thereby does not need the loop around `cond_wait`. However, as mentioned above and by

⁴That is: (1) step the main thread until it reaches `cond_wait`, (2) let one of the other threads finish, and (3) continue stepping the main thread.

⁵For example, it is not possible to wake *a particular* thread without waking all of them and have the threads themselves figure out which one should continue

the comment in the program, note that this type of optimization is often only necessary in “hot code”, that is known to be the bottleneck of a concurrent program.

8.3.1 Summary of the Method

The method above is actually general enough to be applicable to arbitrary problems. As such, we can summarize what we did above into four general steps that we can use to turn a non-synchronized program that utilizes loops to wait into a properly synchronized program that uses locks and condition variables.

1. Examine the program and identify (1) the shared data, (2) the critical sections that manipulate the shared data, and (3) loops that wait for a condition to become true. Note that the condition will involve shared data, and thereby examine the state of the shared data.
2. Use one or more locks to protect the critical sections. Note that all data that are used in the same conditions need to be protected by the same lock (this will likely be the result anyway, but it is worth considering as it may help the implementation).
3. Add a condition variable for each set of variables that the program is waiting for, and call `cond_wait` in the loop that waits until the condition becomes true. Note that the call to `cond_wait` needs to be a part of the same critical section that checks the condition.
4. Find all places where the program changes variables associated with each condition. Call `cond_broadcast` or `cond_signal` after they are updated, but before the lock is released. Using `cond_broadcast` is always safe, but may wake up too many threads. If only one thread needs to wake, and the implementation does not care *which* out of all threads that wait on the same condition, then `cond_signal` can be used.

8.4 A More Complex Condition

To illustrate some more nuances with the method above, we will next look at the program `seats-x.c`, which contains a more complex condition. The first version of the program (`seats-1.c` and Listing 8.7) is again not synchronized at all. The program manages a number of seats (e.g., seats at a cinema). The variable `seat_taken` is an array with `num_seats` elements that stores whether each seat is currently occupied or not. The function `seat_acquire` searches the `seat_taken` array to find a seat that is not occupied, marks it as taken, and returns the index of the seat. If no seat is available, the function waits until one becomes available. Once the user of the seat is finished, they are expected to call `seat_release` to mark it as free again.

The program additionally contains a variable `seat_used_count` that the `thread_fn` (Listing 8.8) function is used. Once it has acquired a seat, it increases the value for that seat, to keep track of how many times each seat has been used. Since `seat_acquire` and `seat_release` have semantics similar to a lock, `thread_fn` uses them to manage access to the used count for the seats,

and this variable therefore does not need additional protection. You may have noted that this program essentially implements a lock that manages access to multiple instances of a resource that are treated as equivalent as discussed in Section 7.1.2.

The rest of the code in the main program initializes the arrays to contain 2 seats that are both initially not taken. It then starts two threads that both call `thread_fn` and then calls `thread_fn` itself. As such, a total of three threads will attempt to acquire a seat concurrently, which means that one of them will have to wait. After the call to `thread_fn`, the main thread waits for all threads to finish and verifies that the two elements in `seat_used_count` sums to 3.

```
bool *seat_taken;
int *seat_used_count;
int num_seats;

int seat_acquire(void) {
    int taken = -1;
    while (taken == -1) {
        for (int i = 0; i < num_seats; i++) {
            if (seat_taken[i] == false) {
                seat_taken[i] = true;
                taken = i;
                break;
            }
        }
    }
    return taken;
}

void seat_release(int id) {
    seat_taken[id] = false;
}
```

Listing 8.7: The initial version of the `seats` program.

```
int seat = seat_acquire();
seat_used_count[seat]++;
seat_release(seat);
```

Listing 8.8: The code that uses `seat_acquire` and `seat_release` in `seats-1.c`. Note that they have similar semantics to a lock.

Even though we can formulate what `seat_acquire` is waiting for (i.e., it is waiting for one seat to be available), the condition is perhaps not immediately apparent in Listing 8.7. To help our understanding, we start by refactoring the code as shown in Listing 8.9. In particular, we move the `for` loop in `seat_acquire` into its own function that we call `grab_seat`. We mark the function as `static` to remind ourselves that the function is not supposed to be called outside of `seat_acquire`.

```
static bool grab_seat(int *taken) {
    for (int i = 0; i < num_seats; i++) {
        if (seat_taken[i] == false) {
            seat_taken[i] = true;
            *taken = i;
            return true;
        }
    }
    return false;
}

int seat_acquire(void) {
    int taken = -1;
    while (!grab_seat(&taken))
        ;
    return taken;
}

void seat_release(int id) {
    seat_taken[id] = false;
}
```

Listing 8.9: Refactoring of the `seats` program to make it easier to see what the program is waiting for.

By moving the `for` loop into `grab_seat` the condition that the program waits for becomes clearer, since we can put the function call in the header of the `while` loop inside `seat_acquire`. Note that we do not need to refactor the code in this way to apply condition variables. It just makes it easier to see what is happening, and helps making the discussion in this section clearer. We will provide a non-refactored version of the solution at the end of the section in addition to the refactored version.

Now that we have refactored the code, we can start apply the steps from Section 8.3.1. The first steps asks us to identify (1) the shared data, (2) the critical sections, and (3) loops that wait for a condition to become true. Our answer to (1) is that the shared data is the elements of the `seat_taken` array, and that (2) they need to be protected in the if-statement in `grab_seat` as well as in `seat_release`. Finally, (3) the program waits for at least one seat to be available in the `while` loop in `seat_acquire`.

Using this information, we can then move on to step 2 from Section 8.3.1 and protect the shared data using locks. Since our initial impression is that seats can be treated separately, we define an array of locks so that we can use one lock for each individual seat. This gives us the program `seats-3.c` (Listing 8.10).⁶

⁶You might have noticed that this disregards an important detail. This is intentional to illustrate *why* this detail is important. We will get back to it!

```

static bool grab_seat(int *taken) {
    for (int i = 0; i < num_seats; i++) {
        lock_acquire(&seat_taken_lock[i]);
        bool take = seat_taken[i] == false;
        if (take) {
            seat_taken[i] = true;
            *taken = i;
        }
        lock_release(&seat_taken_lock[i]);
        if (take)
            return true;
    }
    return false;
}

int seat_acquire(void) {
    int taken = -1;
    while (!grab_seat(&taken))
        ;
    return taken;
}

void seat_release(int id) {
    lock_acquire(&seat_taken_lock[id]);
    seat_taken[id] = false;
    lock_release(&seat_taken_lock[id]);
}

```

Listing 8.10: Protecting the critical sections using an array of locks.

Practice: As we would expect, the program does not yet work as intended since we still have 2 more steps to follow. However, if you run the program (either in the terminal or in Progviz) you will see that the program actually behaves correctly at this point. In particular, the third call to `seat_create` does not cause a deadlock even if all seats are occupied.

Investigate the behavior of the program using Progviz to understand *why* `seats-3.c` does not deadlock even though `wait-for-two-2.c` caused a deadlock in the same situation.

The program is not correct, however. If you select *Run* \Rightarrow *Look for errors...* in Progviz, it will tell you that the program uses *busy wait* (it needs to investigate around 54 000 transitions to reach this conclusion, so it takes a while). What is the problem?

As you have noted above, the current version of the program surprisingly works as expected, even though the same stage of the previous program (`wait-for-two-2.c`) deadlocked in certain situations. The reason for this is that `seats-3.c` does *not* hold a lock throughout the *entire* loop it uses to wait for a seat to become free. Since it makes sure to release the lock periodically, other threads are still able to release their seats. This was not the case in `wait-for-two-2.c`, which did not release the lock inside the loop. We could cause `wait-for-two-2.c`

to behave the same way by inserting `lock_release` directly followed by `lock_acquire` in the body of the `while` loop.

Even though the program does work as expected, it is of course not ideal. As Progviz reports, the program uses *busy wait* to wait for the condition to become true. In this case, if a thread finds that no seats are available, it will keep checking for an available seat until one becomes available, thereby essentially wasting CPU time to do nothing. While this is fine to do for a short amount of time,⁷ it is problematic if it is done for a longer time or at multiple places in the program. Since the scheduler does not know which threads do useful work and which threads are just wasting CPU time to wait, it is unable to prioritize the threads that actually do useful work. As such, threads that busy wait for a considerable amount of time reduce the available CPU time for the threads that do useful work,⁸ which means that it will take longer for them to finish their work. Because of this, it is good practice to use synchronization primitives to wait. They contain logic to inform the scheduler that the thread is waiting, which allows it to not waste CPU time in this way.

To fix this issue, we move to step 3 from Section 8.3.1 and add a condition variable to help us wait efficiently. We would like to call `cond_wait` in the `while` loop in `seat_acquire`. However, this is not currently possible as we do not hold a lock in the while loop. We also need a lock to protect the condition variable itself.⁹ One way to solve the issue is to add a separate lock that protects the condition variable and allows us to call `cond_wait`. We call the condition variable `seat_free_cond` since we wait for a seat to become free, and we call the new lock `seat_free_lock` to help us associate the two together. With these additions we can insert the call to `cond_wait` in `seat_acquire`. While we are at it, we can insert a call to `cond_signal` in `seat_release` as instructed by step 4 as well. This gives us the code in `seats-4.c` (Listing 8.11).

Practice: The program `seats-4.c` is still not correct. There are situations that cause the program to deadlock. Use Progviz to find the issue. If you use *Run* \Rightarrow *Look for errors...*, note that it needs to investigate around 100 000 transitions to find the issue, so checking will take some time.

⁷For example, one implementation of a lock is a *spinlock*, which utilizes busy wait. It is, however, best suited for situations where we know that threads never have to wait for a significant amount of time.

⁸Or prevent the CPU from becoming idle and thereby consume less energy.

⁹This is of course an indication that we have done something wrong, as we shall see.

```

int seat_acquire(void) {
    int taken = -1;
    while (!grab_seat(&taken)) {
        lock_acquire(&seat_free_lock);
        cond_wait(&seat_free_cond, &seat_free_lock);
        lock_release(&seat_free_lock);
    }
    return taken;
}

void seat_release(int id) {
    lock_acquire(&seat_taken_lock[id]);
    seat_taken[id] = false;
    lock_release(&seat_taken_lock[id]);
    lock_acquire(&seat_free_lock);
    cond_signal(&seat_free_cond, &seat_free_lock);
    lock_release(&seat_free_lock);
}

```

Listing 8.11: The `seat_acquire` and `seat_release` functions after adding a condition variable. Available as `seats-4.c`.

As you have probably noted, it is possible for the program to deadlock since `cond_wait` is in a *different* critical section to the code that checks if any seats are available. Since the program does not hold any locks at the end of `grab_seat`, a seat may become available after `grab_seat` returns but before `cond_wait` is called. In particular, assume that all seats are occupied and a thread A calls `seat_acquire`. Thread A calls `grab_seat` which returns `false` since no seats are available. Before it acquires `seat_free_lock`, thread B calls `seat_release` which marks one of the seats as available and signals the condition variable. When Thread A continues again, it acquires the lock and calls `cond_wait`. However, since it called `cond_wait` while a seat was available, the condition variable will not be signaled again, and the thread may have to wait forever.

This observation shows that it is necessary to treat condition variables as any other variables when we look for critical sections and consider which locks to use when protecting shared data. As we saw above, `cond_wait` needs to be called in the same critical section where we verified that the condition was false. The same is often true for calls to `cond_signal` and `cond_broadcast`,¹⁰ even if they are forgiving since an extra wakeup rarely hurts correctness, only performance. This is the reason why we want to treat condition variables *as if* they are not thread-safe, and protect them with a particular lock. This also shows that we should treat the condition variable as a part of the condition we wish to wait for, which is the reason why step 2 in Section 8.3.1 states that “Note that all data that are used in the same conditions need to be protected by the same lock.”

As such, since we have one condition variable that waits for one of multiple

¹⁰In particular, it is only important if we do not always call `cond_signal` or `cond_broadcast`. The logic that determines if we need to signal the condition variable likely depends on shared state, and if we break the critical section, other threads may have modified the state before we inspect it.

seats to become available, we need to revise our synchronization to use a single lock for all seats, rather than separate locks for each seat. This way we can use the same lock to protect the condition variable and all data involved in checking the condition. In a way, this is similar to what happened when we added `accounts_total` to the `bank` program in Section 6.4.2. Since `accounts_total` required stronger guarantees from the rest of the program, it required us to revise our synchronization approach. Here, the requirement existed from the start but it was not immediately obvious until we added the condition variable.

To fix the problem, we make `seat_taken_lock` into a single lock instead of an array and remove `seat_free_lock` since it is no longer necessary. We then revise the critical sections in `seat_acquire` with the assumption that we need to have a consistent view of *all* seats, not just individual seats at a time. If we do this, we will end up with the program in `seats-5.c` (Listing 8.12), which works as intended.

```

bool *seat_taken;
int *seat_used_count;
int num_seats;
struct lock seat_taken_lock;
struct condition seat_free_cond;

static bool grab_seat(int *taken) {
    for (int i = 0; i < num_seats; i++) {
        if (seat_taken[i] == false) {
            seat_taken[i] = true;
            *taken = i;
            return true;
        }
    }
    return false;
}

int seat_acquire(void) {
    int taken = -1;
    lock_acquire(&seat_taken_lock);
    while (!grab_seat(&taken))
        cond_wait(&seat_free_cond, &seat_taken_lock);
    lock_release(&seat_taken_lock);
    return taken;
}

void seat_release(int id) {
    lock_acquire(&seat_taken_lock);
    seat_taken[id] = false;
    cond_signal(&seat_free_cond, &seat_taken_lock);
    lock_release(&seat_taken_lock);
}

```

Listing 8.12: Correct synchronization of the `seats` program. Available as `seats-5.c`.

As mentioned previously in the section, it is not necessary to refactor the

program to add condition variables. The refactoring does, however, make it easier to see what part of the code is the condition, and what part is the loop that waits for the condition to become true. For reference, a non-refactored version of the program is available as `seats-no-refactor.c` and the relevant parts are shown in Listing 8.13. Note that it is significantly harder to follow the control flow of the program to mentally verify that it is correct.

```
int seat_acquire(void) {
    int taken = -1;
    lock_acquire(&seat_taken_lock);
    while (taken == -1) {
        for (int i = 0; i < num_seats; i++) {
            if (seat_taken[i] == false) {
                seat_taken[i] = true;
                taken = i;
                break;
            }
        }
        if (taken == -1)
            cond_wait(&seat_free_cond, &seat_taken_lock);
    }
    lock_release(&seat_taken_lock);
    return taken;
}
```

Listing 8.13: Correct synchronization of the non-refactored version of the `seats` program. Available as `seats-no-refactor.c`.

8.5 Exercises

1. The file `sema.c` contains an unfinished implementation of a semaphore. Apply the method in Section 8.3.1 to add locks and condition variables to make it work properly. You can use the provided main program to test your implementation in Progvis.
2. The file `dispatch.c` contains logic to share data safely between threads. In particular, one thread calls `put` to save a value that other threads may later retrieve with `get`. Data is stored in the variable `stored`. When `stored` contains a negative number, it is considered to be “empty”. The function `get` further lets the caller wait for `stored` to contain either an even or an odd number.

The program is currently not synchronized at all. Use `stored_lock` and `stored_cond` to make the program behave correctly. Remember to think about whether you need to use `cond_signal` or `cond_broadcast`.

3. The file `multi-lock.c` contains an implementation of a set of “locks” that can be acquired and released together without the risk of deadlock. For example, one thread can ask to acquire the lock numbered 0 and 1, while another can ask to acquire the locks numbered 1 and 0 without risking deadlock.

One way to achieve this is to represent each lock as a boolean (`acquired`), and use a global lock and a condition variable to ensure that only one thread attempts to acquire locks at the same time. A starting point of this idea is implemented in `multi-lock.c`, but the synchronization is non-existent.

Use `ack_lock` and `ack_cond` to synchronize the implementation. You can verify your implementation using Progvis.

4. The file `bank.c` contains the final version of the bank program used in Chapters 6 and 7. One shortcoming of the program is that `accounts_total` can not be called from multiple threads concurrently, even though there is no problem in doing so since it only reads from the accounts. Modify the code so that both (1) multiple threads can call `transfer` concurrently and (2) multiple threads can call `accounts_total` concurrently. However, note that `accounts_total` can not be executed concurrently with `transfer`, so your solution must make threads wait for each other if that is the case.

You can verify that your solution behaves correctly using *Run \Rightarrow Look for errors...* This does not verify that `transfer` and `accounts_total` can be executed concurrently. You need to do that manually. Note that the sample solution needs to explore around 170 000 transitions, so verification takes a while, especially if you solution is more complex than the sample solution.

Abstractions and Concurrency

At this point, you hopefully have a good understanding of what problems arise in concurrent programs (i.e., order of execution and shared data), and how to use synchronization primitives (semaphores, locks, and condition variables) to address these issues. However, up until now we have reasoned about the behavior of programs at a quite detailed level. While these detailed insights are crucial to reason about small portions of programs, it is infeasible to reason about large programs at the this level of detail.

It is worth noting that you have probably experienced this situation previously. When learning programming (i.e., “normal” sequential programming), one usually starts by focusing on the semantics of individual constructs in the programming language, and how they can be combined into interesting programs. However, at some point the programs become large enough that we can not treat them as a sequence of language constructs with some behavior. To be able to reason about the program, we need to introduce additional structure by introducing *abstractions*, for example by encapsulating a part of the code into a *function*, bundling associated data into a *data structure*, or creating *classes* or *modules* with associated functionality.

Now that we are at the same point with our understanding of *concurrent programming* and *synchronization*, the aim of this chapter is therefore to investigate how our insights from previous chapters interact with the tools we usually use for abstraction, and to amend these tools so that they work well for concurrent programs as well. As such, the goal is to be able to create and use abstractions correctly in concurrent programs.

A central term when working with abstractions in a concurrent system is *thread safety*. You have likely at least encountered the term previously, for example if you have read the documentation for a function in the standard library,¹ or for some other library. At a high level, functions or data structures that are *thread safe* can be used from multiple threads concurrently. However, as we will see in the rest of this chapter, the concept of thread-safety involves much more nuance than simply stating that some parts of the system are thread safe while others are not. The aim is to illustrate these nuances with examples, and then summarize our findings at the end of the chapter.

¹For example, if you open the manpage for `strtok` (`man strtok`) you will see that `strtok` is marked as `MT-Unsafe` while `strtok_r` is marked `MT-Safe`.

9.1 Functions

The most fundamental abstraction tool we have at our disposal in C and most other language is perhaps the ability to create functions. As such, we will start by examining how the concept of *thread safety* applies to functions with the help of the programs `count-length-n`. All programs contain the function `count_length` that helps programs keep track of the number of characters outputted somewhere.² The most basic version of the implementation is available as `count-length-1.c` (Listing 9.1).

```
int total_length;

int count_length(const char *string) {
    if (string == NULL) {
        total_length = 0;
    } else {
        total_length += strlen(string);
    }
    return total_length;
}

int main(void) {
    count_length(NULL);

    printf("first");
    count_length("first");

    printf("second");
    int total = count_length("second");

    printf("\nTotal: %d\n", total);

    return 0;
}
```

Listing 9.1: The program `count-length-1.c` showing the implementation and usage of `count_length`.

From Listing 9.1 we can see that `count_length` has two different modes of operation. If `NULL` is passed as the parameter to the function, it resets `total_length` to zero. Otherwise, it adds the length of the string to `total_length`. The `main` function illustrates the intended usage of the function. It starts by calling `count_length(NULL)` to reset `total_length` in case some other part of the system `count_length` previously. After that it prints two strings (`first` and `second`) to standard output. To keep track of the number of characters printed, it calls `count_length` after each call to `printf`. The value returned from `count_length` corresponds to the total number of characters printed since the last time `count_length(NULL)` was called. In this case, `main` only does this at the end of the program to find the total number of characters printed. Of course, this

²This is useful for example when printing data in a tabular format, or when it is necessary to store the size of a data block in a binary format.

structure is not very useful in a small and simple program like this. We can quite easily see that `total` will always be 11 in this case. However, imagine that `main` is more complex (e.g., by calling many levels of functions).

While the implementation in `count-length-1.c` works well for sequential programs, it quickly breaks down when we introduce multiple threads. This is illustrated by the program `count-length-2.c` (Listing 9.2), which attempts to use `count_length` from two different threads. Note that for clarity, the calls to `printf` are removed from this point and onward for clarity. Imagine that `main` and `thread` attempt to write characters to different locations, for example two different files.

```
int total_length;

int count_length(const char *string) {
    if (string == NULL) {
        total_length = 0;
    } else {
        total_length += strlen(string);
    }
    return total_length;
}

void thread(void) {
    count_length(NULL);

    count_length("first");

    int total = count_length("second");
    printf("Thread: %d\n", total);
}

int main(void) {
    thread_new(&thread);

    count_length(NULL);

    int total = count_length("main");
    printf("Main: %d\n", total);

    return 0;
}
```

Listing 9.2: The program `count-length-2.c` calling `count_length` from two different threads.

Practice: As you have probably already realized, the program `count-length-2.c` does not behave correctly. In particular, it contains a data race. Where is the data race? You can probably locate it by yourself at this point. You can also use *Run* \Rightarrow *Look for errors...* in Progvis to help you.

9.1.1 Thread Safety

As you have probably noticed, `count-length-2.c` uses the function `count_length` concurrently from multiple threads. This is not a problem in and of itself, but since `count_length` modifies `total_length`, using it in this way leads to a *data race*. Because of this, we say that this implementation of `count_length` is **not thread safe**. This simply means that it is *not* correct to call the function from multiple threads without synchronization to ensure that the calls never happen concurrently.

One large benefit of classifying functions as *thread safe* and *not thread safe* is that we can give a high-level description of the function without giving all the details about the implementation. In this particular case, it is quite easy to see that it is problematic to call `count_length` concurrently since it uses a global variable. However, all that *users* of `count_length` really need to know is that it is *not thread safe*. This is one important step towards managing complexity by hiding details and only considering the high-level important aspects of an abstraction (i.e., a function in this case).

The notion of thread safety is also useful to determine *how* to properly fix a concurrency issue. In this case, if the intention is for `count_length` to *not* be thread safe, then we can quickly conclude that the bug in `count-length-2.c` is that `main` and `thread` call `count_length` without proper synchronization. However, if the intention was that `count_length` *should be* thread safe, then the bug is that `count_length` is not thread safe as we intended.

One way to think of this is to consider the function declaration as a contract between the caller and the callee. Among other things, this contract includes what parameters the caller should pass to the function and what return values to expect. This is the same as in sequential programming, so you are likely already familiar with the idea. What is new to concurrent programming is that this contract also includes whether or not the function is thread safe. As such, if the contract for `count_length` states that it is *not* thread safe, then `main` and `thread` breaks the contract by calling `count_length` without synchronization. On the other hand, if the contract states that `count_length` *is* thread safe, then `count_length` breaks the contract by not being thread safe.

Since the intended behavior of `count_length` is not stated, we first need to determine whether we *wish* `count_length` to be thread safe or not.³ In many cases one option is better than the other. In this case, we will solve the issue by making `count_length` thread safe.⁴

9.1.2 Synchronizing `count_length`

To make `count_length` thread safe, we can apply what we have learned in Chapter 6. First, we look for shared data and find that the variable `total_length` is shared between multiple threads. Secondly, we find places in the code where `total_length` is used (i.e., where shared data is used) and identify critical sections. In this case, the critical section includes writing `total_length` in both

³In other words, we need to decide if we want to blame *users* of `count_length` or if we want to blame `count_length` itself.

⁴As we will see later, making `count_length` thread safe is the best option. Modifying `main` and `thread` to ensure that they work properly with a `count_length` that is *not* thread safe will remove most parallelism in the program. Are you able to see why that is already?

branches of the `if` statement (both setting it to zero and increasing it), as well as reading it as a part of the `return` statement. Once we have identified the critical sections, we can protect them using a lock. In this case there is only one critical section we need to protect. It includes the entire `if` statement in the `return` statement.⁵ and the read of `total_length`. We protect it using the lock `total_lock`. Note that we need to rewrite `count_length` a bit so that we can include the final read of `total_length` in the critical section (we can not release the lock after `return`, since code after `return` is not executed). These changes result in the code in Listing 9.3.

```
int total_length;
struct lock total_lock;

int count_length(const char *string) {
    lock_acquire(&total_lock);
    if (string == NULL) {
        total_length = 0;
    } else {
        total_length += strlen(string);
    }
    int result = total_length;
    lock_release(&total_lock);
    return result;
}
```

Listing 9.3: Synchronized version of `count_length` from `count-length-3.c`.

These changes to `count_length` do indeed avoid data races, and we could therefore consider it to be thread safe. However, as we have seen before, adding locks around the problematic parts of the program do not necessarily make the program behave the way we expect it to. As it turns out, this is true here as well: even though we do not have any data races, the program will not behave as we expect it to.

⁵The the condition (i.e., `string == NULL`) is not technically a part of the critical section. However, since the condition is very cheap, the benefit of keeping the implementation simple vastly outweigh the slight loss of parallelism.

Practice: Examine the code in the functions `main` and `thread` in `count-length-3.c` (depicted below for convenience), and write down your expectations about the value of the variable `l` in each of the two functions should contain, preferably in the form of one assertion for each function (e.g., `assert(l == 0);`).

```

void thread(void) {
    count_length(NULL);

    count_length("first");

    int l = count_length("second");
    printf("Thread: %d\n", l);
}

int main(void) {
    lock_init(&total_lock);
    thread_new(&thread);

    count_length(NULL);

    int l = count_length("main");
    printf("Main: %d\n", l);

    return 0;
}

```

Once you have written down your assertions, consider whether the conditions in the assertions are always true, or if some interleaving causes them to be false. You can use *Run \Rightarrow Look for errors...* in Progviz to verify your findings.

The way the program in `count-length-3.c` is written, especially if we keep in mind that we imagine that the two functions write data to separate files, we would assume that `l` in `main` will always get a value of 4, while `l` inside `thread` always gets a value of 11. However, as you have most likely realized by now, this is not always the case. The reason is that the variable `total_length` contains global state that is unintentionally shared between the two threads. Even if it does not constitute a data race, it makes the program behave incorrectly.

To illustrate a concrete possibility, assume that `main` first executes `count_length(NULL)` to set `total_length` to 0. After that, `thread` gets a chance to execute `count_length("first")` followed by `count_length("first")`. Then, `main` continues and calls `count_length("main")`, which returns the value 9 that is stored in `l`. Finally, `thread` continues to call `add_length("second")`, which returns the value 15 that is stored in `l`. As such, in this particular case the variable `l` in `main` became 9, while the variable `l` in `thread` became 15. Neither of them matched our expectations.

This problem is likely not too surprising in a small program like `count-length-3.c`, where it is quite easy to see the *entire* program at once. However, imagine that `clear_length` and `count_length` are used in a much larger program, where the `main` and `thread` are located in different source files. In such a

program, it is much more difficult to get a good enough overview of the program to understand these unintentional interactions between different parts of the program. Furthermore, since the issue is in a quite small part of the program, the program will likely behave correctly 99% of the time in spite of this problem. This means that when a user of the program eventually reports this issue, it will be very difficult to locate (the `count_length` function will be one of the last places you look, since it has seemed to work well for a long time even though it was always broken).

9.1.3 Solving the Problem

As we have seen above, simply adding locks to functions in an attempt to make them thread safe does not always make programs behave correctly. This is likely not too surprising, since the conclusion is very similar to our findings in Chapter 6. One important question does, however, remain. How *do* we solve the problem in `count-length-n` properly?

All our problems so far has been related to the shared variable `total_length`. First, it was the cause of a data race, and now it makes it possible for the two threads to affect each other. The proper solution to our problems is therefore to make `count_length` thread safe by removing the global variable `total_length` altogether and instead accept a parameter that contains all state that the function needs. This way each thread can supply their own instance of the state (i.e., `count_total`) that `count_length` needs.

As such, to solve the problem we define a structure, `struct state`,⁶ and move the previously global variables into the data structure. Finally, we add a parameter to `count_length` so that the caller can specify which instance of the global state the function should use.⁷

```

struct state {
    int total_length;
};

int count_length(struct state *s, const char *string) {
    if (string == NULL) {
        s->total_length = 0;
    } else {
        s->total_length += strlen(string);
    }
    return s->total_length;
}

```

Listing 9.4: Thread safe version of `count_length` from `count-length-4.c`.

Applying these changes to `count_length` gives us the code in Listing 9.4. Notably, `total_length` is no longer shared, so we no longer need the lock.

⁶Ideally, `count_state` or a more descriptive name, but we keep it short in order to fit the example code better onto the pages.

⁷In this particular case, we could have simply used an integer instead of `struct state`. However, this approach makes it clearer that the same idea also works when the state consists of multiple variables, and hiding the data in a struct makes it clearer for the user of `count_length` that the contents of the struct should not be accessed directly.

However, since we changed the interface of `count_length`, we also need to update the code that calls `count_length`. In particular, each of the two functions `main` and `thread` need a `struct` state variable that they can pass to `count_length`. Since the two functions use *separate* state variables, it is no longer possible for the two functions to affect each other. As such, we now always get the expected result, that 1 in `thread` contains 11 and that 1 in `main` contains 4. These changes are shown in Listing 9.5.

```

void thread(void) {
    struct state s;
    count_length(&s, NULL);

    count_length(&s, "first");

    int l = count_length(&s, "second");
    printf("Thread: %d\n", l);
}

int main(void) {
    thread_new(&thread);

    struct state s;
    count_length(&s, NULL);

    int l = count_length(&s, "main");
    printf("Main: %d\n", l);

    return 0;
}

```

Listing 9.5: Code that uses the thread safe version of `count_length` in Listing 9.4. Available in `count-length-4.c`.

At this point, the function `count_length` is properly *thread safe*, since it is safe to call it from multiple threads without issues. However, using this (arguably simple) definition one could argue that the version of `count_length` in `count-length-3.c` (Listing 9.3) is also thread safe. After all, the lock we used then avoided data races! This observation is indeed true. However, calling a function *thread safe* does not only imply that calling the function from multiple threads avoids data races. It also implies that *it is possible to use the function correctly* from multiple threads. This is the key difference between `count_length` in `count-length-3.c` and `count-length-4.c`. As we have seen, it is not possible to use the version in `count-length-3.c` correctly without additional synchronization. As such, we **do not** call it thread safe. However, the version in `count-length-4.c` is possible (and even quite easy!) to use correctly from multiple threads. Therefore, we say that it **is** thread safe.

9.1.4 Caveats of Thread Safe Functions

Now that we have a version of `count_length` that is thread safe (specifically the version in `count-length-4.c` and Listing 9.4), we might think that it is thread

safe no matter how we use it. This is, however, not the case.

Remember that we did *not* call the `count-length-3.c` version of `count_length` thread safe, even though no data races occurred even if it is used from multiple threads. The reason is that it would not be *useful* to call it thread safe since we saw that it is not possible to use it from multiple threads in practice, even though we avoided data races. This illustrates that the term *thread safe* does not correspond directly to “no data races possible”. Rather, the term *thread safe function* means “this function is possible to use correctly without proper synchronization”. As we have seen, this excludes some functions that fulfill the weaker condition of “no data races possible”.

There are also situations that can lead to data races that we willfully ignore when we are thinking about thread safety. Again, this is to make the concept of *thread safety* more useful. However, these implicit assumptions might be surprising at first, since it seems intuitive that a *thread safe function* is safe to call from multiple threads, regardless of *how* it is called.

```

void thread(struct state *s, struct semaphore *done) {
    count_length(s, NULL);

    count_length(s, "first");

    int l = count_length(s, "second");
    printf("Thread: %d\n", l);

    sema_up(done);
}

int main(void) {
    struct state s;
    struct semaphore done;
    sema_init(&done, 0);

    thread_new(&thread, &s, &done);

    count_length(&s, NULL);

    int l = count_length(&s, "main");
    printf("Main: %d\n", l);

    sema_down(&done);
    return 0;
}

```

Listing 9.6: The functions `thread` and `main` from `count-length-5a.c`. Note that the two threads share one `struct state` instance.

To illustrate these situations, consider the program in `count-length-5a.c`. This program is similar to `count-length-4.c`, but as shown in Listing 9.6, the two threads share the same `struct state` instead of having their own instance. The semaphore is needed to avoid Progviz complaining about using `s` after `main` terminates. This is a problem, but not the main point of this discussion. As

such, you can mostly ignore the semaphore for now.

Practice: The program `count-length-5a.c` is not correct. Use the same approach as earlier to find the problem: First write down your assumptions about the value of `l` in both `thread` and `main` as `assert` statements. Then find an interleaving that causes the assertions to be false, either by yourself or using *Run \Rightarrow Look for errors...* in Progvis.

As you have likely seen from above, `count-length-5a.c` has the same problem that we saw in `count-length-2.c`. We could, of course, add a lock to `struct state` and synchronize the critical section inside `count_length`, but that will still not solve the problems in `count-length-3.c`.

In spite of these problems, we call this version of `count_length` *thread safe*. The reason is that our implementation of `count_length` makes an implicit assumption: it assumes that *no other thread uses the same `struct state` at the same time*. This assumption is very common for functions that accept a parameter that represents some internal state, and is not always stated explicitly for each individual function. The reason for this is that it is still useful to think of `count_length` as thread safe, since the expectation is that different parts of the code use different `struct state`. One way to think of it is that `count_length` does not introduce any data races that are not apparent from its parameters.

```
char str[] = "main";

void thread(void) {
    str[2] = 0;
}

int main(void) {
    thread_new(&thread);

    struct state state;
    count_length(&state, NULL);

    int l = count_length(&state, str);
    printf("Main: %d\n", l);

    return 0;
}
```

Listing 9.7: The functions `thread` and `main` from `count-length-5b.c`. Note that `thread` modifies the contents of the string that `main` passes to `count_length`.

A second situation that is similar but not exactly the same is illustrated by the program `count-length-5b.c`. This program also uses the same version of `count_length` as in `count-length-4.c` (Listing 9.4). In `main`, the program calls `count_length` as before. However, this time it uses a global string, that the function `thread` (which runs concurrently with `main`) modifies. As you might guess, this does not end well since `strlen` called by `count_length` will read memory that is concurrently being modified by another thread, which constitutes a data race.

Again, in spite of this problem, we consider this version of `count_length` to be thread safe. Just like before, this is because we make implicit assumptions about the parameters passed to `count_length`. While this is overall similar to what we saw regarding `struct state` in the previous example, there are some differences in nuance that are worth discussing closer. The main difference between the `struct state` parameter and the string parameter is that we consider `struct state` and `count_length` to belong together. That is, `struct state` was made specifically for use with `count_length`. In an object oriented language (e.g., C++), we might have chosen to implement `count_length` as a member function of `struct state`. This is not true for the string parameter: there is no special relation between `count_length` and the string datatype other than that `count_string` wishes to take a string as an input parameter.

This nuance might initially seem insignificant, and to some extent it is: it is usually a good idea to document the what expectations functions make about their arguments regardless. However, this nuance affects our expectations of what is “normal” to do with the function. In the case of `count_length`, we expect it to modify `struct state` anyway it pleases since the two are connected tightly. This is even more true if you treat the contents of `struct state` as being something that users of `count_length` should not know about (i.e., `private` in languages like C++). Because of all of this, as a user of `count_length` we have to be pessimistic and assume that `count_length` both reads and writes the contents of the `struct state` we give it. As such, we do not expect it is possible to share `struct state` between different threads calling `count_length`. We have different expectations for the string. First and foremost we do not expect `count_length` to *change* the contents of our string. This is further emphasized by using `const` in the type of the parameter. This signals that `count_length` only *reads* from the data. As such, we expect that we pass any string to `count_length` as long as we do not *modify* the string while `count_length` is executing.

As you can see, much of the discussion around what is thread safe and what is not thread safe is based on whether or not some particular function *deviates from our expectations* of how we can use the function.⁸ This is perhaps what makes thread safety difficult. It is important to know the “reasonable expectations” so that we are able to interpret the available documentation in the proper context. This is why the distinction above is important.

In addition to the communication aspect covered in the previous paragraph, there is also a technical difference between the two parameters to `count_length`. Because `struct state` is tightly connected to `count_length`, it is possible that `count_length` contains logic for synchronizing access to `struct state` (e.g., by including some elaborate locking scheme that we do not know about). If this is the case, you will likely find documentation that explicitly states that it is possible to share `struct state` between threads even if we would not expect it to be possible. This is not possible for the string, since there is no standard way to associate the string representation (i.e., `const char *`) with a lock. As such, even if `count_length` *wanted* to avoid the issue in `count-length-5b.c`, there is no way we can do so, since there is no way of knowing which (if any) lock

⁸At least in “normal” documentation. The situation is slightly different in formal specifications, for example. It is of course always good to be explicit, but stating the same “default” constraints for all functions in your program will likely bore you quite quickly, so it rarely happens in practice.

that is used to protect the string that is given to us. As such, it is *necessary* for `count_length` to assume that no other thread modifies the string while it is running.

9.2 Abstract Data Types

At the end of the last section we saw that there are many nuances in the interaction between functions and data structures, even for functions that we consider thread safe. The goal of this section is therefore to extend the discussion to also include data in the context of abstract data types. As described in Section 2.8, an abstract data type is a type in our language that represents some data. However, we do not necessarily know the exact representation of the data, and we can therefore only manipulate the data through the functions associated with the data type. If you are familiar with object oriented languages like C++ and Java, an abstract data type is essentially a class that has no public data members. Since C does not provide special constructs for implementing abstract data types, we will use the conventions described in Section 2.8 to implement abstract data types in C.

To illustrate the assumptions typically made about abstract data types in concurrent systems, we will use the abstract data type `array` introduced in Section 2.8. The version of the data type we will use here is available in the file `array-1.c` and shown in Listing 9.8.

As can be seen from the figure, the type itself is named `struct array`. This implementation of the data type contains two members, `data` and `count`. The first, `data`, is a pointer to an array that space for `count` elements. The data type uses the *internal allocation* allocation strategy and therefore provides the functions `array_create` and `array_free` for memory management. As implied by the names, `array_create` creates a `struct array` that can store up to `count` elements. Similarly, `array_free` deallocates the memory associated with the array.

Apart from memory management, the data type provides the following functions for manipulating the contents of the array (remember, since `struct array` is an *abstract* data type, the intention is that `data` and `count` in `struct array` are only manipulated by these functions, even if we do not enforce it):

```
array_count(&array)
```

Get the number of elements in the array.

```
array_get(&array, pos)
```

Get the element at position `pos` in the array. If `pos` is out of bounds, 0 is returned.

```
array_set(&array, pos, value)
```

Set the element at position `pos` in the array to `value`. If `pos` is out of bounds, no change is made to the array.

```

struct array {
    int *data;
    int count;
};

struct array *array_create(int count) {
    struct array *a = malloc(sizeof(struct array));
    a->data = malloc(sizeof(int) * count);
    a->count = count;
    return a;
}

void array_free(struct array *a) {
    free(a->data);
    free(a);
}

int array_count(struct array *a) {
    return a->count;
}

int array_get(struct array *a, int pos) {
    int out = 0;
    if (pos >= 0 && pos < a->count)
        out = a->data[pos];
    return out;
}

void array_set(struct array *a, int pos, int value) {
    if (pos >= 0 && pos < a->count)
        a->data[pos] = value;
}

```

Listing 9.8: Implementation of the abstract data type `array` from `array-1.c`.

The file `array-1.c` also contains the function `use_array` (see Listing 9.9) that uses `struct array`. As can be seen in the figure, the function simply creates an array with enough space for 2 elements using `array_create`, sets both elements using `array_set`, and finally prints the array by calling `array_get` in a loop.

The file also contains a `main` function that calls `use_array`. The `use_array` function is called twice: first through `thread_new`, and secondly directly by `main`. As such, `use_array` will be executed from two threads concurrently.

Practice: Consider how `array-1.c` uses `struct array`. Does `main` and `use_array` use the data structure correctly? Why? You can use *Run* \Rightarrow *Look for errors...* in Progvis to check if there are any concurrency errors in the program.

```

void use_array(void) {
    struct array *arr = array_create(2);
    array_set(arr, 0, 10);
    array_set(arr, 1, 20);

    for (int i = 0; i < array_count(arr); i++)
        printf("%2d: %2d\n", i, array_get(arr, i));

    array_free(arr);
}

int main(void) {
    thread_new(&use_array);
    use_array();
    return 0;
}

```

Listing 9.9: The function `use_array` and `main` from `array-1.c` which uses the array from two threads.

Hopefully, your conclusion from the previous page is that the program in `array-1.c` uses `struct array` correctly. Even though the program uses the data structure from different threads, the two threads operate on *different instances* of the data structure. As such, the threads do not share data, and therefore they do not affect each other in any way. This is quite clearly visible if you run the program in Progvis, as shown in Fig. 9.1.

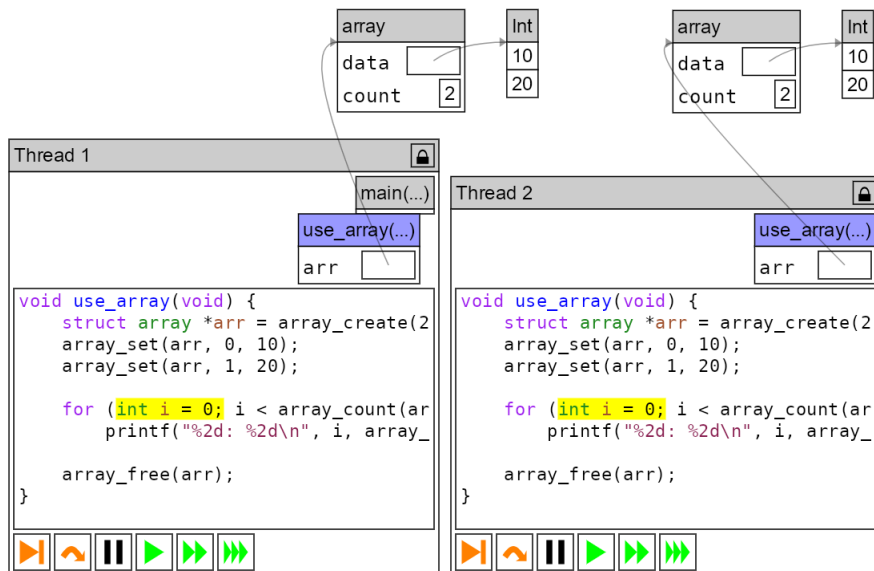


Figure 9.1: The program `array-1.c` running in Progvis.

This observation actually generalizes to a large number of data structures. As long as the data structure does not rely on any global state (e.g., in the

form of global variables), we get this thread safety without doing anything in particular. Because of this, this type of thread safety is usually something that we implicitly expect when we are working with concurrent programs. Namely, we assume that it is safe to manipulate *different instances* of the same data structure from different threads without synchronization.

With this assumption in mind, we can briefly look back to our thread safe version of `count_length`. What we did to make `count_length` thread safe was to encapsulate its state (in this case a single integer) into an abstract data type that we call `struct state`, and then treat `count_length` as an operation of that data type. As such, without additional synchronization, `count_length` is thread safe as long as we follow the expectation above: namely that we do not manipulate `struct state` from multiple threads.

To further explore the expectations we have on abstract data structures in general, we move on to `array-2.c`. This program contains the same implementation of `struct state` as before. The difference is how the program uses the data structure. This part of the program is shown in Listing 9.10 below.

```
void print(struct array *arr, struct semaphore *done) {
    for (int i = 0; i < array_count(arr); i++)
        printf("%2d: %2d\n", i, array_get(arr, i));

    sema_up(done);
}

int main(void) {
    struct semaphore done;
    sema_init(&done, 0);

    struct array *arr = array_create(2);

    thread_new(&print, arr, &done);

    array_set(arr, 0, 10);
    array_set(arr, 1, 20);

    sema_down(&done);
    array_free(arr);
    return 0;
}
```

Listing 9.10: The part of `array-2.c` that uses the `array` data type.

As you can see from the figure, the `main` function first creates a semaphore that it will use to wait for the second thread at a later point. After this, it creates an array with space for two elements, starts a second thread, sets the two elements of the array, waits for the second thread, and finally deallocates the array. The second thread will simply print the contents of the array.

Practice: This program uses `struct` `array` quite differently compared to the previous one. Do you think this way of using the data structure is correct? You can use *Run* \Rightarrow *Look for errors...* in Progviz to verify your answer.

As you have likely discovered, the way `array-2.c` uses the data structure is *not* safe. Since the two threads share data, there is a risk that they access the same data in a way that leads to a data race. In this case, both `main` and `print` will access *all* elements in the data structure, which means that there is a data race in the code. For example, the one illustrated by Progviz in Fig. 9.2.

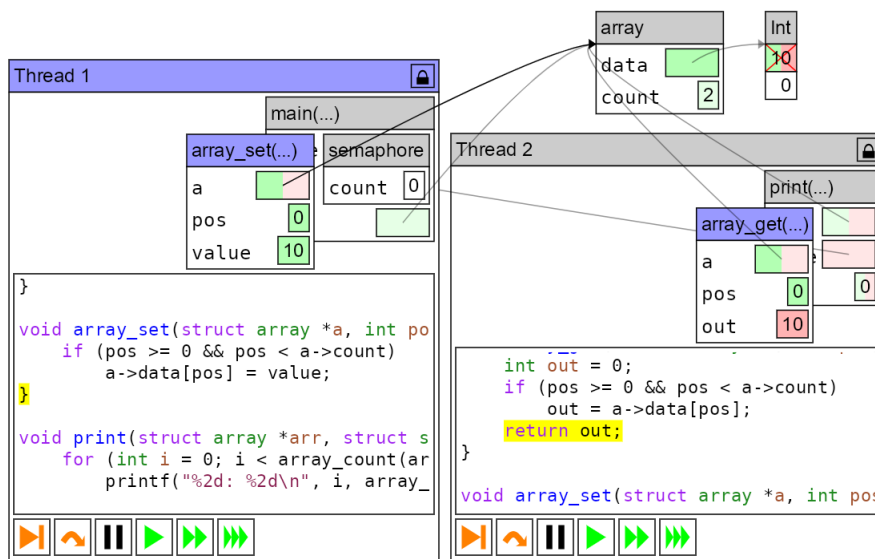


Figure 9.2: The program `array-2.c` running in Progviz, which displays a concurrency error.

Now that we have both observed what we consider *safe* and *unsafe* usage of a data structure, we can start to generalize our findings. In general, we expect that abstract data structures are *as thread safe as* primitive data types. In other words, when we see an abstract data type, we want to be able to think of it as a unit that behaves like an `int` for example. In particular, this means that we expect to be able to manipulate two different instances of the data structure from different threads without issue as we saw above with `array-1.c`. However, we assume that we *need to* synchronize access to the data structure if it is shared between threads as is the case in `array-2.c`.

The generalization above form our default expectations of an abstract data type if nothing else is specified. However, since abstract data types are more complex than simple primitive types, there are numerous exceptions. For example, the synchronization primitives we have introduced so far (i.e., semaphores, locks, and conditions) are examples of such abstract data types. However, all data structures that deviate from the default expectations outlined above typically have documentation that indicates how they differ from “normal” data structures.

Before we continue, it is worth pointing out that our default expectation that abstract data types behave like primitive types is a simplification of reality so that we do not have to consider the exact implementation of all data structures all of the time. For example, assume that one instance of `struct array` is shared between two threads. One thread sets element 0 to 10, while the other thread sets element 1 to 20. The analogy to primitive types tells us that we should synchronize access to the instance, since two threads are modifying it concurrently. However, with our particular implementation of the data type, this situation happens to not lead to a data race since the two threads will write to different locations in the underlying array. However, since this deviation from our expectations is not documented, we should take the safe route and synchronize access to the shared data structure.⁹ The takeaway is that while it is convenient to treat an entire data structure as a unit, it will sometimes lead us to synchronize too much. However, too much synchronization is usually much better (and easier to debug!) compared to a program that crashes once a month due to a hidden synchronization error.

9.2.1 Thread Safe Abstract Data Types

As we just concluded, our default expectations of thread safety of abstract data structures are easy to meet. As long as the implementation of the data structure does not use any global variables it fulfills the expectations without issues. However, as we have seen before, there are situations where we need to share data between threads to allow the threads to communicate with each other and coordinate their work. In this situation it is convenient to use *thread safe abstract data types*, sometimes simply called *thread safe data structures*. These data structures go beyond our basic expectations (i.e., what `array-1.c` and `array-2.c` does), and explicitly allow the data structure to be used from multiple threads without incurring data races or other problematic behavior in the data structure itself.

To illustrate the idea, we will make `struct array` thread safe. To do this, we will need to protect access to shared data with a lock. As in Chapter 6, we start by finding shared data. This step is easy, since we only need to consider the variables in `struct array`. Next, we examine how the shared data is used in the functions that are associated with the data structure. First, we can quickly see that the contents of the `data` array is both read and written by the `array_get` and `array_set` functions. As such we need to protect access to the contents of `data` using a lock.

Next, we need to consider `count` and the pointer `data` itself (i.e., what memory `data` is pointing to, not the content of `data[0]`, `data[1]`, etc.). For both variables we can see that they are written once in `array_create`, and then only ever read from. The fact that they are written to in `array_create` in particular is important. Since `array_create` always allocates a *new* `struct array` using `malloc`, we are certain that *only* the thread that called `array_create` can access the new instance until it is returned from `array_create`. Once `array_create`

⁹One important reason is that while the observation that concurrent writes to different elements is true for *this particular* implementation, it is not true for all possible implementations. For example, if `struct array` was implemented as a linked list, and `array_set` was implemented by removing the old element and inserting the new element, we would have a data race.

returns, we do not know what the caller does with the pointer, so from that point onward we must account for the fact that the data structure might be shared.

Since both `count` and `data` are initialized while we *know* that the newly created array instance is *not* shared, we do not need to protect these writes. Furthermore, we also know that after initialization, the variables are only ever read from. This means that it is safe to not protect them with locks. As such, we only need to protect the memory that `data` points to with a lock. Adding a lock to the data structure makes it look like in Listing 9.11.

```
struct array {
    int *data;
    int count;
    struct lock data_lock;
};

struct array *array_create(int count) {
    struct array *a = malloc(sizeof(struct array));
    a->data = malloc(sizeof(int) * count);
    a->count = count;
    lock_init(&a->data_lock);
    return a;
}

void array_free(struct array *a) {
    free(a->data);
    free(a);
}

int array_count(struct array *a) {
    return a->count;
}

int array_get(struct array *a, int id) {
    int out = 0;
    if (id < a->count) {
        lock_acquire(&a->data_lock);
        out = a->data[id];
        lock_release(&a->data_lock);
    }
    return out;
}

void array_set(struct array *a, int id, int value) {
    if (id < a->count) {
        lock_acquire(&a->data_lock);
        a->data[id] = value;
        lock_release(&a->data_lock);
    }
}
```

Listing 9.11: Thread safe implementation of `struct array` from `array-3.c`.

Note that while the data structure itself is now thread safe in the sense that data races never occur, programs that use the data structure do not automatically behave correctly. For example, the program in `array-3.c` does not always print that element 0 contains 10 and element 1 contains 20 since there is no guarantee that the main thread executes both `array_set` before `print` executes `array_get`.

At this point, it is worth clarifying some details that may not immediately be obvious from before. First, we concluded that we are able to more or less disregard the contents of `array_create` when considering what data needs to be protected by locks. The reason we can do this is that we know that the newly created instance is not shared with other threads yet. Note, however, that there are cases where this happens *before* the `create`-function returns. For example, it is sometimes useful to keep track of all instances of a data structure in a global list (e.g., for caching). In such cases, the `create`-function will add the newly created instance to some data structure that is accessible by other threads. As such, from that point onward, we need to consider the data structure to be shared, even if the `create` function has not yet returned to the caller.

We also ignored the code inside `array_free` in our analysis of which variables need to be protected. The reason for this is that we assume that the thread that calls `array_free` is the only thread that will access the array passed to `array_free`. As such, we are in a similar situation to `array_create`: we know that the array is not shared at that point, so we do not need any synchronization. However, the *reason* we can make this assumption is different from `array_create`. In the case of `array_create` it is obvious from the code that no other threads can possibly have access to the newly allocated array until `array_create` returns. This is not true for `array_free`. The reason we can still assume that the array is not shared at that point in the program, we need to consider that the purpose of `array_free` is to deallocate the data structure, which makes it unusable from that point onward in the program.

To illustrate why we make the assumption that `array_free` must not be called when other threads are using the array, we will have a closer look at the problems that arise if we *allow* `array_free` to be called while other threads are using the data structure. To do this, we can simply remove the semaphore from `array-3.c`, which gives us the program in `array-3-nosync.c`, depicted in Listing 9.12.

```

void print(struct array *arr) {
    for (int i = 0; i < array_count(arr); i++)
        printf("%2d: %2d\n", i, array_get(arr, i));
}

int main(void) {
    struct array *arr = array_create(2);
    thread_new(&print, arr);

    array_set(arr, 0, 10);
    array_set(arr, 1, 20);

    array_free(arr);
    return 0;
}

```

Listing 9.12: The part of the program that uses `struct array` in `array-3-nosync.c`. Note that the semaphore is removed. The remainder of the program is the same as in `array-3.c`.

Practice: As mentioned above, the program `array-3-nosync.c` removes the semaphore from `array-3.c`, which makes it possible for the main thread to call `array_free` while the second thread calls `array_set`. Does this cause any problems? Start by trying reasoning about the program yourself. Then you can use *Run* \Rightarrow *Look for errors...* in Progvis to verify your findings.

If the removal of the semaphore causes any problems, is it possible to solve the problems by modifying *the data structure* (i.e., `struct array` and/or the functions that start with `array_`)? Remember that the data structure should be usable by other programs than the one in `array-3-nosync.c`. In particular, this means that it is *not* possible to have `array_free` wait until `array_count` is called three times. This would not work for programs that call `array_count` one time before calling `array_free` for example.

We can draw two important conclusions from our analysis of `array-3-nosync.c`. First and foremost, if we do not follow the assumption that `array_free` is not called while other threads use the data structure causes the program to misbehave. This is not surprising in and of itself. However, the *reason* the program fails is important. If we load the program in Progvis, and let the main thread run until completion, we will see a situation similar to the one shown in Fig. 9.3. From the figure, we can see that the main thread (Thread 1) has just deallocated the shared instance of `struct array`. This will cause Thread 2 to access memory that was deallocated, which is undefined. If we are *lucky*, the program crashes immediately, which makes us notice the issue so we can fix it. However, most likely the program will appear to run normally, perhaps failing unexpectedly a month later when we add a new feature to the program elsewhere.

Now that we know that the problem is not caused by a data race per se, but rather that a *use after free* might happen, we can start to think about

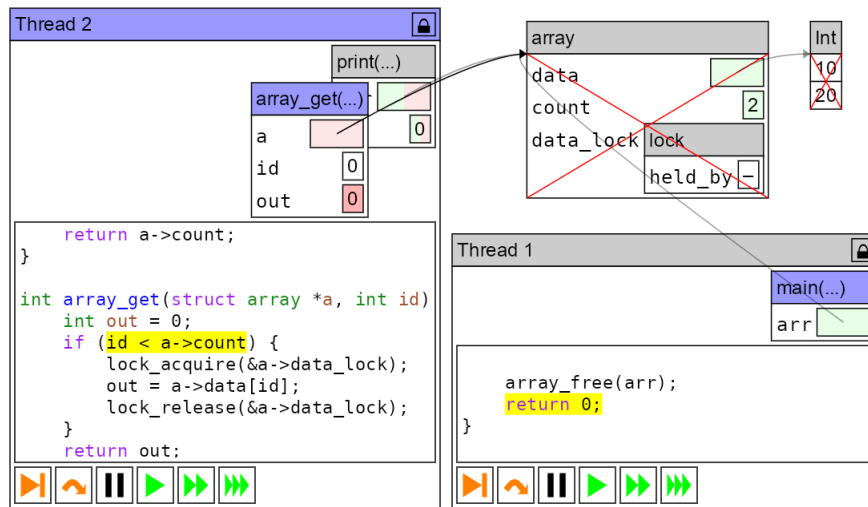


Figure 9.3: The program `array-3-nosync.c` running in Progvis, illustrating a situation where Thread 2 (left) will read memory that was just deallocated by Thread 1 (right).

how to solve the problem. This is our second important conclusion: we *can not* solve the problem by modifying the data structure itself in a generic way. To solve the problem, we must ensure that the calls to `free` inside `array_free` is the last thing that happens for each instance of the data structure. This is not possible to do without intricate knowledge about how the data structure is used. In this particular case, we could have the call to `array_free` wait until `array_count` has been called exactly three times. However, this would not work for the vast majority of programs. For example, it would fail if we would make the seemingly inconsequential change to `print` illustrated below:

```
void print(struct array *arr) {
    int count = array_count(arr);
    for (int i = 0; i < count; i++)
        printf("%2d: %2d\n", i, array_get(arr, i));
}
```

As such, the conclusion is that we generally *need to* leave this responsibility to the user of the data structure.¹⁰ Therefore, almost all data structures implicitly assume that the user of the data structure does not attempt to destroy it while some other thread is using it, which is a very reasonable assumption once we spell it out!

Finally, a brief note on the differences between what we call *internal allocation* and *external allocation*. The data structure we used to illustrate the assumptions used *internal allocation*, which makes it obvious that the instance created in `array_create` is only accessible from the current thread, and that we

¹⁰In certain cases it *is* possible to structure the data structure in such a way that the data structure can deallocate itself at a suitable point. This does, however, often involve other constraints on some of the functions. For example, that it is only possible to call a particular function *once* for each instance of the data structure.

get a *use after free* if we call `array_free` while the data structure is used. The corresponding functions for *external allocation* make the same assumptions, but the errors are perhaps less obvious at first. In the case of `struct array`, the corresponding functions would look like below:

```
void array_init(struct array *a, int count) {
    a->data = malloc(sizeof(int) * count);
    a->count = count;
}

void array_destroy(struct array *a) {
    free(a->count);
}
```

The key insight here is that we consider `array_init` as the only proper way to initialize a `struct array`, and `array_destroy` as the only proper way to destroy it. As such, if we create a variable of the type `struct array` somewhere in our program, we consider it to be in an invalid state until we have called `array_init` to initialize it. Obviously, passing an invalid `struct array` to any function other than `list_init` is something the data structure assumes will not happen. Because of these implicit rules regarding initialization (and the assumption that we only initialize each instance once), it follows that `array_init` needs to be called by one thread, *before* the data structure is shared with other threads. This is illustrated below:

```
struct array a;
// 'a' exists, but is in an invalid state. As such, we
// may not call functions that use 'a'.
array_get(&a, 0); // Incorrect.
// Initialize 'a'. From this point onwards it is valid.
array_init(&a, 2);
// We may now call functions that use 'a'.
array_get(&a, 0);
// We may not attempt to initialize it again.
array_init(&a, 3); // Incorrect.
```

The reverse is true for `array_destroy`. This function accepts an initialized instance of a `struct array` and makes it invalid by destroying it. Conceptually, this is almost identical to `array_free`, but with the exception that we do not deallocate the `struct array` itself. The only real difference is that it is less likely that the program will misbehave, so we need to be extra careful. It is worth noting, that it *is* possible to initialize a variable that we have previously destroyed if we want to re-use it. However, we must still make sure that `array_destroy` and `array_init` are only called once, and that no other threads call any other functions while the variable is in an invalid state. In practice, this means that we must ensure that no other threads access the variable before we call `array_destroy`. As such, once we are done with `a` from the example above, we can destroy it as follows:

```

// We assume that 'a' is initialized. As such, we may
// call functions that use 'a'.
array_get(&a, 0);
// Once we have ensured that no other threads access
// 'a', we can destroy it:
array_destroy(&a);
// At this point we may no longer call any functions
// that use 'a':
array_get(&a, 0); // Incorrect.
// However, we may initialize 'a' to make it valid again
// if we wish. We can do this from a different thread
// if we are careful.
array_init(&a, 2); // OK.
// Of course we need to destroy 'a' again eventually.

```

You probably recognize this pattern from the synchronization primitives that were introduced in previous chapters (i.e., `struct semaphore`, `struct lock`, and `struct condition`). They are all thread safe as long as you follow the assumptions outlined above. This is the reason why we need to make sure that the initialization of our synchronization primitives are guaranteed to happen *before* any thread may try to use them.

9.2.2 Global State and Abstract Data Types

In the previous section, we noted that it is easy to make an data structure follow the basic assumptions as long as it does not depend on any global state. At the start of this chapter, we examined the program `count-length` and found that the proper solution was to avoid global state altogether. Based on this, one might conclude that global state is always problematic in concurrent programs. While it is true that it is a good idea to avoid global state wherever possible, it *is* possible to use global state safely as long as the abstractions that rely on global state are designed with care.

To illustrate how we can manage global state safely, we will examine the program `array-pool`. The program is based on the initial version of our implementation of `struct array`. However, one of the calls to `malloc` has been replaced with a call to `alloc_array`, which allocates memory from a pre-allocated pool of memory in a global array. This is a technique that is often used in applications that need to maximize performance. While `malloc` is not particularly expensive, it is *much* more expensive than `alloc_array`. The first version of the implementation is available in `array-pool-1.c`, and the modified parts are depicted in Listing 9.13.

Practice: The program `array-pool-1.c` is not correct. Find the problem and attempt to fix it using suitable synchronization primitives. You can use *Run \Rightarrow Look for errors...* in Progviz both to find the problem and to verify your solution.

Once you have come up with a solution, compare your approach to synchronization to the one used in `count-length-3.c`. Do you see any similarities between the two? If so, since the approach in `count-length-3.c` was problematic, do you think that your solution to `array-pool-1.c` is correct?

```

int pool[4];
int pool_first_free = 0;

int *pool_alloc(int count) {
    int *alloc = &pool[pool_first_free];
    pool_first_free += count;
    return alloc;
}

struct array {
    int *data;
    int count;
};

struct array *array_create(int count) {
    struct array *a = malloc(sizeof(struct array));
    a->data = pool_alloc(count);
    a->count = count;
    return a;
}

```

Listing 9.13: The implementation and usage of the pool allocator in `array-pool-1.c`.

As mentioned above, the program `array-pool-1.c` is not correct. The issue is that `pool_alloc` uses the global variable `pool_first_free` without synchronization. Since it is used by `array_create`, this means that `array_create` fails to meet the assumptions we make when we are working with an abstract data type, namely that we assume that it is safe to call the `array_`-functions from different threads as long as they operate on different instances of the data structure.

We can solve the problem by applying the approach from Chapter 6 once more. We already found that the problematic data is `pool_first_free` (not `pool`: if `pool_first_free` is synchronized correctly, the rest of the program uses different elements of `pool`). The critical section we need to protect therefore spans the two lines in `pool_alloc` where `data_first_free` is used. We can protect them by adding a global lock, `pool_lock`, and using it as depicted in Listing 9.14 (also in `array-pool-2.c`).

If we compare this implementation to our attempt at synchronizing the program `count-length-3.c` you will see that they are similar. In both cases, we ended up adding a lock to protect access to the global variables. As such, we might suspect that our synchronized version of `pool_alloc` has similar issues to `count_length` in `count-length-3.c`.

This is indeed the case. Depending on the order in which the two threads in `array-pool-2.c` execute `pool_alloc` (it can not happen at the same time due to the lock), the pointers returned from will be different. This effect is visible if we run the program in Progvis. Figure 9.4a shows what happens when the main thread executes `pool_alloc` first, and Fig. 9.4b shows what happens when

```
int pool[4];
int pool_first_free = 0;
struct lock pool_lock;

int *pool_alloc(int count) {
    lock_acquire(&pool_lock);
    int *alloc = &pool[pool_first_free];
    pool_first_free += count;
    lock_release(&pool_lock);
    return alloc;
}
```

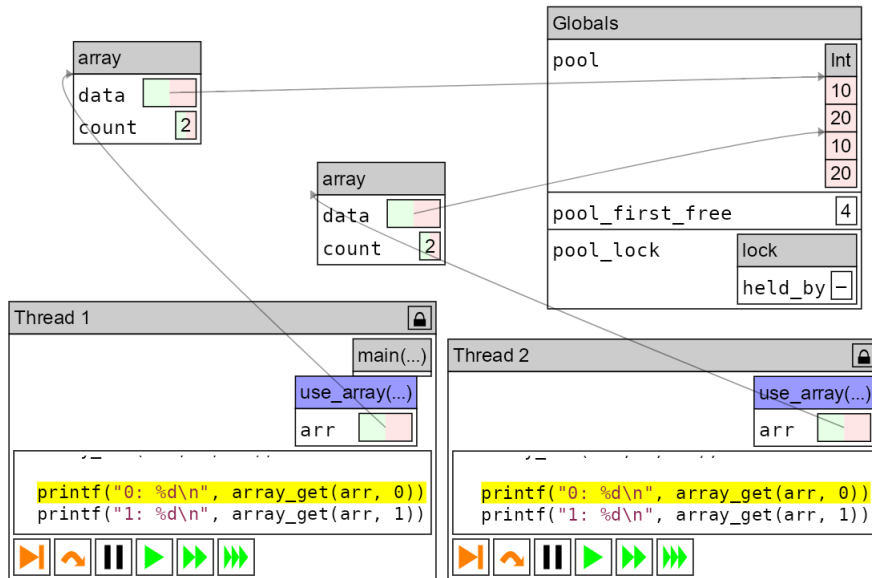
Listing 9.14: Synchronized version of `pool_alloc` from `array-pool-2.c`.

the second thread calls `pool_alloc` first. Note that the only difference between the two figures is that the `data` variable in the `struct array` instance used by Thread 1 refers to the first element in the `pool` array in Fig. 9.4a, and to the second element in Fig. 9.4b. The reverse is true for the `data` variable in the `struct data` instance used by Thread 2.

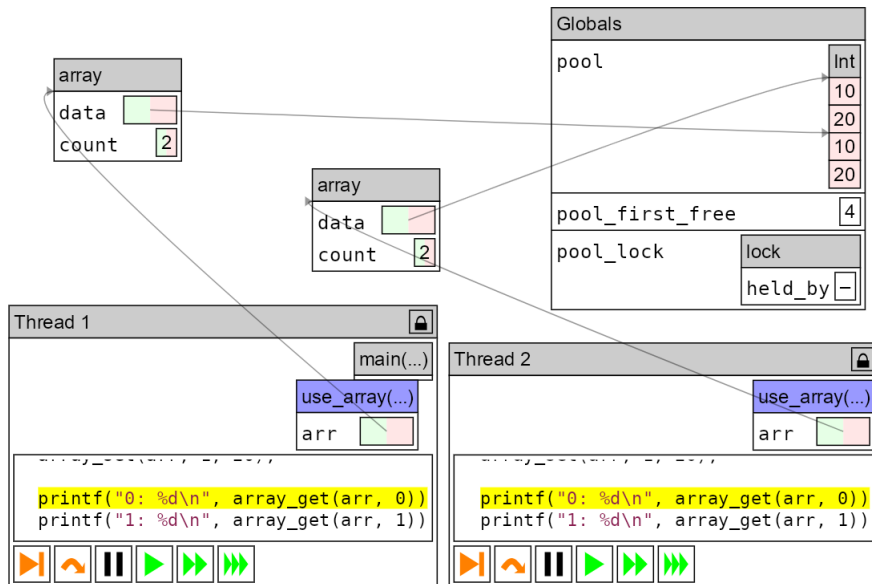
While it is true that the behavior of the program depends on the order in which threads execute `pool_alloc`, there is a big difference between this program and `count-length-3.c`. In this case, the exact address returned by `pool_alloc` does not matter. As long as the implementation ensures that each `struct array` receives a *different* set of elements from the `pool` array, the program works correctly.

It is useful to generalize this insight, so that we can use it to design functions that utilize global state while being thread safe. The key difference between the two is that `count_length` is designed in such a way that it needs to be called multiple times in a row to be useful. Since it is important that the state in the global variable is preserved between the calls, it is problematic if another thread interrupts the sequence by calling `count_length`, even if `count_length` itself is properly synchronized.¹¹ This is not true for `pool_alloc`. It is enough with a single call to `pool_alloc` to allocate a chunk of memory for the array. This is true as long as the user of `pool_alloc` do not assume that two subsequent calls to `pool_alloc` will return pointers to elements in `pool` that are after each other. As such, as long as the user make the same assumptions about `pool_alloc` as they do about `malloc`, we can consider `pool_alloc` to be thread safe.

¹¹This is very similar to what we have seen before. We can indeed remove all data races by protecting all individual accesses using a lock. However, sometimes this is not enough to ensure that the program behaves correctly. Sometimes we need to ensure that a *sequence* of accesses occur as a unit for the program to behave correctly.



(a) The main thread calls pool_alloc first.



(b) The second thread calls pool_alloc first.

Figure 9.4: Progvis running the program array-pool-2.c.

To illustrate usage that seemingly contradicts that `pool_alloc` is thread safe, consider the implementation of `array_create` in `array-pool-faulty.c`, which modifies `array_create` to allocate the `data` array one element at a time (see Listing 9.15). This approach works as long as no other thread calls `pool_alloc` while `array_thread` is running. However, we would consider this code incorrect regardless, since it breaks the contract with `pool_alloc` by assuming that the pointer returned by the second call to `pool_alloc` always points to the element after the pointer returned by the first call (this is why it is ignored).

```

struct array *array_create(int count) {
    struct array *a = malloc(sizeof(struct array));
    a->data = pool_alloc(1);
    for (int i = 1; i < count; i++)
        pool_alloc(1);
    a->count = count;
    return a;
}

```

Listing 9.15: The function `array_create` from `array-pool-faulty.c`, illustrating incorrect usage of `pool_alloc`.

While the example above is perhaps a bit extreme (you would not assume that the code above would work if we replaced `pool_alloc` with `malloc`!), it illustrates two important ideas. First, it is not too difficult to use global state correctly, as long as the interface for interacting with the global state can be completed *as one operation*. That makes it possible to protect the global state using locks so that modifications to the global variables become *atomic*.

The second idea is that, even if we are using a function or data structure that is documented to be thread safe, this only holds if we use it according to its intended usage. Otherwise, we might encounter situations that are not thread safe regardless. For example, as we saw above, by observing that two subsequent calls to `pool_alloc(1)` are equivalent to one call of `pool_alloc(2)`. This is not a part of the documented behavior of `pool_alloc`, and as we saw, it breaks down in concurrent systems. Similarly, even though we think of a lock as being thread safe, it is not correct to acquire the lock from one thread while another thread is initializing it.

9.3 Summary: Sequential Abstractions in Concurrent Systems

Up to this point in the chapter, we have examined how introducing multiple threads to a program affects existing abstractions (i.e., functions and abstract data types). Before we move on to explore how concurrency allows us to create *new* kinds of abstractions compared to in sequential systems, we will take a brief moment to summarize our findings so far.

Overall, we found that many abstractions designed to be used in sequential programs work well for concurrent programs without additional modifications, given that they are used properly. Abstractions in a concurrent system as-

sume the following from the program that uses them unless the documentation indicates otherwise:

- It is safe to call a function (both functions associated with an abstract data type (e.g., member functions in C++) and other functions) from multiple threads, as long as the caller ensures that no other threads may to access the parameters of the function during the call. If one or more of the parameters are functions (or references), this also applies to data referenced by the parameters that the function may use. We say that functions that fulfill this property are *thread safe*.

This rule is similar to what we learned about operators for primitive data types that are provided by the language. For example, we learned that the code that uses the `+=` operator needs to ensure that both operands are not accessed by other threads to avoid a data race. Similarly, if a library provides a function `increment(X *, X)` for some data type `X`, the code that calls `increment` needs to ensure that both parameters are not accessed by other threads during the call to `increment`. As noted above, this not only applies to the pointer passed as the first parameter, but also the instance pointed to by the pointer.

- As long as all operations for an abstract data type (i.e., `struct` in C, classes in other languages, etc.) are thread safe according to the definition above, it is safe to use multiple instances of the same data structure in concurrent programs as long as an individual instance is not used by two or more threads concurrently. Note that we do *not* call these data structures thread safe. This is simply our default expectation.

Again, this rule is similar to our observations regarding primitive types, such as integers and pointers. We have seen that we can use these types from multiple threads concurrently, as long as the different threads do not access the same data concurrently, as that causes a data race. If this is the case, we need to protect the shared data using locks for example. As such, the observation above states that we can treat instances of abstract data types in the same way. One difference is that we can no longer distinguish between reads and writes as we could with primitive types since we do not generally know the exact implementation of every operation. As such, we just assume that all operations may involve a write and disallow concurrent access to the data structure altogether.

- Each instance of an abstract data type needs to be initialized before it is used and destroyed after we finish using it. It is not possible to call any other functions that use a data structure before it has been initialized and after it has been destroyed. Furthermore, initialization and destruction may not occur concurrently with any other operation, just as described above.

As you can see, these assumptions match our observations regarding primitive data types (e.g., integers) well. This is convenient since it means that we generally do not have to distinguish the two when looking for shared data structures in our program.

Furthermore, implementing functions and abstract data types that are safe to use as stated above is very similar to how we would implement them in sequential programs. In fact, all abstractions for sequential programs will behave correctly under the assumptions above, as long as they *do not* use global state (e.g., global variables, CPU state, calling other functions that affect globals etc.). Functions (or data structures as a whole) that use globals need extra care when using them in a concurrent system since they will not automatically behave correctly under the assumptions above. To achieve this, the accesses to global state needs to be protected in a suitable way. However, this is not always trivial to do. Below is an outline of the available approaches:

- If a function that uses globals can be used correctly by calling the function once, it is usually enough to simply add locks or other suitable synchronization primitives to the function's implementation to ensure that two threads do not access the global state at the same time. One example of this is `pool_alloc` in `array-pool-2.c`.
- If a function that uses globals needs multiple calls to happen in sequence, more work is involved. One option is to re-structure the function so that it can complete its purpose in a single call and synchronize the new function. Another option is to encapsulate the state needed by the function as an abstract data structure and let the user of the function manage the state. An example of this is `count_length` in `count-length-4.c`.
- If neither of the above options are possibly, we can document the relevant function(s) as *not thread safe* to inform the user of the functions that they are an exception to the expectations above.

Note that since our default assumption is to *not* be able to share instances of abstract data types between threads, we *do not* call them thread safe by default. If it *is* safe to use a data structure between multiple threads, we call that data structure *thread safe*. To make an abstract data type thread safe, the operations that access the data structure needs to use synchronization primitives internally to ensure that no data races are possible even though multiple operations are called concurrently. An example of this is `array-3.c`. Note, however, that simply using a thread safe data structure does not necessarily mean that all programs that use it automatically behave as we intend them to. Sometimes additional synchronization is necessary by the user of the data structure.

9.4 Concurrent Abstract Data Types

Now that we have covered how abstractions from sequential programming behave in concurrent systems, we are ready to examine what *additional* abilities abstractions have in a concurrent system. In particular, since there are multiple threads executing concurrently in a concurrent system, it becomes useful for abstractions to *wait for other threads* in various ways. As such, abstractions may not only determine *what* value is returned, but also *when* a result is returned. This is often used to implement *thread safe data structures* that are specifically designed to allow different threads to communicate in a safe and easy to use way.

In fact, we have already encountered concurrent abstract data structures (we will use the term *concurrent data structures* for simplicity), namely the synchronization primitives `struct semaphore`, `struct lock`, and `struct condition`. All of them are thread safe since they allow multiple threads to interact with them concurrently. Furthermore, they are typically only useful in concurrent programs since their sole purpose is to wait for other threads in order to allow the threads to communicate safely. The goal of this section is therefore to illustrate that we can build our own abstractions in terms of concurrent data structures, that are not only thread safe, but can be used to influence thread execution similarly to semaphores.

9.4.1 Future

The first concurrent data structure we will cover is a *future*. Essentially, a *future* represents some data that will be produced by a thread at some point in the future. The future itself then helps manage access to that data. In particular, it ensures that the data is not accessed before it is produced. To illustrate the idea in more detail, we will both examine an example of applying a future as well as an implementation of a simple future.

```
int result;
struct semaphore has_result;

void expensive(const char *x) {
    // We use 'strlen' as a stand in for some more
    // complex and time-consuming operation.
    result = strlen(x);
    sema_up(&has_result);
}

int main(void) {
    sema_init(&has_result, 0);
    thread_new(&expensive, "example");
    // 'main' can do other things here!
    sema_down(&has_result);
    printf("Result: %d\n", result);
    return 0;
}
```

Listing 9.16: The program in `future-1.c`.

We will use the program `future-1.c` as a starting point. As can be seen in Listing 9.16, the program contains two functions: `expensive` and `main`. As the name implies, we imagine that the former performs some form of expensive computations. For that reason, `main` wishes to run `expensive` on a separate thread, so that it is able to do other things while it waits for the expensive computations to complete. As you can see, the `expensive` function only calls `strlen` to keep the code simple. For the purposes of this discussion, imagine that `strlen` is expensive.¹²

¹²While `strlen` is not the cheapest thing one can do in terms of performance, it is neither particularly expensive in the grand scheme of things.

The program uses the same approach we used in Chapter 4 to let the main thread wait for the second thread to finish in a safe way. In particular, we consider `has_result` to count whether `result` contains a valid result or not. The `expensive_result` function writes the result of the expensive computations to the `result` variable and then calls `sema_up`. Similarly, the main thread calls `sema_down` before reading `result` to ensure that there is a result to read.

As we have seen before, this approach works well for small programs. However, assume that we have more than 3 cores available and we wish to execute two calls to `expensive` in parallel. If we modify `main` to start another thread that also executes `expensive`, we get the program in `future-2.c` (Listing 9.17).

```
int main(void) {
    sema_init(&has_result, 0);
    thread_new(&expensive, "example");
    thread_new(&expensive, "second");
    // 'main' can do other things here!
    sema_down(&has_result);
    printf("Result: %d\n", result);
    return 0;
}
```

Listing 9.17: The program `future-2.c` starts two threads that each run the same function.

As you have likely realized already, the program in `future-2.c` does not work as expected. Since both threads use the same `result` variable and the same semaphore, they will overwrite each other's result, and it will not be possible to determine which of the two threads that finished when the call to `sema_down` returns.

The solution to this problem is of course to have two sets of variables, one for each thread. To make `expensive` use the correct variables, one option is to have two versions of `expensive`, one for each thread. Another option is to have `expensive` accept pointers to the variables it should use. Regardless, the implementation quickly gets difficult to understand and maintain. So, to keep the code easy to understand and to avoid repeating the waiting and signaling logic, we encapsulate the logic into a concurrent data structure named `struct future`. Just like when we create other abstractions, the first step is to figure out which operations the data structure needs to support.

If we examine the synchronization-related tasks performed by the program `future-1.c`, we find two operations: the last thing that `expensive` does is to *post* a result to the main thread. The main thread in turn *gets* the result, which implies that it has to wait for the other thread to post it. Apart from that, we will also need initialization and destruction. In summary, our data structure needs to support the following operations:

```
struct future *future_create(void)
```

Allocate and create an instance of `struct future`.

```
void future_free(struct future *f)
```

Destroy an instance of `struct future`.

```
void future_post(struct future *f, int result)
```

Post a result to the future. The future stores the result and makes it available using the `future_get` operation. We assume that this only happens once for each instance.

```
int future_get(struct future *f)
```

Get the value from the future. Waits until a result has been posted to the future `f` and returns it. If a value has already been posted, the function returns immediately. It is possible to call `future_get` multiple times for each future.

Note that in contrast to abstractions used in sequential programs, the description of `future_get` specifies *when* the function returns, not only *what* it returns. In this case, it specifies that the function may *wait* for something else to happen within the program. This kind of behavior is what is unique to concurrent data structures. After all, waiting for another part of the program to do something is not very useful without multiple threads that execute concurrently.¹³

The next step is to consider what `struct future` needs to contain. In this case, we can quite easily conclude that we can simply move the global variables `result` and `has_result` into the data structure. As before, we treat these variables as if they are private to the data structure, even if our implementation does not prevent other code from accessing them.

Finally, we can implement the operations above. As a first attempt, we simply move the relevant pieces of code from `future-1.c` into the corresponding operation. This gives us the implementation in `future-3.c` (Listing 9.18).

We also need to modify the `expensive` and `main` functions accordingly. In `expensive`, we replace the assignment to `result` with the initialization of a local variable, and the call to `sema_up` is replaced with a call to `future_post`. However, we also need an instance of a future to pass as the first parameter to `future_post`. To avoid creating two copies of `expensive`, we also add an additional parameter so that it accepts a pointer to the `struct future` to be used.

The `main` function also needs to be modified accordingly. We remove the call to `sema_init` and replace it with two calls to `future_create` to create two futures and save them in two variables. Secondly, we remove the call to `sema_down` and instead replace `result` as the parameter to `printf` with a call to `future_get`. We also duplicate the `printf` line to print the result from the second thread. The end result is the code in `future-3.c` (Listing 9.18).

¹³One exception is perhaps when a sequential thread waits for some external event (e.g., hardware, network). However, one could argue that this is an example of a concurrent system. After all, whatever triggers the external event is probably doing some form of computation that the program is waiting for.

```
struct future {
    int result;
    struct semaphore has_result;
};

struct future *future_create(void) {
    struct future *f = malloc(sizeof(struct future));
    sema_init(&f->has_result, 0);
    return f;
}

void future_free(struct future *f) {
    free(f);
}

void future_post(struct future *f, int result) {
    f->result = result;
    sema_up(&f->has_result);
}

int future_get(struct future *f) {
    sema_down(&f->has_result);
    return f->result;
}

void expensive(struct future *f, const char *x) {
    // We use 'strlen' as a stand in for some more
    // complex and time-consuming operation.
    int result = strlen(x);
    future_post(f, result);
}

int main(void) {
    struct future *f1 = future_create();
    struct future *f2 = future_create();
    thread_new(&expensive, f1, "example");
    thread_new(&expensive, f2, "second");
    // 'main' can do other things here!
    printf("Result 1: %d\n", future_get(f1));
    printf("Result 2: %d\n", future_get(f2));
    future_free(f1);
    future_free(f2);
    return 0;
}
```

Listing 9.18: The previous program rewritten to use `struct future`. Available as `future-3.c`.

Practice: Consider the programs below that use `struct future` and answer the following two questions:

- Which of the programs use the future according to the specification?
- Does any of the programs that use the future according to the specification still not work correctly with the current implementation of the future?

The functions `thread1` and `thread2` execute concurrently in two different threads. Both have access to the same future `f`. The future `f` has been initialized before both threads start executing. The programs are available as `future-practice-n.c` if you wish to examine them using Progviz.

1:	<pre>void thread1() { future_post(f, 1); }</pre>	<pre>void thread2() { future_get(f); }</pre>
2:	<pre>void thread1() { future_post(f, 1); }</pre>	<pre>void thread2() { future_post(f, 2); }</pre>
3:	<pre>void thread1() { future_post(f, 1); future_get(f); }</pre>	<pre>void thread2() { future_get(f); }</pre>
4:	<pre>void thread1() { future_post(f, 1); future_free(f); }</pre>	<pre>void thread2() { future_get(f); }</pre>
5:	<pre>void thread1() { future_post(f, 1); }</pre>	<pre>void thread2() { future_get(f); future_free(f); }</pre>

Answers are provided in a footnote on the next page.

Now that we have implemented the future and spent some time familiarizing ourselves with its semantics, we can make some observations so far. First and foremost, we can see that our `struct future` can be used to coordinate multiple threads, just like semaphores. This is perhaps not too surprising since it is using a semaphore internally. Furthermore, the data structure is thread safe in the sense that it allows multiple threads accessing it concurrently. This is nice since we do not have to remember to protect access to the data structure if we wish to share it between multiple threads. However, as we have seen from above, this does not mean that it is safe to use the data structure any way we

please. We still need to respect the constraints defined by the data structure, and sometimes it is not immediately obvious if those constraints are respected in a complex concurrent system.

For example, even though we consider `struct future` to be thread safe, program 2 (`future-practice-2.c`) still causes a data race as it violates the assumption that `future_post` is only called once. Furthermore, programs 4 and 5 both look suspicious since they attempt to free the data structure. However, since the future ensures that `future_post` waits until `future_get` is called, program 5 is actually safe since this means that `thread1` has always finished using the future once `future_free` is called. This example illustrates that it is important to take the guarantees and limitations provided by concurrent abstract data types into account when reasoning about the behavior of programs. In this way, they are much like our synchronization primitives!

Finally, you have most likely noticed that program 3 (`future-practice-3.c`) uses the operations of the future correctly, but does not work correctly. This again shows the benefit of having a well-defined interface between the user of the abstraction and the abstraction itself. That way we can conclude that the problem in program 3 is due to a bug in the implementation of the future rather than the future being used incorrectly.¹⁴

```
int main(void) {
    struct future *f1 = future_create();
    struct future *f2 = future_create();
    thread_new(&expensive, f1, "example");
    thread_new(&expensive, f2, "second");
    // 'main' can do other things here!
    printf("Result 1: %d\n", future_get(f1));
    printf("Result 2: %d\n", future_get(f2));
    printf("Sum: %d\n", future_get(f1) + future_get(f2));
    future_free(f1);
    future_free(f2);
    return 0;
}
```

Listing 9.19: A revised `main` function that illustrate an issue with the `struct future` implementation.

The same problem can be illustrated in `future-3.c` by adding another `printf` statement that prints the sum of the two calls to `expensive` as shown in Listing 9.19. Luckily, it is fairly easy to solve the problem. We can find the cause of the issue by considering what resource that `has_result` is counting. As before, we think of the semaphore as counting whether `result` contains a value or not. We initialize the semaphore to 0, and increment it in `future_post`. We then decrement the semaphore in `future_get` in order to wait for the

Answers to the Practice box on the previous page: Programs 1, 3, and 5 use `struct future` correctly. Program 2 calls `future_post` more than once. Program 4 is allows `thread1` to free the future before `thread2` finishes its call to `future_get`. Furthermore, program 3 does not currently work correctly with the current implementation of `struct future`.

¹⁴A third option is of course also possible: a bug in the specification.

semaphore to be incremented to 1. However, since the value in `result` is not *removed* by `future_get`, the counter in the semaphore no longer reflects the fact that we still have a result. One way to solve the problem is to simply add a call to `sema_up` directly after `sema_down` as is the case in `future-4.c` (Listing 9.20).

```
int future_get(struct future *f) {
    sema_down(&f->has_result);
    sema_up(&f->has_result);
    return f->result;
}
```

Listing 9.20: Revised version of `future_get` in `future-4.c`.

While this change fixes the issue, it is not necessarily optimal in the case where many threads retrieve the result from the future. Since each thread decrements the semaphore to 0, it will cause other threads to wait in `sema_down` for a short while. Since we know that threads will always wake up fairly quickly in this case (all threads execute `sema_up` right after `sema_down` after all), we usually do not consider this as if the threads are waiting for something else. We will discuss this distinction in more detail in the next section. It is possible to solve this slight performance problem (e.g., using *atomic operations* that are discussed in Chapter 10), but since it is fairly minor we will not discuss it in further detail here.

9.4.2 Bounded Queue

Next we will examine another useful concurrent abstract data type, both to introduce the data structure itself, but also to illustrate some more nuances of abstractions in concurrent systems. We will refer to the data structure as a *bounded queue*, but be aware that you might find it under different names. For example, *pipes* in UNIX and Windows NT and *channels* in Go are implemented like bounded queues. Both are used to make communication between different threads or processes more convenient.

As the name implies, a bounded queue is implemented as a queue that has a limited size. We can conveniently implement this data structure using an array of a pre-determined size and two variables, `head` and `tail`, that keep track of the start and end of the elements in the queue. To avoid problems when the queue reaches the end of the array, we treat the array as if it is *circular*. That is, once we get to the end of the array, we start over from the beginning. You have most likely encountered using a *circular array* to implement a queue previously.

Since the data structure is a queue, we can easily imagine that it needs two operations, push and pop (or equivalently enqueue or dequeue), in addition to the ability to create and destroy instances of the queue. We can implement them as shown in Listing 9.21 (`queue-1.c`).

From the implementation in Listing 9.21 we can see that `struct queue` contains 5 data members. The variable `data` stores a pointer to an array that is used to store the elements in the queue. The array is of a fixed size, and since it is not possible to retrieve the size of a dynamically allocated array, we also store the size of the array in `capacity`. Even though the array has a specific

```

struct queue {
    int *data;
    int capacity;
    int filled;
    int head;
    int tail;
};

struct queue *queue_create(int capacity) {
    struct queue *q = malloc(sizeof(struct queue));
    q->data = malloc(sizeof(int) * capacity);
    q->capacity = capacity;
    q->filled = 0;
    q->head = 0;
    q->tail = 0;
    return q;
}

void queue_free(struct queue *q) {
    free(q->data);
    free(q);
}

void queue_push(struct queue *q, int value) {
    assert(q->filled < q->capacity); // Need free space!
    q->data[q->tail] = value;
    if (++q->tail >= q->capacity)
        q->tail = 0;
}

int queue_pop(struct queue *q) {
    assert(q->filled > 0); // Need something to pop!
    int result = q->data[q->head];
    if (++q->head >= q->capacity)
        q->head = 0;
    return result;
}

```

Listing 9.21: An initial attempt at implementing a bounded queue. Available in `queue-1.c`.

capacity, the queue does not always store `capacity` number of elements. As such, it is also convenient to store the number of elements actually used. We use the variable `filled` for this. Finally, we store the index of the next element to remove in `head`, and the index of the element after the end of the queue in `tail` (i.e., the index where we will store the next element).

The queue is currently implemented as a normal, non-thread safe data structure. One important design decision when implementing a bounded queue (or a bounded stack) is to decide what should happen when `queue_pop` is called on an empty queue, and when `queue_push` is called on a full queue. Popping an empty queue likely indicates an error in the program that uses the queue,

so we signal an error using `assert`.¹⁵ Since we are implementing a bounded queue, we also assert when pushing to a queue that is full. However, we could attempt to resize the queue in this case if we did not care about the fact that the queue is bounded.

However, if we aim to implement a concurrent data structure we have other options! In particular, we can choose to let the calling thread *wait* in both of these situations. That is, if one thread calls `queue_pop` on a queue that is empty, we can let the thread *wait* until another thread pushes a value to the queue. Similarly, when a thread calls `queue_push` and finds that the queue is full, we let it *wait* until another thread pops a value to make space for the new one. As you can see, this behavior only makes sense in a concurrent system where multiple threads may access the data structure. It is not very useful to have the current thread wait for another thread if there is only one thread in the application after all. Of course, this means that we need to make sure that `struct queue` is thread safe.

Before starting to make the queue thread safe, we start by writing down the expected behaviors of the operations. We will pay special attention to the behavior in a concurrent system.

```
struct queue *queue_create(int capacity);
```

Create a queue that is able to store up to `capacity` integers. The returned queue is safe to use from multiple threads concurrently without additional synchronization.

```
void queue_free(struct queue *q);
```

Free the memory associated with a queue that was previously created using `queue_create`. We assume that the queue is not used by any other threads when it is freed.

```
void queue_push(struct queue *q, int value);
```

Push an element to the tail end of the queue. If the queue already contains `capacity` elements, the calling thread will wait until an element is removed using `queue_pop`.

```
int queue_pop(struct queue *q);
```

Pop an element from the head end of the queue. If the queue does not contain any elements, the calling thread will wait until an element is added using `queue_push`.

Based on the more detailed description of the semantics, we can proceed to synchronize the implementation. In this example we will use semaphores, and as such we start by identifying which resources the semaphores should count so that the threads will wait at the appropriate times. In this case, we need threads to wait when the queue is empty and when it is full. To get a thread to wait when the queue is empty, we can simply use a semaphore that counts the number of elements currently in the queue (just like the variable `free`). In

¹⁵Other good options include returning a type that can indicate failure (e.g., `Maybe<int>`) or throwing an exception, depending on what is convenient in the language in question.

light of this, we can think of `queue_push` as producing an element in the queue, and we should therefore call `sema_up` once we have added an element. Similarly, `queue_pop` consumes an element in the queue, and we should therefore call `sema_down` before we try to remove an element from the queue. Crucially, since the semaphore is counting the number of elements in the queue, and the semaphore will wait when trying to decrement the internal counter below zero, we will get the desired behavior.

It is not possible to use the same semaphore to wait when the queue is full: when we need to wait, the counter in the semaphore that replaces `filled` will be equal to `capacity`. As such we need another semaphore for this purpose, and this semaphore needs to count some quantity that becomes zero when we wish to wait. In this case, we can count the number of *free elements* in the array (i.e., we count `capacity - filled`). Calling `queue_pop` produces one free element, `queue_push` consumes one element, and crucially the number of free places will be zero exactly when we need `queue_push` to wait.

Based on this reasoning, we can conclude that we need two semaphores, `filled` that counts the number of elements in the queue, and `free` that counts the number of unused elements in the queue. Interestingly, we no longer need `int filled` since all information is stored in the two semaphores, and therefore we remove it. This results in the implementation in Listing 9.22, which is available in `queue-2.c`.

The updated behavior once again makes the queue into concurrent data structure. Just like with `struct future`, this means that we can use the queue to coordinate the work done by different threads. While a bounded queue can be used as a simple future (in cases where it is not important to be able to retrieve the result multiple times), the ability for bounded queues to store multiple values makes it useful in more situations. One common example is to parallelize some long-running task into multiple distinct steps, and have one thread perform each step in order.¹⁶ If we use bounded queues to pass elements from one step to the next, it means that the whole pipeline does not need to stop if one step needs a bit more time to process some elements. The storage capacity in each bounded queue is able to store some additional elements to let the earlier steps continue regardless.

The last part of the program `queue-2.c` illustrates this using the functions `thread` and `main`. Here, we assume that `thread` is the first step in the pipeline. It produces data for the next step by calling `expensive` and pushes them onto a queue. We assume that the function `expensive` performs some interesting but time-consuming computations. For clarity it is implemented as the identity function, however. The queue used by `thread` is shared with `main`, which is the second step of the pipeline. It pops elements from the queue one by one and prints them. The relevant part of `queue-2.c` is shown in Listing 9.23.

As long as the implementation of `struct queue` adheres to the specification above, the program in Listing 9.23 will work correctly. Since `main` calls `queue_pop` the main thread will wait until `thread` has pushed at least one element to the queue. Furthermore, `thread` is able to continue computing and producing elements without waiting for `main` to print each element since the queue has room for multiple elements. However, this is only true to a certain degree since

¹⁶That is, much like a pipeline of processes: `a | b | c`.

```

struct queue {
    int *data;
    int capacity;
    int head;
    int tail;

    struct semaphore filled;
    struct semaphore free;
};

struct queue *queue_create(int capacity) {
    struct queue *q = malloc(sizeof(struct queue));
    q->data = malloc(sizeof(int) * capacity);
    q->capacity = capacity;
    q->head = 0;
    q->tail = 0;
    sema_init(&q->filled, 0);
    sema_init(&q->free, capacity);
    return q;
}

void queue_free(struct queue *q) {
    free(q->data);
    free(q);
}

void queue_push(struct queue *q, int value) {
    sema_down(&q->free);
    q->data[q->tail] = value;
    if (++q->tail >= q->capacity)
        q->tail = 0;
    sema_up(&q->filled);
}

int queue_pop(struct queue *q) {
    sema_down(&q->filled);
    int result = q->data[q->head];
    if (++q->head >= q->capacity)
        q->head = 0;
    sema_up(&q->free);
    return result;
}

```

Listing 9.22: Version of the queue that handles over- and underflow by waiting rather than asserting. Available in `queue-2.c`.

the queue has a limited capacity. Since the capacity is set to 2 in this case, `thread` can not push its third element to the queue before `main` has popped at least one element. As such, `thread` will eventually have to wait for `main` if the output is slow.

The fact that `thread` has to wait if `main` is too far behind (determined by the size of the bounded queue) is important in a pipeline like the one illustrated

```

void thread(struct queue *q, int start) {
    for (int i = 0; i < 6; i++) {
        int to_push = expensive(i + start);
        queue_push(q, to_push);
    }
}

int main(void) {
    struct queue *q = queue_create(2);
    thread_new(&thread, q, 1);

    for (int i = 0; i < 6; i++) {
        int item = queue_pop(q);
        printf("Got item: %d\n", item);
    }

    queue_free(q);
    return 0;
}

```

Listing 9.23: Program that uses the queue to pass data between two threads. Available in `queue-2.c`.

by `queue-2.c`. Imagine that the loops in both threads run forever (e.g., the program continuously monitors some sensor). If the output of the program is saved to a file or sent over the network, `printf` may occasionally need much more time than `expensive` (e.g., other programs use the disk, the network is unstable). If an unbounded queue was used in such cases, `expensive` would still be able to continue and make the queue grow indefinitely, eventually causing the program to run out of memory. A bounded queue solves this problem by causing `expensive` to wait in such situations. This behavior is often called *back pressure* since it is useful in pipelines. At a large scale, it means that all parts of a pipeline will end up processing data at roughly the same speed, regardless of the computations involved. In our case, the program works correctly regardless of whether `expensive` is faster or slower than `printf`.

Practice: Our intention was to make `struct queue` into a thread safe data structure. Given that we successfully used it to coordinate the work between threads in `queue-2.c`, it seems like we have succeeded. However, the data structure does not currently fulfill the specification presented earlier in this chapter. In particular, it makes one assumption about `queue_push` and `queue_pop` that is not stated in the specification. Which one is it?

As a clue, `queue-3.c` contains the same implementation of `struct queue` as in `queue-2.c`, but some modifications to `main` (Listing 9.24) that reveals a problem. As usual, you can use *Run* \Rightarrow *Look for errors...* in Progviz to help you find the problem. Use that insight to help you formulate the extra assumption the data structure makes about how it is used!

```

void thread(struct queue *q, int start) {
    for (int i = 0; i < 3; i++) {
        int to_push = expensive(i + start);
        queue_push(q, to_push);
    }
}

int main(void) {
    struct queue *q = queue_create(2);
    thread_new(&thread, q, 1);
    thread_new(&thread, q, 4);

    for (int i = 0; i < 6; i++) {
        int item = queue_pop(q);
        printf("Got item: %d\n", item);
    }

    queue_free(q);
    return 0;
}

```

Listing 9.24: The `thread` and `main` functions from `queue-3.c`.

As you have likely concluded from above, the current implementation of `struct queue` assumes that each of `queue_push` and `queue_pop` are only used by one thread at a time. In `queue-2.c`, only `thread` calls `queue_push` and only `main` calls `queue_pop`. However, `queue-3.c` starts *two* threads that execute the `thread` function, and therefore *two* threads call `queue_push` on the same queue concurrently. However, `queue-3.c` starts *two* threads that execute `thread`, and thereby *two* threads call `queue_push` on the same queue concurrently. Since the capacity of the queue is 2, the semaphore lets both of them to continue, and they may therefore access the `tail` variable concurrently which leads to a data race. A similar issue exists in `queue_pop` even though it is not visible given how the data structure is used in `queue-3.c`.

To solve the problem we need to avoid concurrent access to the `head` and `tail` variables. As before, we do this quite conveniently using a lock. Since `head` and `tail` are not used together, we can quite easily conclude that we do not need to synchronize them as a unit.¹⁷ Therefore we use two locks, `head_lock` and `tail_lock` to synchronize the data structure as shown in Listing 9.25 (`queue-4.c`). With these changes the data structure works according to the specification and according to our expectations.

The way that the program `queue-4.c` uses `struct queue` illustrates how bounded queues can be used to coordinate expensive computations performed by multiple threads. In this case, the program calls `expensive` a total of 6 times. One way to do this would of course be to create 6 futures and start 6 threads that post their result to its own future. However, this does not scale well to large computations (e.g., millions of calls) since it is fairly expensive to create and maintain a large number of threads. The program `queue-4.c` instead takes

¹⁷That is, we never need to make decisions based on the value of `head` *in relation to* the value of `tail`.

```

struct queue {
    int *data;
    int capacity;
    int head;
    int tail;

    struct semaphore filled;
    struct semaphore free;
    struct lock head_lock;
    struct lock tail_lock;
};

struct queue *queue_create(int capacity) {
    struct queue *q = malloc(sizeof(struct queue));
    q->data = malloc(sizeof(int) * capacity);
    q->capacity = capacity;
    q->head = 0;
    q->tail = 0;
    sema_init(&q->filled, 0);
    sema_init(&q->free, capacity);
    lock_init(&q->head_lock);
    lock_init(&q->tail_lock);
    return q;
}

void queue_free(struct queue *q) {
    free(q->data);
    free(q);
}

void queue_push(struct queue *q, int value) {
    sema_down(&q->free);
    lock_acquire(&q->tail_lock);
    q->data[q->tail] = value;
    if (++q->tail >= q->capacity)
        q->tail = 0;
    lock_release(&q->tail_lock);
    sema_up(&q->filled);
}

int queue_pop(struct queue *q) {
    sema_down(&q->filled);
    lock_acquire(&q->head_lock);
    int result = q->data[q->head];
    if (++q->head >= q->capacity)
        q->head = 0;
    lock_release(&q->head_lock);
    sema_up(&q->free);
    return result;
}

```

Listing 9.25: The updated implementation of `struct queue` from `queue-4.c`.

a different approach. It spawns a fixed number of threads (typically equal to the number of available CPU cores), and distributes the 6 calls to `expensive` between the two threads by asking the first thread to compute values 1–3 and the second to compute values 4–6. The threads then post their results to a queue so that the main thread can process them as they arrive, regardless of in which order in which the results are completed.

Distributing work between threads ahead of time, as is the case in `queue-4.c`, works well if each call to `expensive` takes approximately the same amount of time (or is otherwise easily predictable). If this is not the case, one option is to use another bounded queue to distribute work to the threads (which are sometimes referred to as a *thread pool*). Since the threads share this work queue, individual items do not belong to a particular thread. Instead, each thread simply pops the next available item from the queue whenever they have completed the previous one. This can be implemented like in Listing 9.26 (`queue-5.c`).

```

void worker(struct queue *in, struct queue *out) {
    int input = 0;
    while (input >= 0) {
        input = queue_pop(in);
        int to_push = expensive(input);
        queue_push(out, to_push);
    }
}

int main(void) {
    struct queue *in = queue_create(2);
    struct queue *out = queue_create(2);
    thread_new(&worker, in, out);
    thread_new(&worker, in, out);

    queue_push(in, 10);
    queue_push(in, 8);
    // Tell both threads it is time to exit.
    queue_push(in, -1);
    queue_push(in, -1);

    for (int i = 0; i < 4; i++) {
        int item = queue_pop(out);
        printf("Got item: %d\n", item);
    }

    queue_free(in);
    queue_free(out);
    return 0;
}

```

Listing 9.26: Using two queues to distribute a (possibly dynamic) amount of work between a fixed number of threads in a *thread pool*.

Before we conclude the chapter, there are a few details of the implementation that we want to explore in further detail. In Listing 9.25 (`queue-4.c`), we

added locks to `queue_push` and `queue_pop` to avoid data races when the operations are used concurrently from multiple threads. However, locks avoid data races by making threads *wait* for each other. In particular, if two threads call either of `queue_push` or `queue_pop` at the same time, they will have to wait for each other. This is not a part of the specification! The specification only states that `queue_push` may wait for `queue_pop` and vice versa. What is the reason for this?

To understand why we treat the two situations where threads may wait differently, we need to consider what threads are waiting *for*. In our implementation we use the semaphore `filled` to allow threads calling `queue_pop` to wait for another thread to call `queue_push`. Similarly, the semaphore `free` is used to allow threads calling `queue_push` to wait for another thread to call `queue_pop`. Note that both of these descriptions require another thread to use the interface of the queue in some way. In particular, if one thread waits for something to happen to a queue that is not shared with other threads, it will wait forever. For this reason, this kind of waiting behavior needs to be specified in detail in the specification, since it affects how other parts of the program uses the data structure.

This is not true for the locks `head_lock` and `tail_lock` we added in Listing 9.25 (`queue-4.c`). They are only used to enforce mutual exclusion to the `head` and `tail` variables. As such, if `lock_acquire` causes a thread to wait, it means that another thread is *currently executing* the code in the corresponding critical section. Since the code inside the critical sections does not involve waiting for some other operation (either directly or indirectly), we know that the thread that holds the lock will eventually release it, and the waiting thread may continue. The conclusion is that neither of the locks cause threads to wait for some other thread to *do* something with the queue. Therefore, there is no risk that the locks will cause a thread to wait forever. This is why we typically do not mention this type of waiting in the specification. It is an implementation detail that is needed by this particular implementation to allow multiple threads to call `queue_push` and `queue_pop` concurrently.¹⁸

To illustrate this idea even clearer, we add the operation `queue_count` to the data structure as shown in Listing 9.27. Note that the addition of the `count` variable means that the two critical sections in `queue_push` and `queue_pop` now need to be protected by the same lock since they both need to update `count`. For this reason the two locks `head_lock` and `tail_lock` are now replaced with a single lock that is named `lock`.

Since `queue_count` acquires a lock, it may need to wait for some other thread(s) to finish executing `queue_push`, `queue_pop`, and/or `queue_count`. However, `queue_count` will always complete eventually (as long as the data structure is not deallocated). Therefore we would simply describe it as returning the current number of elements in the queue, and perhaps reminding the reader that it is safe to call from multiple threads since the queue is thread safe.

Again, note that this is very different from `queue_push` and `queue_pop`. While we know that `queue_count` will always finish eventually, there are situations where `queue_push` and `queue_pop` (when the queue is full and empty respectively) wait forever if no other thread interacts with the queue. This is why we

¹⁸There are indeed other implementations of queues that do not need locks, even though one could argue that they also need to wait in similar situations as well.

```

struct queue {
    int *data;
    int capacity;
    int count;
    int head;
    int tail;

    struct semaphore filled;
    struct semaphore free;
    struct lock lock;
};

void queue_push(struct queue *q, int value) {
    sema_down(&q->free);
    lock_acquire(&q->lock);
    q->count++;
    q->data[q->tail] = value;
    if (++q->tail >= q->capacity)
        q->tail = 0;
    lock_release(&q->lock);
    sema_up(&q->filled);
}

int queue_pop(struct queue *q) {
    sema_down(&q->filled);
    lock_acquire(&q->lock);
    q->count--;
    int result = q->data[q->head];
    if (++q->head >= q->capacity)
        q->head = 0;
    lock_release(&q->lock);
    sema_up(&q->free);
    return result;
}

int queue_count(struct queue *q) {
    lock_acquire(&q->lock);
    int count = q->count;
    lock_acquire(&q->lock);
    return count;
}

```

Listing 9.27: A variant of the bounded queue that also provides `queue_count`.

ignore the waiting behavior of `queue_count`, but carefully note the potentially infinite waiting behavior of `queue_push` and `queue_pop`.

As you can see, a rule of thumb is that we can ignore situations where it is possible to use locks, but need to include situations where we need to use semaphores or condition variables. Note, however, that since we can replace the lock with a semaphore, it is not enough to look at what was used in a particular implementation. This is why the above rule of thumb is phrased as “is it possible to use a lock?” rather than “was a lock used?”. Another implicit

assumption is that locks are acquired and released in the same function (as is good practice), and that threads never wait for semaphores while they are holding a lock.

```

void thread(struct queue *q, int start) {
    for (int i = 0; i < 3; i++) {
        int to_push = expensive(i + start);
        queue_push(q, to_push);
    }
}

int main(void) {
    struct queue *q = queue_create(2);
    thread_new(&thread, q, 1);

    while (queue_count(q) > 0) {
        int item = queue_pop(q);
        printf("Got item: %d\n", item);
    }

    queue_free(q);
    return 0;
}

```

Listing 9.28: The program in `queue-6.c` that uses the `queue_count` function in Listing 9.27. Note, however, that the program does *not* work as expected.

It is worth noting that it is usually not a good idea to add `queue_count` to a bounded queue for the same reason it is not a good idea to be able to read the counter in a semaphore. Namely, it is usually not meaningful to inspect the number of elements in the queue, since another thread may add or remove elements at any point in time. For example, it might be tempting to modify the `main` function as shown in `fig:abstractions-queue-6b` in order to accept an arbitrary number of elements from the other thread. It does, however, not work as expected. For example, if `main` calls `queue_count` before `thread` pushes an element, it will incorrectly determine that no more elements will be available and it therefore terminates.

9.5 Summary: Thread Safety

We conclude the chapter with an overview of the important points related to thread safety that we covered in this chapter. The bullet points below cover our overarching observations and expectations of abstractions in a concurrent system. As we have noted before, individual abstractions may deviate from our expectations below, but if they do so we expect them to highlight these deviations explicitly (i.e., these expectations apply unless otherwise noted).

- We say that a function is thread safe if it is possible to call it from multiple threads concurrently. As long as functions do not use global data (e.g., global variables), they are thread safe by default. If they use global data, additional synchronization is needed.

- When calling functions (even thread-safe functions), it is generally expected that the parameters to the function are not used by other threads during the time of the call.
- In general, we expect that abstract data types can be used from multiple threads as long as each instance is not shared between multiple threads. This follows from the two points above: the individual operations are typically thread safe functions since all state exists in the abstract data type itself (assuming no static/class variables are used), and they expect that parameters to the function are not shared. As such most abstract data types behave like this without additional work.

This means that our expectations of abstract data types are the same as for built in types (e.g., integers, pointers, etc.). That is, we can safely use multiple instances concurrently, but if we share an instance, we need to be careful to avoid data races.

- We say that an abstract data type is thread safe if it is possible to use it from multiple threads concurrently without further synchronization from code that uses them. This typically requires that the implementation of the abstract data type uses some form of synchronization to avoid data races. Since thread safe data types require extra work (both from the author and during runtime), we expect that it is explicitly noted if an abstract data type is thread safe.
- Some data structures (here called concurrent abstract data types), are thread safe abstract data types that may also cause threads to wait for other operations to be called on the data structure. In this regard, they are very similar to the synchronization primitives introduced in this book. Because of this, concurrent abstract data types may be used to synchronize access to other data (e.g., the future and the bounded buffer). This also means that it is possible to use them in a way that causes the program to deadlock.

Note that we only expect these data types to note situations where they wait for some other operation to be called (e.g., `sema_down` may wait for `sema_up`, or `queue_pop` may wait for `queue_push`). We are usually not concerned with situations where a thread needs to sleep temporarily to wait for some other operation to complete, since this does not affect how the abstraction is used.

9.6 Exercises

1. The file `random-1.c` implements the function `random()`, which is a (very poor) generator of pseudo-random numbers. Its interface resembles the `rand()` function in the C standard library. As you can probably see, it is *not* thread safe.

Use appropriate synchronization mechanisms to make it thread-safe. You can use Progis to verify your solution.

2. To avoid getting the same sequence of number each time, it is common to be able to *seed* random number generators (i.e., set the internal state of the number generator). The C standard library provides the function `srand` for this purpose.

The file `random-2.c` extends the program in `random-1.c` to also provide a `seed()` function that can be used to initialize the stream of random numbers. In this case, both threads seed the random number generator to get a predictable stream of numbers. However, since the current implementation is not thread-safe, the program does not work as expected.

Make `random` and `seed` thread-safe. Note that since there are now multiple functions that need to share data, you will need to add parameters to `random` and `seed` (perhaps introducing a new ADT). You can use Progis to verify your solution.

3. The file `barrier.c` contains an outline of an implementation of a *barrier*. This implementation has one operation apart from its initialization: `barrier_sync`. The usage is illustrated in the `work` function. The idea is that n threads all execute code in the same function, but the threads sometimes need to wait for each other to coordinate their work. When this is needed, the threads call `barrier_sync`. The `barrier_sync` function will simply wait until all threads have called `barrier_sync`. Once that happens, all threads are allowed to continue. The number of threads that `barrier_sync` should wait for is specified as a parameter to `barrier_create`.

Complete the implementation of `struct barrier`. You will need to introduce additional variables, both numbers and synchronization primitives. You can verify your solution with Progis.

Note: It is probably easiest to use locks and condition variables to implement the barrier.

Note: It should be possible to use the barrier more than once.

Atomics

In Chapter 3 we saw that programs that utilize multiple threads need to carefully coordinate the threads, both to avoid *data races* and to ensure that things happen in the right order. To achieve this, we have used *synchronization primitives*, particularly *semaphores* (Chapter 4), *locks* (Chapter 5), and *condition variables* (Chapter 8). So far we have considered them to be special in that they possess the unique ability to make threads *wait*, and that they can be shared between multiple threads without protection (i.e., they are thread safe).

Even though we have treated the synchronization primitives as special, it turns out that it is possible to implement them in terms of lower-level operations that the system provide. One class of such operations, *atomic operations*, is particularly interesting since we can use them separately in certain situations in order to reduce the run-time cost of synchronization.

As such, this chapter will first examine how locks can be implemented. This leads us to atomic operations, and the remainder of the chapter will explore how atomic operations can be used in isolation to protect shared data and coordinate threads. However, as you will see it is not easy to use atomic operations correctly. As such, they are typically only used to implement central synchronization mechanisms (e.g., synchronization primitives or concurrent abstract data types), or in simple cases where the usage is obvious. Since this book is an introduction to concurrent programming, this chapter only focuses on the fundamentals of atomic operations. This is enough to use atomic operations in the obvious places (i.e., where it is easy to use them correctly), and to know the important pitfalls when using atomic operations. For the curious reader, the last section contains a very brief overview of more advanced usage of atomic operations.

10.1 Implementing a Lock

To introduce atomic operations, we will first examine how locks are implemented. Up until now, we have treated locks, and semaphores as a special case. In Chapter 3, we concluded that data shared between threads must not be accessed concurrently. However, even though locks and semaphores are evidently some form of data (they are implemented as a `struct` after all), this rule does not seem to apply to them. One explanation is that locks and semaphores

are implemented as thread-safe (and concurrent) abstract data types. However, as we saw in Chapter 9, we typically need to use either a lock or a semaphore to implement a thread-safe abstract data structure. This seems problematic.

To illustrate the problem, we will attempt to implement a simple lock. To keep the implementation simple, we will not provide the error checking described in Chapter 5. That is, we will not verify that a thread that acquires the lock does not already hold the lock, nor that a thread that attempts to release the lock actually holds the lock. It is worth noting that the same problems we will see with the lock also arises when trying to implement a semaphore, and the same idea can be applied to solve the problems. The reason we examine locks rather than semaphores is simply that the implementation is a bit simpler since locks only need to keep track of two states (*locked* and *unlocked*) in contrast to the integer in a semaphore.

```
struct my_lock {
    bool acquired;
};

void my_lock_init(struct my_lock *l) {
    l->acquired = false;
}

void my_lock_acquire(struct my_lock *l) {
    while (l->acquired)
        ;
    l->acquired = true;
}

void my_lock_release(struct my_lock *l) {
    l->acquired = false;
}
```

Listing 10.1: Initial attempt at implementing a lock.

An initial attempt at implementing a lock is shown in Listing 10.1, and in the file `lock-1.c`. As shown in the listing, the lock (`struct my_lock`) contains a single variable, `acquired`, that stores whether the lock is currently acquired by some thread or not. In `my_lock_init`, the `acquired` variable is initialized to `false` as we would expect. The function `my_lock_acquire` acquires the lock in two steps. It first waits for `acquired` to become false by repeatedly checking it in a `while` loop. As soon as `acquired` is `false`, the function concludes that it can acquire the lock and therefore sets `acquired` to `true` to disallow other threads from acquiring the lock. The lock can then be released by `my_lock_release`, which simply sets `acquired` to `false` again.

Based on our knowledge from earlier chapters, we can quickly see that this implementation of `struct my_lock` will not work correctly. If multiple threads call `my_lock_acquire` and/or `my_lock_release` concurrently, they will access `acquired` concurrently which constitutes a data race. As such `struct my_lock` is not currently thread safe. This is a problem since we need locks to be thread safe in order to be able to use them for mutual exclusion, so that we can make other data structures thread safe.

Fortunately, we know how to make `struct my_lock` thread safe! Since we have a non-synchronized version that uses loops to wait for a condition, we can simply apply the method presented in Section 8.3.1. In particular, we note that the entire `my_lock_acquire` and `my_lock_release` manipulate `acquired`, and use a lock to synchronize them. After that, we note that the `while` loop in `my_lock_acquire` waits for `acquired` to become `false`. Therefore we add a call to `cond_wait` inside the loop, and a call to `cond_signal` after setting `acquired` to `false` in `my_lock_release`. This results in the code in Listing 10.2.

```
struct my_lock {
    bool acquired;
    struct lock acquired_lock;
    struct condition acquired_cond;
};

void my_lock_init(struct my_lock *l) {
    l->acquired = false;
    lock_init(&l->acquired_lock);
    cond_init(&l->acquired_cond);
}

void my_lock_acquire(struct my_lock *l) {
    lock_acquire(&l->acquired_lock);
    while (l->acquired)
        cond_wait(&l->acquired_cond, &l->acquired_lock);
    l->acquired = true;
    lock_release(&l->acquired_lock);
}

void my_lock_release(struct my_lock *l) {
    lock_acquire(&l->acquired_lock);
    l->acquired = false;
    cond_signal(&l->acquired_cond, &l->acquired_lock);
    lock_release(&l->acquired_lock);
}
```

Listing 10.2: Thread safe version of the lock implementation in Listing 10.1, available as `lock-2.c`.

While this implementation does indeed work correctly (you can verify this using the code in `lock-2.c`), it does not solve our problem. Remember that our goal was to implement a lock. However, we ended up needing a lock in order to implement a lock. If we look back to what happened, this is not too surprising. After all, locks are an abstract data type like any other. Since the lock needs to contain at least one variable to store if the lock is acquired or not, we will need to ensure that the shared variable is not accessed concurrently somehow. So far, the only tools we know that can help us are locks and semaphores. Since we can not implement a lock by using a lock, we need some other way to access shared data safely in order to implement the lock.

10.2 Atomic Operations

All multicore CPUs have a set of specialized instructions that make it possible to manipulate shared data safely. These instructions are unique since the hardware ensures that the possibly multiple steps involved in the instructions (e.g., load from memory, update, store to memory) are executed as a unit that other threads are not able to interrupt. This means that even if multiple threads access the same data, these accesses do not cause a data race. Since these instructions can not be interrupted during execution, they are often called *atomic instructions*.

Since different CPU architectures provide different atomic instructions with slightly different semantics it is not practical to use atomic instructions directly, particularly when writing portable software. For this reason, many programming languages provide a set of higher-level *atomic operations* in the form of library functions that are implemented using relevant atomic instructions. This makes it possible for the language to provide a set of atomic operations with consistent semantics across different CPUs. Since the atomic operations are known to the compiler, it can make sure to not reorder or remove atomic operations during optimizations. In particular, this prevents both the compiler and the CPU from transforming the program in ways that would cause it to misbehave. As such, shared data accessed *only* through atomic operations *do not* constitute a data race.

The most fundamental atomic operations are perhaps `atomic_read` and `atomic_write` that allow reading and writing to shared data from multiple threads. In the threading library provided with this book, they are generic, meaning that they are usable for all integer and pointer types. To illustrate this, we denote an unknown type using the letter `T`. The semantics of the two functions are as follows:

```
T atomic_read(T *value);
```

Reads the value stored at `value` using an atomic instruction and returns it.

```
void atomic_write(T *value, T write);
```

Writes the value `write` to the location pointed to by `value` using an atomic instruction.

It is important to note that even though we refer to these operations as *atomic operations*, not all memory accesses are atomic. In general, only *one* of the values accessed by an atomic operation is performed in atomic manner. All other inputs and outputs are accessed normally, which means that they may cause data races.

Since the goal of atomic operations is to access memory in an atomic manner, we need to pass the address of the memory that should be accessed to the function.¹ The reason is that parameters are always passed by value in C, so any non-pointer variables will be accessed *before* the call to the function, before the function has a chance to execute any atomic instructions.

To illustrate this in more detail, consider the program `atomic-access.c` depicted in Listing 10.3. The program defines two integers, `a` and `b`, calls

¹In languages that support references, we could also pass a reference.

```

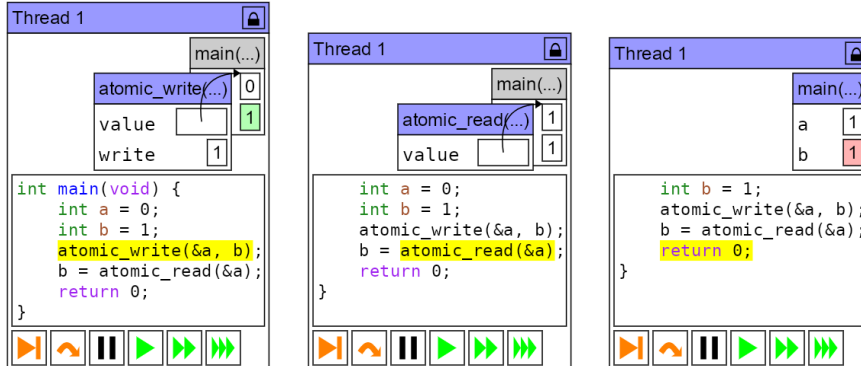
int main(void) {
    int a = 0;
    int b = 1;
    atomic_write(&a, b);
    b = atomic_read(&a);
    return 0;
}

```

Listing 10.3: Example illustrating that only the pointer parameters are accessed atomically.

`atomic_write(&a, b)` to store the value of `b` into `a`, and then calls `atomic_read(&a)` to read the value of `a` and store it into `b`. As mentioned before, only accesses to `a` are atomic even though we use atomic operations.

We can see the reason for this by running the program in Progvis. If we step the program until we reach the start of `atomic_write`, we will see the representation in Fig. 10.1a. At this point, `main` has just computed the values of both parameters and passed them to `atomic_write`. As we can see from the stack, the first parameter (`value`) is a reference to `a`, which does not involve reading `a`. However, the second parameter (`write`) is a *copy* of `b`. As such, `main` needs to read `b` *before* the call to `atomic_write`, and as usual `main` will perform a normal, non-atomic memory access.



(a) Start of `atomic_write`. (b) Start of `atomic_read`. (c) After `atomic_read`.

Figure 10.1: Progvis' visualization of different steps in the program `atomic-access.c`, using `Run ⇒ Report data races ⇒ Only statements` for clarity.

A similar thing happens with the call to `atomic_read`. As shown in Fig. 10.1b, `atomic_read` only gets a pointer to the value it should read (`from`). As such, it will read the value atomically and return it. It is then up to `main` to store the value in `b`, which it will again do using a normal, non-atomic memory access, as can be seen in Fig. 10.1c. It is worth noting that the variable `a` was only highlighted red when it was initialized to zero, and not when atomic operations updated it. The reason is that Progvis do not track atomic accesses to data since they are known to be safe.

Now that we know how to use `atomic_read` and `atomic_write` to manipulate shared data without causing data races, we might be tempted to insert `atomic_read` wherever we read from `acquired` and `atomic_write` wherever we write to `acquired`. This gives us the implementation in `lock-3.c` (Listing 10.4).

```

struct my_lock {
    bool acquired;
};

void my_lock_init(struct my_lock *l) {
    l->acquired = false;
}

void my_lock_acquire(struct my_lock *l) {
    while (atomic_read(&l->acquired))
        ;
    atomic_write(&l->acquired, true);
}

void my_lock_release(struct my_lock *l) {
    atomic_write(&l->acquired, false);
}

```

Listing 10.4: Attempting to use atomic operations to eliminate the data races in the lock implementation.

Practice: The implementation in `lock-3.c` (Listing 10.4) does indeed not contain any data races since we use atomic operations to access `acquired` (apart from in `my_lock_init`, which is not a problem). Does this mean that the implementation works as expected?

As a first step, try to reason about the code yourself. In particular, think back to the contents of Chapter 6 (critical sections). You can use *Run ⇒ Look for errors...* in Progviz to verify your answer with the test program in `lock-3.c`.

Since we use atomic operations to access `acquired` in `lock-3.c` we eliminate the data races in the previous version. This means that the program is well-defined according to the C language. However, in spite of this the implementation in `lock-3.c` does not meet our expectations. In particular, it is possible for two threads to acquire the lock at the same time. Assume that two threads, thread 1 and thread 2, call `my_lock_acquire` to acquire the same lock at the same time. It is then possible for both threads to execute `atomic_read` before the other thread updates `acquired`. As such, the call to `atomic_read` will return `false` for both threads, and both will proceed to acquire the lock by calling `atomic_write` and returning from `my_lock_acquire`.

Hopefully this problem looks familiar. In fact, it is the same kind of problem that was covered Chapter 6. The difference we used locks to synchronize access to shared data in Chapter 6, but now we are using atomic operations. The conclusion in Chapter 6 was that it is not enough to simply surround each access to shared variables with calls to `lock_acquire` and `lock_release` individually,

but that we need to protect the entire critical section as a whole, unbroken region. The same is true when using atomic operations. We can not simply replace all accesses to shared variables with `atomic_read` or `atomic_write`. Since the language only guarantees that atomic operations are atomic in isolation, this is equivalent to protecting individual statements with a lock.²

Just as we did with locks in Chapter 6, the solution to the problem is to ensure that all operations that depend on each other happen as a unit. With locks, this was quite easy. We could just move the calls to `lock_acquire` and `lock_release` as we wish. This is not the case when using atomic operations, since each atomic operation starts and ends their own atomic operation (i.e., they have calls to `lock_acquire` and `lock_release` built-in). Since neither of `atomic_read` and `atomic_write` allows us to both read *and* write *acquired* we need to find a different approach.

10.2.1 Test and Set

Luckily, there are atomic operations that *both* read and write to memory. The simplest one is perhaps *test and set*. It behaves like the code below, but ensures that the two accesses to the memory `value` points to occurs atomically.

```
T test_and_set(T *value) {
    T old = *value;
    *value = 1;
    return old;
}
```

In essence, `test_and_set` reads the current value of `*value`, replaces it with the number 1, and returns the old value. It is worth noting that boolean variables are just integers that contain the value 0 or 1 in C, so it is possible to use `test_and_set` for booleans as well.

It seems likely that `test_and_set` is able to solve our problems since it is able to both read and write to memory. Furthermore, the value that `my_lock_init` wishes needs to write to memory is 1, just what `test_and_set` is able to provide! The question that remains is: how do we achieve this? The calls to `atomic_read` and `atomic_write` in `lock-3.c` do not occur together. Furthermore, we do not *always* write `true` to `acquired`, we only do this once we have observed that `acquired` is `false`.

To replace the two atomic operations in `my_lock_acquire` with a single `test_and_set` we need to be a bit creative. Our goal is to refactor the code so that `atomic_read` and `atomic_write` end up next to each other in a way that resembles how `test_and_set` works. One way to do this is shown in Fig. 10.2.

In step 1, we start by moving the call to `atomic_write` inside the loop. To avoid multiple statements in the condition of the `while` statement,³ we introduce a control variable, `was_acquired`. This makes it straight-forward to rewrite the loop as shown in step 1 in Fig. 10.2.

To get to step 2, we need to make a slight adjustment to the behavior of `my_lock_acquire`. So far, we only set `acquired` to `true` if it was previously `false`. However, we can make one important observation. Regardless

²For the purposes of this discussion, we can think of all atomic operations as if they acquire a global lock at the start and release the lock at the end.

³It *is* possible with the comma operator, but it is *not* readable.

```

1: void my_lock_acquire(struct my_lock *l) {
    bool was_acquired;
    do {
        was_acquired = atomic_read(&l->acquired);
        if (!was_acquired)
            atomic_write(&l->acquired, true);
    } while (was_acquired);
}

2: void my_lock_acquire(struct my_lock *l) {
    bool was_acquired;
    do {
        was_acquired = atomic_read(&l->acquired);
        atomic_write(&l->acquired, true);
    } while (was_acquired);
}

3: void my_lock_acquire(struct my_lock *l) {
    bool was_acquired;
    do {
        was_acquired = test_and_set(&l->acquired);
    } while (was_acquired);
}

```

Figure 10.2: Step by step rewrites of `lock-3.c` to get `atomic_read` and `atomic_write` next to each other so that we can replace them with `test_and_set`. Available as `lock-4-1.c`, `lock-4-2.c`, and `lock-4-3.c` respectively.

of whether `acquired` is `true` or `false` at the start of the loop, we know that `acquired` will always be `true` at the end of the loop. This means that it does not hurt if we *always* set `acquired` to `true` in the loop (that is, assuming we manage to do it atomically with the read). This observation allows us to remove the `if` statement in the loop and always execute `atomic_write`.

At this point, the code in step 2 in Fig. 10.2 is almost identical to the pseudocode for `test_and_set`: we first read the current value of `acquired` and store it in `was_acquired`, then we set it to `true`. As such, we can simply replace `atomic_read` and `atomic_write` with `test_and_set` as shown in step 3 of Fig. 10.2. Since `test_and_set` manages to both read and write to `acquired` as a single atomic unit, we have now reached our goal of including both operations in the critical section, and therefore the implementation now works as intended.

To verify that our implementation is correct, we use *Run \Rightarrow Look for errors...* in Progviz. As we expect, it does not find any data races, neither from the lock itself or because `struct my_lock` behaves incorrectly. It does, however, find a different issue: that `my_lock_acquire` uses busy wait when threads need to acquire the lock. As we saw in Chapter 8, we usually want to avoid busy wait since it is wasteful to let threads wait by repeatedly reading and checking the value of a shared variable. A better option would be to inform the scheduler that the thread has no useful work to do, so that other threads that might

have useful work to do can be scheduled.

The difference is quite clearly visible in Progvis. Whenever a thread needs to wait to acquire a normal `struct lock`, Progvis adds a box with the title *Waiting* and an arrow to the thing the thread is waiting for. Furthermore, it is not possible to step the thread while the *Waiting* box is visible. This is not what happens when a thread needs to wait to acquire a `struct my_lock`. In this case, no *Waiting* box appears, and it is still possible to step the thread, even if it is stuck spinning in the loop inside `my_lock_acquire`. For this reason, locks like `struct my_lock` are sometimes called *spinlocks*.

Since atomic operations only modify memory, we can *not* use them to make threads wait properly. To do this, it is necessary to communicate with the scheduler in some way. The exact mechanism for doing this depends on the operating system, so we will not cover it in depth here. It is, however, worth knowing that at least Linux and Windows provide system calls that resemble `cond_wait` and `cond_signal`, but that let threads wait for values in memory to change. At the time of writing, these are not exposed in a platform-independent way in C.⁴

Even though spinlocks rely on busy wait, they are useful in certain situations. It is comparatively time-consuming to switch between threads,⁵ so if we are running on a system with multiple CPU-cores, know that the thread we are waiting for is currently running on a different CPU core, and know that the other thread will release the lock in a short amount of time (i.e., the critical section is short, perhaps only a 50 instructions or so), it is actually cheaper to use busy wait than to wait properly. As you might imagine, it is hard to make sure that all of these conditions are met in large systems (especially outside of the kernel), so most lock implementations combine the two approaches. They first spin in a busy-wait loop until either the lock becomes available, or a maximum number of iterations have been reached. If the lock becomes available during this time, we avoid an expensive thread switch. If the maximum number of iterations was reached, the lock assumes that the thread needs to wait for a long time, and asks the scheduler to switch to another thread to avoid wasting CPU time. This way we can get most benefits from both approaches.

With the discussion about busy wait and spinlocks out of the way, the conclusion is that we can not avoid busy wait in our implementation of `struct my_lock` while using only atomic operations.⁶ As such, we tell Progvis that the busy wait in the implementation is intentional by adding a call to the function `intentional_busy_wait()` inside the loop in `my_lock_acquire`. We can also remove the variable `was_acquired` again since the loop only contains one statement. This gives us the final implementation in `lock-5.c` (Listing 10.5), which works correctly according to Progvis.

⁴Although C++ 20 introduces `wait()`, `notify_one()`, and `notify_all()` in `std::atomic<T>`, which provides this functionality

⁵It requires executing a couple of hundred or thousand instructions, and typically incurs cache and TLB misses which are costly.

⁶We can of course use a semaphore once we decide to wait, but then we are back to implementing locks/semaphores using locks/semaphores again.

```

struct my_lock {
    bool acquired;
};

void my_lock_init(struct my_lock *l) {
    l->acquired = false;
}

void my_lock_acquire(struct my_lock *l) {
    while (test_and_set(&l->acquired))
        intentional_busy_wait();
}

void my_lock_release(struct my_lock *l) {
    atomic_write(&l->acquired, false);
}

```

Listing 10.5: Final implementation of `struct my_lock` using `test_and_set`. Note that we use `intentional_busy_wait()` to inform Progviz that we are aware of the busy wait in the code.

10.2.2 Atomic Swap

Most languages provide additional atomic operations that are more flexible than `test_and_set`. The operation `atomic_swap` is one of them. Similarly to `test_and_set`, it reads a value from memory and replaces it with another value. The big difference is that rather than always replacing the value with 1, it uses an user-supplied value.

The operation behaves like the code below. As with `test_and_set`, the accesses through the pointer `value` are atomic, and both happen as a unit that other threads are not able to interrupt.

```

T atomic_swap(T *value, T replace) {
    T old = *value;
    *value = replace;
    return old;
}

```

Before we examine how `atomic_swap` can be used, it is worth noting that there are many different names for the `atomic_swap` operation. The name used here is chosen to highlight the similarity to *compare and swap* or *CAS*, which is quite well known. However, since the interface of `atomic_swap` differs from a typical `swap` function that swaps two elements in memory, some languages prefer to use the word *exchange* instead.⁷

To illustrate the difference, consider the program in Listing 10.6. It starts by defining a typical `swap` function that accepts two pointers to integers and swaps the values pointed to. It is used in `main` to swap the values of the variables `x` and `y`. However, we can not use `atomic_swap` in the same way since `atomic_swap` accepts one pointer and one value. The reason for this difference is that atomic

⁷For example, `std::atomic` in C++ uses the names `exchange` and `compare_exchange` for `atomic_swap` and `compare_and_swap` respectively.

```
void swap(int *a, int *b) {
    int tmp = *a;
    *a = *b;
    *b = tmp;
}

int main(void) {
    int x = 1;
    int y = 2;

    swap(&x, &y);

    y = atomic_swap(&x, y);

    return 0;
}
```

Listing 10.6: A program that illustrates the different interfaces of `swap` and `atomic_swap`. Available as `swap.c`.

operations are typically only able to access *one* value in memory atomically. Therefore, `atomic_swap` only accepts *one* of its parameters as a pointer in order to make it obvious which value is accessed atomically.

Of course, it is possible to implement a version of `atomic_swap` that accepts two pointer parameters. A possible implementation of this is shown in Listing 10.7. However, as the name of the function (`bad_atomic_swap`) implies, this is not a good idea. Since the implementation uses `atomic_swap`, the function `bad_atomic_swap` still only accesses `*a` in an atomic manner. In particular, we can clearly see that `*b` is accessed twice: once to read the value as a parameter to `atomic_swap`, and written once to store the value returned from `atomic_swap`. However, this fact is *not* apparent from the functions interface. Rather, since the function accepts two parameters, the user is led to believe that *both* values are accessed atomically and will therefore most likely use `bad_atomic_swap` incorrectly at some point. This is the reason why `atomic_swap` is different from a normal swap, to highlight an important detail in its functionality!

```
void bad_atomic_swap(int *a, int *b) {
    *b = atomic_swap(a, *b);
}
```

Listing 10.7: A possible implementation of a version of `atomic_swap` with the same interface as `swap`. As the name implies, this is not a good idea.

After the slight digression into the importance of clear API design, we will continue to examine how `atomic_swap` can be used. Since `atomic_swap` is essentially a generalization of `test_and_set`, we can quite easily use it to implement a lock by simply replacing `test_and_set(...)` with `atomic_swap(..., 1)` as shown in Listing 10.8.

Since `atomic_swap` allows us to write any value atomically, it is more versatile than `test_and_set`. For example, we can use it to implement `atomic_`

```

void my_lock_acquire(struct my_lock *l) {
    while (atomic_swap(&l->acquired, true))
        intentional_busy_wait();
}

```

Listing 10.8: Using `atomic_swap` to implement a lock. Available as `lock-6.c`.

`write(value, write)` by simply calling `atomic_swap(value, write)` and discarding the result.

10.2.3 Compare and Swap

Another common atomic operation is `compare_and_swap` (often abbreviated CAS), which provides even more flexibility compared to `test_and_set` and `atomic_swap`. The previous two operations allow reading and writing to memory as a single atomic operation. As its name implies, compare and swap adds a comparison into the atomic operation. That is, it first reads the current value in memory and compares it to an expected value. If they are equal, the value in memory is replaced with a new value. As before, the entire operation (i.e., read, compare, and possibly write) is atomic. The behavior can be expressed as the code below.

```

T compare_and_swap(T *value, T compare, T replace) {
    T old = *value;
    if (old == compare)
        *value = replace;
    return old;
}

```

Since `compare_and_swap` includes a comparison, it is quite natural to implement a lock using `compare_and_swap`. Recall that the goal of the `my_lock_acquire` function was to check `acquired`. If `acquired` was `false`, we wished to set it to `true`. If not, we did not want to modify `acquired` at all, and just try again later. This can be expressed as a single call to `compare_and_swap(&acquired, false, true)`. The only detail that might be counter-intuitive is that `compare_and_swap` returns the *old* value of `acquired`. This means that we have successfully acquired the lock if `compare_and_swap` returns `false` (i.e., the same value as the second parameter). As such, the `my_lock_acquire` function can be implemented as in Listing 10.9.

```

void my_lock_acquire(struct my_lock *l) {
    while (compare_and_swap(&l->acquired, false, true))
        intentional_busy_wait();
}

```

Listing 10.9: Using `compare_and_swap` to implement a lock. The full source code is available as `lock-7.c`.

It is worth noting that there are a number of different variations of `compare_and_swap`. First and foremost, languages that use the name `exchange` for atomic_

`swap` often name it `compare_exchange` for consistency. Another variation is that some implementations do not return the old value that `*value` pointed to, but instead return a boolean that indicates whether the value was updated or not. This makes it a bit more convenient to use in some situations (e.g., the ones covered in Section 10.3.2). We can implement it using the `compare_and_swap` presented above as shown in Listing 10.10.

```
bool cas_bool(T *value, T compare, T replace) {
    int old = compare_and_swap(value, compare, replace);
    return old == compare;
}
```

Listing 10.10: Implementation of a variant of `compare_and_swap` that returns a boolean rather than the old value of `*value`. Named `cas_bool` for brevity.

The variant `cas_bool` from Listing 10.10 is actually less powerful than the `compare_and_swap` that returns the old value. For example, we can implement `atomic_read` by calling `compare_and_swap(&value, 0, 0)`. This is not possible with `cas_bool`. Since we can implement `cas_bool` using `compare_and_swap`, we can conclude that `compare_and_swap` is strictly more powerful than `cas_bool`. This is a big reason why we (similarly to the compiler intrinsics provided by e.g., GCC) use this version of compare and swap.

10.2.4 Properties of Locks

To conclude the section about implementing locks using atomic operations, we will examine what we typically expect from locks and other synchronization primitives in more detail. We will also see that our simple lock implementation does *not* fulfill all of our expectations. This illustrates why it is often a good idea to rely on pre-existing synchronization libraries, but also highlights some aspects that are important to pay attention to when working with atomic operations in general.

When using a lock, we expect it to satisfy the following properties:

Mutual exclusion We expect that the lock provides *mutual exclusion* to whichever resource the lock is protecting. That is, once a thread has successfully achieved the lock, no other thread should be able to acquire the lock until the first thread has released it.

Progress If multiple threads are trying to acquire the same lock concurrently, we expect the implementation to decide which of the threads should be allowed to acquire the lock within a bounded amount of time.

Bounded waiting If a thread wishes to acquire the lock, it should be able to acquire the lock within a bounded amount of time.

As mentioned before, we also expect locks to wait efficiently by avoiding busy waiting. However, this is typically not included in the list of formal properties. The reason is that busy wait only makes the program less efficient, it does not make the program behave incorrectly. As we will see, breaking one of the three properties above may lead to unwanted behavior.

The first property, *mutual exclusion*, is perhaps the most unsurprising one. The reason why we use locks is that we wish to avoid data races by introducing mutual exclusion. As such, it is little use if the lock does not manage to enforce mutual exclusion. However, this is not always easy as we saw in our initial attempt in `lock-1.c`. Our finished implementations (`lock-5.c`, `lock-6.c`, and `lock-7.c`) do however satisfy this requirement.

The second property, *progress*, is perhaps less obvious. It states that if two or more threads wish to acquire the same lock, the implementation needs to decide which of the threads should “win” and acquire the lock within a bounded time. This should hold regardless of how the threads started waiting for the lock. Perhaps they called `lock_acquire` at the same time, perhaps multiple threads were waiting for a lock that was released by some other thread. In particular, this means that it should *not* be possible for the implementation to delay the decision of which thread should be allowed to acquire the lock for an infinite time. While it is not stated, we would expect that the decision should be reached quickly as well (i.e., we do not want to wait for multiple seconds for the system to decide).

All of our finished lock implementations satisfy this property. Interestingly, they satisfy the property without us having to do anything special. It could, however, be a concern if we aimed to reach mutual exclusion by always preferring to let other threads acquire lock before them.

The final property, *bounded waiting*, is related to *progress*. However, instead of looking at the system as a whole, it looks at a single thread. That is, once a thread starts waiting to acquire a thread, it should eventually be able to acquire the lock. To understand the property it might be useful to imagine the lock as a queue of threads that are waiting to acquire the lock. *Progress* means that the lock needs to pick a thread from the queue in finite time. *Bounded waiting*, on the other hand, means that the lock should pick threads so that no individual thread has to wait forever (i.e., other threads should not be able to skip the line “too much”). Note that it is implied here that we assume that all threads that have successfully acquired a lock will eventually release it.

This property is *not* satisfied by our lock implementations. In our implementations, whichever thread happens to execute the atomic operation inside `my_lock_acquire` first wins and gets to acquire the lock. As such, if we have two threads that acquire and release the same lock in a loop, it is possible (but unlikely) for one thread to always be faster than the other, and thereby the other thread has to wait forever.

Practice: The programs `lock-bounded-a.c` and `lock-bounded-b.c` contains a small test program to illustrate the issue with bounded waiting. Both programs contain the code below that acquires and releases a lock in an infinite loop. The difference is that `lock-bounded-a.c` uses our custom lock implementation while `lock-bounded-b.c` uses the lock implementation from the threading library.

```
int shared;
struct my_lock shared_lock;

void use_lock(void) {
    while (true) {
        my_lock_acquire(&shared_lock);
        shared += 1;
        my_lock_release(&shared_lock);
    }
}

int main(void) {
    my_lock_init(&shared_lock);
    thread_new(&use_lock);
    use_lock();
    return 0;
}
```

Using Progviz, find a situation in `lock-bounded-a.c` where one of the threads has to wait forever to acquire the lock. Remember that the scheduler needs to allocate *some* CPU time to all threads eventually. As such, while you may choose in what order and how often you step each of the two threads, you need to give at least some CPU time to both threads eventually (i.e., simply deciding to never step one of the threads is not a valid solution).

Once you have found an issue in `lock-bounded-a.c`, try to do the same thing in `lock-bounded-b.c`. Note that this is not possible since `struct lock` ensures bounded waiting.

As you have noted above, since our implementation does not fulfill *bounded waiting*, it is possible (albeit unlikely) that some thread may have to wait for a long time, or even forever. This is called *starvation* since the thread is essentially starved for CPU time. Since it is likely that only a few of the threads that use the lock are starved, the symptoms of starvation could be that a part of a larger system appears to freeze while other parts may work normally.

It is worth noting that *bounded waiting* does not imply that locks need to be entirely fair. For example, in `lock-bounded-a.c`, the lock does not need to ensure that thread 2 always acquires the lock after thread 1 and vice versa. It is acceptable for thread 1 to acquire the lock hundreds of times before thread 2 gets its chance, as long as the number of times that thread 1 can skip the queue ahead of thread 2 is bounded (and ideally not too large). In practice, it is often beneficial to have locks *not* be entirely fair, since it potentially avoids thread switches and reduces the need to synchronize caches between multiple threads.

Finally, it is worth noting that even though the properties *mutual exclusion*, *progress*, and *bounded waiting* are often expressed in terms of locks, they apply to the other synchronization primitives as well. Similarly, they are useful to consider when working with condition variables or when implementing custom synchronization using atomic operations. For example, to avoid starvation.

Progress and *bounded waiting* are straight-forward to apply to both semaphores and condition variables. Progress simply means that whenever the primitive needs to wake one out of potentially many eligible threads, it needs to decide which one to wake up in a bounded amount of time. Similarly, bounded waiting, means that it should not be possible for other threads to always get priority over some other thread (e.g., one of the threads waiting in `sema_down` waits forever even though `sema_up` is called multiple times).

The final property, *mutual exclusion*, is perhaps less clear how to apply to other situations, particularly since semaphores and condition variables are not always used to create mutual exclusion. For them, it is perhaps more useful to think that mutual exclusion should apply to the data inside the synchronization primitive. That is, it should appear *as if* all operations on the synchronization primitive are executed in some order.

10.3 Lock-Free Access to Shared Data

The main reason why languages provide atomic operations is not to allow programmers to build their own locks. As we saw in the previous section, we are often better off using the synchronization primitives provided by the language directly.⁸ Rather, the main usage of atomic operations is to allow threads to share data without incurring the costs involved with locks. Even though it is relatively cheap to acquire and release a lock, the cost can easily be larger than the actual work in a short critical section.

This is one reason why *lock-free data structures* are used in certain situations. A *lock-free data structure* is simply a thread-safe data structure that does not use lock to ensure thread safety. As such, the main selling point of lock-free data structures is that they are more performant than other data structures.

Many lock-free data structures have been developed and used so far, such as lock-free queues and hash tables. However, the goal of this chapter is to provide an introduction to atomic operations. Therefore we will not cover them in detail. Rather, we will complement expand on what we have already used to implement locks to cover many commonly occurring cases. That is, the goal is to cover the fundamentals, both so that you can use atomic operations in simple and obvious cases to improve performance, but also to give you a foundation for looking into more details on your own.

10.3.1 Single-Step Operations

One observation from our effort in implementing a lock is that it is fairly simple to use atomic operations to update shared data in a single step. That

⁸One use-case is of course to make it possible to implement the synchronization primitives as a part of the runtime library, rather than having to embed them in the compiler. But this is only relevant to language designers and implementers of threading libraries.

is, in cases where all accesses to shared data in a critical section matches what some atomic operation does, we can easily replace the critical section with that atomic operation. For example, `my_lock_release` only assigned `false` to a variable, and it is thereby trivial to replace it with `atomic_write`. Similarly, if some part of the program only needs to read from a variable, we can easily use `atomic_read`.

The difficult part about using atomic operations is when the critical section contains both reads writes and conditionals. We saw this when we refactored `my_lock_acquire` so that we could use atomic operations. As such, it is useful to be aware of a decent number of atomic operations. That way we can use atomic operations without the need to refactor in more cases.

In this subsection we will introduce the atomic operations `atomic_add` and `atomic_sub`, and how they can be used to implement *reference counting*. As the name implies, `atomic_add` and `atomic_sub` adds and subtracts a value from some variable in an atomic manner. They behave like the code below, but ensures that accesses to `value` occur atomically.

```
T atomic_add(T *value, T add) {
    T old = *value;
    *value += add;
    return old;
}

T atomic_sub(T *value, T sub) {
    T old = *value;
    *value -= sub;
    return old;
}
```

These atomic operations are particularly useful to implement *reference counting* in a concurrent system. The core idea behind reference counting is to keep track of the number of references (i.e., pointers in C) that refer to allocations. Once the number of references reaches zero, we know that no part of the program can use the object and we can therefore safely free the memory.

To illustrate the concept, we will add reference counting to the `struct future` we implemented in Chapter 9. The changes are shown in Listing 10.11 (`refcount-1.c`). As shown in the figure, we add an integer variable `refs` that stores the number of references to the future. It is initialized to 1 in `future_create`, since the caller of `future_create` will receive a reference. Furthermore, `future_free` is replaced with `future_retain` and `future_release`. The function `future_retain` is called if we wish to store another reference to the future somewhere. As the implementation shows, `future_retain` simply increases the reference count. Conversely, `future_release` is called before we remove a reference to the future (e.g., before a `struct future` variable goes out of scope). It decreases `refs` by one, and if it reaches zero we know it is safe to deallocate the `struct future`.

The program `refcount-1.c` also contains a simple program that uses `struct future` to post a result to two threads to illustrate how refcounting can be used. As can be seen in Listing 10.11, the main function first creates a future and starts two threads that will get the result from the future. Importantly, since each thread will receive a pointer to the future, the main program calls `future_`

```
struct future {
    int result;
    struct semaphore has_result;
    int refs;
};

struct future *future_create(void) {
    struct future *f = malloc(sizeof(struct future));
    sema_init(&f->has_result, 0);
    f->refs = 1;
    return f;
}

void future_retain(struct future *f) {
    f->refs++;
}

void future_release(struct future *f) {
    if (--f->refs == 0) {
        free(f);
    }
}

// future_post and future_get are the same as before.

void thread(struct future *f) {
    int value = future_get(f);
    printf("Value: %d\n", value);
    future_release(f);
}

int main(void) {
    struct future *f = future_create();
    future_retain(f);
    thread_new(&thread, f);
    future_retain(f);
    thread_new(&thread, f);

    future_post(f, 9);
    future_release(f);

    return 0;
}
```

Listing 10.11: Changes to `struct future` to add reference counting. Available as `refcount-1.c`.

`retain` before each of the threads are created to keep the reference counter accurate. Thus, just before `future_post` is called, the reference count will be 3.

To ensure that the future is deallocated properly, both the main function and the two threads call `future_release` to indicate that their reference is about to disappear. This scheme means that the future will be deallocated regardless of in which order the three threads finish executing.

As you have likely realized already, the implementation is not correct since `ref` is accessed by multiple threads. Progviz will confirm this. As such, we need to protect it either using locks, or atomic operations. Fortunately, we can quite easily replace the `refs++` and `--refs` with `atomic_add` and `atomic_sub`. The only thing we need to be careful about is that both of them return the *old* value of the shared variable, which is not what we want in `future_release`. After doing this, we get the code in Listing 10.12 (`refcount-2.c`), which is correct.

```
void future_retain(struct future *f) {
    atomic_add(&f->refs, 1);
}

void future_release(struct future *f) {
    if (atomic_sub(&f->refs, 1) == 1) {
        free(f);
    }
}
```

Listing 10.12: The reference counting from Listing 10.11 updated to use atomic operations.

Using atomic operations for reference counting is fairly common. First and foremost, it is quite easy to replace non-atomic code with `atomic_add` and `atomic_sub`. Secondly, the *retain* and *release* operations are used frequently, and therefore we will greatly benefit from not having to use locks to protect the `refs` variable. In essence, acquiring and releasing the lock will require at least two atomic operations (e.g., using our lock implementation from before), but using atomic operations directly we can do everything in *one* atomic operation. This is therefore a good example of where we can easily see that it is beneficial to use atomic operations: 1) it is a frequent operation, so performance likely matters, 2) the critical section is short, and 3) we can trivially update the code to use available atomic operations.

10.3.2 Multi-Step Operations

It is not as easy to use atomic operations in cases that require multiple step to update the shared data. Initially, it might seem that we need to fall back on using locks in such cases. However, if we are a bit creative, we can actually use `compare_and_swap` to make arbitrary modifications to a single shared variables in an atomic manner. However, this comes with some caveats as we shall see.

To illustrate the idea, we will use the function `append_digit` in `append-digit-1.c` as an example (Listing 10.13). The function treats the value in the variable `shared` as a string of digits, and appends the digit passed as a

```
int shared;

void append_digit(int digit) {
    shared = shared * 10 + digit;
}
```

Listing 10.13: The function `append_digit` from `append-digit-1.c`.

parameter to the end of the number. For example, calling `append_digit(1)` followed by `append_digit(2)` makes `shared` contain the number 12.

This particular function is likely not very useful in real programs. However, it is simple enough to not obscure the usage of atomic operations while being complex enough to illustrate the idea. In particular, `append_digit` needs to first read `shared`, multiply it by 10, add `digit`, and then write the result back to `shared`. To do this as an atomic operation, we would need an atomic operation that performs *both* a multiplication and an addition. Sadly, we do not have such an operation available.

In general, the process of updating a single shared variable can be decomposed into three steps as follows:

1. Read the shared variable, store it as `current`.
2. Compute the new value based on `current`. Store it as `next`.
3. Store the `next` back into the shared variable.

In the case of `append_digit`, step 2 consists of two steps: a multiplication and an addition. If we examine the steps closer, we can see that only steps 1 and 3 access shared data. As such, we can eliminate data races by implementing step 1 with `atomic_read` and step 3 with `atomic_write`. Using this idea for `append_digit` makes it look like in Listing 10.14.

```
void append_digit(int digit) {
    int current = atomic_read(&shared);
    int next = current * 10 + digit;
    atomic_write(&shared, next);
}
```

Listing 10.14: Using `atomic_read` and `atomic_write` to eliminate data races. Available as `append-digit-2.c`.

However, as we saw from earlier in the chapter, simply replacing reads and writes in this manner does not produce the desired results. If two threads call `append_digit` at the same time, and both threads finish executing `atomic_read` before any of the threads has executed `atomic_write`, one of the additions to the number gets lost. This is visible in the program that uses `append_digit` in `append-digit-2.c`. The function `append_digit` is called twice in total, so we would expect `shared` to contain a 2-digit number at the end of the program (either 12 or 21 depending on which thread executed `append_digit` first). However, with the current implementation `shared` may only contain 1 digit. You can observe this behavior using `Run ⇒ Look for errors...` in Progviz.

To solve the problem, we can modify step 3 a bit. The key insight is that the implementation in `append-digit-2.c` works correctly in *some* cases. In particular, it works fine as long as no other thread updates `shared` after we have read it using `atomic_read`. This causes us to overwrite any changes the other thread made to the variable. As such, if we are able to *detect* that the shared data was modified before overwriting it, we could simply abort our attempt at updating the variable and go back to step 1 to try again. The tricky part is that we need to perform this check together with the write as an atomic operation.

Luckily, we can achieve precisely this using `compare_and_swap`. We can ask `compare_and_swap` to compare the value of `shared` to the value we read in step 1 before replacing the value with the new updated value. That way, we only update `shared` if it still contains the value we expect it to. Hence, if another thread updated the value while we were executing step 2, `compare_and_swap` will not update the value, and we can retry the update. As such, the revised steps become as follows:

1. Read the shared variable, store it as `current`.
2. Compute the new value based on `current`. Store it as `next`.
3. Check the value stored in the shared variable. If it is equal to `current` (i.e., no other thread has modified it), store `next` into it. Otherwise go to step 1.

If we update the implementation of `append_digit` according to the revised step 3, it will look like in Listing 10.15 (`append-digit-3.c`).⁹ This makes the function behave according to our expectations, which you can see by using *Run ⇒ Look for errors...* in Progviz.

```
void append_digit(int digit) {
    int current;
    int next;
    do {
        current = atomic_read(&shared);
        next = current * 10 + digit;
    } while (compare_and_swap(&shared, current, next)
            != current);
}
```

Listing 10.15: Revised implementation of `append_digit` that works properly. Available as `append-digit-3.c`.

The most interesting part of this approach is that it is general enough to work for any update to the variable, as long as we can break it down into the three steps above. Note that this includes operations that examine the current

⁹It is not necessary to re-read `shared` inside the loop. We could restructure the program to only call `atomic_read` before the loop, and then use the return value from `compare_and_swap` in subsequent iterations. However, this makes the program flow less clear, which is why this is not done in `append-digit-3.c`.

state of the shared variable and decide to *not* update it at all. That is, it is acceptable if step 2 aborts the update altogether and never executes step 3.

While the approach is powerful, it has an important drawback. Note that our modifications to step 3 introduced a loop (perhaps more clearly visible in Listing 10.15). This means that it is possible that a single thread has to re-try the update multiple times before it succeeds.

To illustrate the impacts of this, consider a program where 2 threads call `append_digit` in a loop. Assume that the scheduler happens to schedule the threads so that thread 1 always gets interrupted just after `atomic_read` has finished executing. In this situation, thread 1 will read the current value of `shared`, then it will be interrupted by thread 2 which will update `shared` with some new value. Once thread 1 resumes, it will execute `compare_and_swap` and find that the update failed since thread 2 just updated `shared`. As such, it will try again, but after executing `atomic_read` to retrieve the new value it will be interrupted by the scheduler once again, and thread 2 will update `shared` once more, so that thread 1 fails once again. As long as this continues, thread 1 will never succeed.

The example above illustrates that this approach does not provide *bounded waiting* and thereby suffers from *starvation*, just like our lock implementation. It does, however, provide *progress*, since threads only have to retry an update if some other thread succeeded. In other words, this approach guarantees that at least *one* thread will always succeed, but since we have no way of controlling *which one*, it is possible that one of the threads is unlucky and loses the race every time. To observe this in practice, however, one would need a fair number of threads that constantly call `append_digit` in a tight loop. It is, however, worth being aware of.

10.4 Performance

Our conclusion so far in the chapter is that atomic operations allow us to implement locks. However, we have seen that it is not trivial to implement a good lock, so we rather use the implementation provided by the language if possible. This leaves the second use-case: using atomic operations to avoid using locks altogether to improve performance. As such, it is worthwhile to investigate how atomic operations affect performance, so that we can better judge when it is worthwhile to use atomic operations over locks and vice versa. Of course, it is nearly impossible to accurately predict the exact runtime of two different implementations without measuring. Even so, having a rough idea helps us to pick a suitable starting point.

To compare the performance of different approaches, we will use the program `performance.c`. It starts by creating an array of 10 000 000 integers. The contents of each element is random, but the elements are created in such a way that the sum of all of them is always 10 000 000. The program then computes the sum of all elements in the array using 5 different synchronization approaches and records how much time each approach required.

To understand the overall approach used by all versions, we will start by looking at the non-synchronized version in Listing 10.16. As we can see from the figure, the implementation is fairly simple. The function simply adds elements 0 to `data_count / 2` to the global variable `result`.

```

void no_lock(int *start) {
    for (int i = 0; i < data_count / 2; i++) {
        result += start[i];
    }

    sema_up(&done);
}

```

Listing 10.16: Non-synchronized version of the summation function.

Since the function only sums *half* of the elements in the array, we need to call the function twice to sum the elements in the entire array. In an attempt to speed up the summation, we start two threads that sum two halves of the array in parallel. This means that `result` needs some form of synchronization is needed in order to avoid data races. As such, the difference between the five versions is how they avoid data races when updating the `result` variable.

The program `performance.c` is intended to be executed from the command line. You can choose to run it either with or without optimizations (use `make OPT=1 performance` to enable optimizations). When you run the program, it will first create the array, run all of the different approaches, and finally print a summary of the time for each approach. The first part of the output looks something like below:

```

Running no_lock:
    sequential time:      10.3 ms
    threaded time   :      7.6 ms
    threaded result: 5134884

```

The first line simply states which of the different approaches was used. In this case, `no_lock` was used. The first indented line, *sequential time*, shows how much time was needed to sum the array sequentially, using a single thread that calls `no_lock` twice. This is included both as a baseline so that we can see how much faster the multi-threaded implementation is, but also to make it possible to assess the cost of synchronization in the best case, when threads never have to wait.

The next indented line, *threaded time*, shows how much time was needed to sum the array using two threads. That is, one thread sums the first half of the array, and the other sums the second half of the array. In this case, we can see that it was indeed quicker to split the work across two threads. However, due to the cost of creating and waiting for the threads, the threaded time (7.6 ms) was more than half of the time of the sequential time (10.6 ms) as we would have expected.

The last line, *threaded result*, shows the value of the `result` variable after both threads finished. In this case, the result was 5 134 884, which is quite far from the expected result of 10 000 000. Of course, this is because the code contains a data race.

One way to avoid the data race and make the code produce the correct result is to use a lock to ensure that `result` is never accessed by the two threads concurrently. This gives us the code shown in Listing 10.17.

```
void lock(int *start) {
    for (int i = 0; i < data_count / 2; i++) {
        lock_acquire(&result_lock);
        result += start[i];
        lock_release(&result_lock);
    }

    sema_up(&done);
}
```

Listing 10.17: Summation function synchronized using locks.

If we look at the next block of data in the output from the program we will see output similar to below:

```
Running lock:
sequential time:    221.9 ms
threaded time  :   669.9 ms
threaded result: 10000000
```

A first observation is that the sequential time is about 20 times longer than the sequential time without any synchronization. There are two main reasons for this increase. The first one is of course the fact that we need to call `lock_acquire` and `lock_release`. Since the logic inside the critical section runs quickly, the difference is quite large. The cost is also slightly higher here than what it could have been since `lock_acquire` and `lock_release` are implemented in a different translation unit, meaning that they can not be inlined. The second reason is that the use of synchronization primitives prevents the compiler from making certain assumptions. For example, without synchronization, the compiler is allowed to read `result` *once* and then just update it as it progresses throughout the loop.¹⁰ However, after the call to `lock_acquire`, the compiler needs to assume that `result` has been changed, and must therefore read it in each iteration of the loop.¹¹

If we look at the *threaded time*, we get another surprise. In this case, adding a second thread made the summation take more than twice the time. The main reason for this is *lock contention*. Since the critical section covers the entire body of the loop, the only thing the two threads can do without holding the lock is to increment the loop counter. As such, the lock is almost always held by one of the threads, and the other one frequently needs to wait. The reason for the increase in time is thus a combination of poor opportunities for threads to run in parallel, combined with the overhead of pausing and resuming the two threads occasionally.

¹⁰In this particular case, it can not simply write `result` once at the end due to aliasing. But this is outside the scope of this book.

¹¹This is actually the case for all function calls that the compiler cannot see the implementation of. Since the compiler does not know what they do, it cannot rule out the possibility that the function modifies `result`, for example.

```
void atomics(int *start) {
    for (int i = 0; i < data_count / 2; i++) {
        atomic_add(&result, start[i]);
    }

    sema_up(&done);
}
```

Listing 10.18: Summation function synchronized using atomic operations.

Since locks did not work very well in this situation, let's try using atomic operations instead. In this case, we can easily replace += with a call to `atomic_add`, which gives us the code in Listing 10.18. If we look at the next block of output from the program, we see something similar to below:

```
Running atomics:
sequential time:    55.7 ms
threaded time  :   213.9 ms
threaded result: 10000000
```

From the *sequential time* we can see that using atomic operations is more expensive than no synchronization. However, the total runtime was much shorter than when using locks. This time it only took 5 times longer than the non-synchronized version (i.e., 55.7 ms vs. 10.3 ms). This is not a surprising result. We expect using atomic operations should incur *some* cost. After all, similar to locks they disallow some optimizations, and the hardware needs to do extra work to ensure that the operation becomes atomic.

Looking at the *threaded time* does however reveal something surprising. We would have expected that using two threads would be faster than using one thread. However, the total runtime is almost four times the runtime of using a single thread, even though the threads do not need to wait for each other. This result reveals another interesting property about atomic operations. As we saw for the sequential case, atomic operations are typically quite fast when only a single thread accesses the same data. However, if multiple threads access the same data using atomics, they are often more expensive. The main reason for this is that the hardware needs to keep caches that are not shared between CPU cores synchronized (e.g., the L1 cache is often not shared). For example, when thread 1 executes `atomic_add`, the hardware might discover that thread 2 recently updated `shared` but has not yet updated the value in memory. This means that thread 1 needs to either wait for thread 2 to store the data to memory (or some shared cache), or fetch it directly from thread 2's cache. Both of these operations take additional time, which is one reason why using two threads is much more expensive than using one thread.

So far, all of our attempts at improving performance have failed. Even though we have seen that atomic operations are indeed faster than locks, both were surprisingly slower when we added more threads. How do we actually make our program faster?

The key insight from both cases is that managing access to shared data is expensive, regardless of how we choose to do it. As such, to actually get a speedup we wish to avoid synchronization as much as we can. This is actually

a good advice in general. As we have seen in this book so far, synchronizing access to shared data is both difficult to get right and expensive. As such, we both simplify our work as a programmer and make our programs perform better if we structure our programs to avoid shared data as much as possible.

```
void lock_separate(int *start) {
    int local = 0;
    for (int i = 0; i < data_count / 2; i++) {
        local += start[i];
    }

    lock_acquire(&result_lock);
    result += local;
    lock_release(&result_lock);
    sema_up(&done);
}
```

Listing 10.19: Summation function synchronized using locks, but using a separate variable.

```
void atomics_separate(int *start) {
    int local = 0;
    for (int i = 0; i < data_count / 2; i++) {
        local += start[i];
    }

    atomic_add(&result, local);
    sema_up(&done);
}
```

Listing 10.20: Summation function synchronized using atomics, but using a separate variable.

In the case of our summation function, this is quite easy to do. We can simply let each thread keep track of the sum of their half of the array in a local variable, and add the local variable to the global one once at the end of the function. This works since it does not matter in which order we add the different elements of the array, and since we are only interested in the final result of the summation.

We can implement this as shown in Listings 10.19 and 10.20. The difference between the two implementations is that one uses a lock to make sure that there is no data race when updating `result`, while the other uses atomic operations instead.

If we look at the final part of the output, we see something like below:

```
Running lock_separate:
  sequential time:    3.5 ms
  threaded time  :    2.3 ms
  threaded result: 10000000
Running atomics_separate:
  sequential time:    6.5 ms
  threaded time  :    2.5 ms
  threaded result: 10000000
```

From the output we can see that this implementation is actually faster when we use two threads compared to when we only use one thread. The main reason for this is that the two threads can work independently for most of their runtime. As we saw, synchronization is expensive, mostly since threads need to wait for each other. Furthermore, we can see that the sequential time is actually lower than the sequential time of the non-synchronized version. One reason for this is that we use a local variable to store the partial sum rather than a global variable. Since the scope of local variables are limited, the compiler know more about how they are used and can more easily place them in registers to improve performance.

We can see another reason for this discrepancy if we compare the sequential time for `lock_separate` and `atomics_separate`. Since the two functions are almost the same, we would expect their runtime to be almost the same as well. However, in this case `atomics_separate` takes almost twice as long compared to `lock_separate`. This shows the problems with this kind of micro-benchmarks. There are many factors that impact the runtime of a program. For example, what parts of the array are in the cache, alignment of the machine code, among others. All of this is further influenced by what other processes in the system are doing at the time. As such, while a difference between 10, 50, and 500 ms is no coincidence, a difference between 3 and 6 ms might very well be. To get a more definitive answer, we would need to either run the test more times, or use a much larger array as input.

We will not do that here, since the goal if this section is to build an *intuition* about the cost of the different options. The results here are accurate enough for that purpose. We have seen that atomic operations are indeed cheaper than locks. However, we also found that synchronization is always fairly expensive, so the key takeaway is that the best course of action is to avoid synchronization as much as possible. If we manage to minimize the places where synchronization is needed (which is not always possible), it becomes less important exactly which approach we choose.

10.5 Advanced Usage

To conclude the chapter, we will take a brief look at some more nuances of atomic operations. To understand these nuances, we will also need to examine some nuances in the memory model that we have ignored so far. It is worth noting that even though the memory model that has been presented so far in this book is a simplification of the memory model used by C, it is still correct in the context of what has been presented in the book, and it is sufficient for

the majority of concurrent programming. However, if you are interested in continuing to study concurrent programming, it is worth to be aware of which simplifications we have done, so that you know where to look to learn more.

10.5.1 Memory Ordering

To understand the nuances of atomic operations, we must first have a closer look at different *memory orderings*. A good starting point is the page about *memory order* at cppreference.com.¹² It is worth noting that even though the page describes C++, the parts about memory ordering in concurrency is almost identical in C.

As you hopefully remember from the start of the book, we discovered that a major reason why we need to avoid data races is that both the compiler and the hardware may reorder the memory accesses made by our program to make it execute more efficiently. While care is taken to ensure that we can not observe these changes in a single-threaded program, they may be visible in a concurrent program. As such, by default we have no guarantees as to when and in what order stores from one thread become visible to another thread.

By using synchronization primitives or atomic operations (we call them *synchronization events* for simplicity), we can tell the compiler that we need stronger guarantees in some places since we are trying to coordinate work between multiple threads. The descriptions in the book so far have implied that these synchronization operations provide what is called *sequential consistency*. This means two things. First, it means that neither the compiler nor the hardware is allowed to move loads or stores across the synchronization event. That is, stores that happen before the synchronization event in the source code may not be delayed so that they occur after it. Similarly, loads that occur after it in the source code may not occur before the synchronization event. Secondly, it means that all threads in the program observe that all synchronization events occur in the same order.

This is very close to Progvis' visualizations. Even though we control in which order different threads execute, we only have one representation of the “one and true” state of memory that is observed by all threads. In particular, this means that all threads observe all synchronization events in the same order. As a sidenote, the button “take larger steps” in Progvis steps the program to the next synchronization event (but also to the next loop iteration or return from a function for additional clarity).

Synchronization events may have weaker memory orderings as well. Some examples are *acquire*, *release*, and *acquire-release*. These affect how much liberty the compiler has in reordering memory operations.

A synchronization event with *acquire* semantics disallows the compiler and the hardware from moving loads that occur after the event to before the event. As such, one way to think of it is that an acquire event marks the earliest possible time where subsequent loads may occur.

Similarly, a synchronization event with *release* semantics disallows the compiler and the hardware from moving stores that occur before the event to after the event. As such, it marks the point where all stores must be complete.

¹²https://en.cppreference.com/w/cpp/atomic/memory_order.html

Finally, *acquire-release* simply means both acquire and release. That is, neither loads nor stores may be moved across the synchronization event. This is equivalent to two synchronization events after one another, first an *acquire* event and then a *release* event. It is worth noting that while an acquire-release event is similar to a sequentially consistent event, there is a slight difference: namely that if two threads perform two acquire-release events on different variables, these two events do not necessarily have a global order. That is, other threads in the system may observe them to have happened in a different order. This nuance is rarely important unless you are working with low-level atomic operations, but it is worth noting that there *is* a difference between the two.

The names *acquire* and *release* may seem arbitrary at first. However, they are chosen to match the semantics required by `lock_acquire` and `lock_release`. That is, `lock_acquire` typically has *acquire* semantics, since we only really need to disallow that accesses to variables that are protected by the lock are loaded before the lock is actually acquired. It is fine if a store to a variable that should happen before we acquired the lock actually occurs after acquiring the lock. Similarly, `lock_release` typically has *release* semantics, since it is important that any stores to the variables protected by the lock are completed before we release the lock, but it is fine if we start loading some variable that is needed after the lock is released a bit too early.

This is one example of a simplification made in the book: we implicitly considered both `lock_acquire` and `lock_release` to be sequentially consistent even though their actual semantics are a bit weaker. However, as long as we follow the rule that we protect shared data with the *same* lock everywhere, this difference does not matter. We can reason correctly about our program regardless.

10.5.2 Relaxed Atomics

This brings us to the relation to atomic operations. Atomic operations constitute a synchronization event, which means that they have a memory ordering. By default, the ordering is strong enough to be equivalent to sequential consistency (most operations are sequentially consistent, except the ones that only read or only write). This means that they act as described in this chapter by default.

However, there is often some way to weaken the memory order to gain some performance. For example, in C++ it is possible to specify a memory order as a (compile-time constant) parameter to the atomic operations. That way we can tell the compiler that we only need *acquire* semantics on the `compare_and_swap` in our `my_lock_acquire` function for example. We can even tell the compiler that we do not care about the ordering in relation to other memory accesses, only that it is atomic. This is called a *relaxed atomic operation* in C++ and in other places. One example where this might be useful is in our summation example. Here, it is not important that the `atomic_add` operations occur in some global order, only that all of them are eventually completed (which might be problematic since we need to know that they are all done at some point). As you might imagine, it is quite tricky to get this right. Particularly since even an incorrect program may appear to work correctly on some systems (e.g., because the particular CPU does not provide the exact requested semantics, so the compiler used a slightly stronger semantics). Since

the performance gains are often rather small, this kind of atomic operations are rarely used outside of high-performance concurrent data structures.

10.6 Exercises

1. The file `future.c` contains the implementation of `struct future` from Chapter 9. The semaphore is removed and replaced with a boolean variable.
 - a) Use relevant atomic operations to make the implementation work as expected. You can use Progvis to verify your solution.
 - b) In your solution to a), you needed to use `intentional_busy_wait` to make Progvis ignore the busy-wait loop in the solution. Use one of the available synchronization primitives (i.e., semaphores, locks, and/or condition variables) to make your implementation wait efficiently. The goal is to utilize atomic operations to avoid the situation where threads need to wait for each other in `future_get` even though `future_post` has been called.

You can use Progvis to verify that your solution is correct. However, Progvis will not check that the threads never have to wait for each other.
2. The file `semaphore.c` contains an incomplete implementation of a semaphore. Use only atomic operations to finish the implementation. You can use Progvis to verify your solution.

Threading Library Reference

This chapter contains reference to all threading-related functions introduced in this book. All of the functions are introduced in more detail where they are first used. The goal of this chapter is therefore to list all functions in one place to make them easier to find.

The functionality described here is available both in Progis and in C through the threading library provided with this book. The names and semantics of the synchronization primitives are inspired from their appearance in Pintos.

A.1 Thread Management

```
#include "thread.h"
```

These functions and types are defined in the file `thread.h`.

```
tid_t
```

A type that is used to store thread identifiers. A thread identifier may be any type, so it is only safe to assume that it is possible to compare two `tid_t` values with the `==` and `!=` operators. For the purposes of this book you can mostly ignore this type. The book mostly uses `thread_new` below to start threads. It is seldom necessary to keep track of threads by their identifier.

```
tid_t thread_new(thread_func *fn, void *data, ...);
```

Creates a new thread. The new thread will execute the function passed as `fn`. Note that it is not well-defined *when* the new thread will start execute `fn` in relation to when the call to `thread_new` returns. The function returns the thread-id of the newly created thread.

For convenience, this function is implemented specially to allow a variable number of parameters to be passed to the function in the thread. As such, you will not find the above definition in the header files in the threading library.

The function can be used as follows:

```
void zero(void);
void one(int *x);
void two(int *x, struct my_data *y);

int main(void) {
    int a = /* ... */;
    struct my_data b = /* ... */;

    thread_new(&zero);
    thread_new(&one, &a);
    thread_new(&two, &a, &b);

    return 0;
}
```

Note that `thread_new` accepts a variable number of parameters. However, the implementation requires that all parameters are pointers. As such, passing `b` by value to a function would fail.

```
tid_t thread_current(void);
```

Retrieve the identifier of the current thread.

```
void thread_yield(void);
```

Asks the scheduler to let other threads run for a while. This is just a hint, and it is entirely up to the scheduler to determine which other threads to run and for how long.

```
void timer_msleep(int ms);
```

Pause execution of the current thread for *at least* the specified number of milliseconds. Note that it is acceptable for the implementation to wait for *more* than the specified time.

```
void prevent_optimization(void);
```

A utility function that does nothing but preventing the compiler from performing certain kinds of optimizations. If you look at the function you will see that it is empty. The reason it works is that it is implemented in another *translation unit*, which means that the compiler cannot see what the implementation actually does, and therefore has to make pessimistic assumptions.¹ This is not a general way of preventing optimizations that break programs. It is only used in this book to force the compiler to take certain decisions and illustrate certain problems. It is not a replacement for proper synchronization.

¹Unless link-time optimizations are used.

Progviz-specific Notes:

- `thread_new` is implemented as a special form. It is thereby possible to pass an arbitrary number of parameters to functions in Progviz. Progviz is also not restricted to pointer types, but can pass arbitrary data types. This also means that the calls are properly type-checked. However, be aware of the limitations in the C version noted above if you aim to write portable code!
- `tid_t` is implemented as a special type. It is only possible to compare them to other `tid_t` variables.
- `thread_yield` and `timer_msleep` have no effect at all. These functions are used to make certain problematic interleavings more likely to occur when running programs outside of Progviz. Since Progviz allows single stepping anyway, they provide no benefit inside Progviz.

C-specific Notes:

- `thread_new` supports up to 5 parameters. Passing more than 5 parameters will fail with an error message. Also note that parameters are not type-checked since all parameters are cast to `void *`. Make sure that the parameter count and types passed to `thread_new` match the parameter count and types accepted by the function!

The C version also contains the following functions that might be useful to force interesting thread behaviors:

```
void thread_exit(void);
```

Immediately terminates the current thread. This is not used in the book, as we can simply return from the entry-point of the thread instead.

A.2 Synchronization Primitives

```
#include "synch.h"
```

These functions and types are defined in the file `synch.h`.

```
struct semaphore;
```

Datatype that represents a semaphore. The datatype is fully defined in the header file, so it can be stack-allocated and used inside other data structures. Even though the data structure is available in the header file, it should be considered as a “black box”, and it should only be accessed through the functions below. Note that it is generally not safe to copy a semaphore once it has been initialized, so make sure to pass the semaphore or the data structure that contains it by pointer.

```
void sema_init(struct semaphore *s, int value);
```

Initialize the semaphore `s` with a counter of `value`. Assumes that `value` ≥ 0 .

```
void sema_down(struct semaphore *s);
```

Try to decrement the counter in the semaphore `s`. If the counter would become negative, the calling thread waits until another thread calls `sema_up` and decrements the counter when it wakes up.

```
void sema_up(struct semaphore *s);
```

Increments the counter inside the semaphore `s`. This may wake one thread that is waiting inside a call to `sema_down`. Note that it is not well-defined *which* thread wakes up if multiple threads are waiting.

```
void sema_destroy(struct semaphore *s);
```

Deallocates any resources allocated by `sema_init`. This is not crucial to proper operation of the program, but leads to resource leaks on some platforms. Progviz does not consider it essential to deallocate semaphores, so this book considers it to be optional to destroy semaphores.

```
struct lock;
```

Datatype that represents a lock. The datatype is fully defined in the header file, so it can be stack-allocated and used inside other data structures. The datatype is fully defined in the header file, so it can be stack-allocated and used inside other data structures. Even though the data structure is available in the header file, it should be considered as a “black box”, and it should only be accessed through the functions below. Note that it is generally not safe to copy a semaphore once it has been initialized, so make sure to pass the semaphore or the data structure that contains it by pointer.

```
void lock_init(struct lock *l);
```

Initialize the lock `l`.

```
void lock_acquire(struct lock *l);
```

Try to acquire the lock `l`. If another thread currently holds the lock, the calling thread will wait until the lock is released.

```
void lock_release(struct lock *l);
```

Release the lock `l`. Assumes that the current thread currently holds the lock (i.e., the current thread has previously called `lock_acquire` for the same lock). If one or more threads are waiting to acquire the lock, this will cause one of them to wake up eventually. Note that it is not well-defined *which* thread wakes up.

```
void lock_destroy(struct lock *l);
```

Deallocates any resources allocated by `lock_init`. This is not crucial to proper operation of the program, but leads to resource leaks on some platforms. Progviz does not consider it essential to deallocate semaphores, so this book considers it to be optional to destroy semaphores.

```
struct condition;
```

Datastructure that represents a condition variable. The datatype is fully defined in the header file, so it can be stack-allocated and used inside other data structures. Even though the data structure is available in the header file, it should be considered as a “black box”, and it should only be accessed through the functions below. Note that it is generally not safe to copy a semaphore once it has been initialized, so make sure to pass the semaphore or the data structure that contains it by pointer.

This implementation of condition variables is *not* thread-safe in and of itself (in contrast to e.g., pthreads). As such, each condition variable needs to be protected with *one* lock. As such, the implementation expects that this lock is held when `cond_wait`, `cond_signal`, or `cond_broadcast` is called, and that the lock is passed as the parameter `l`.

```
void cond_init(struct condition *c);
```

Initialize the condition variable `c`.

```
void cond_wait(struct condition *c, struct lock *l);
```

Releases the lock `l` and causes the current thread to wait until another thread calls `cond_signal` or `cond_broadcast`. When the thread resumes, it acquires `l` before returning from `cond_wait`.

```
void cond_signal(struct condition *c, struct lock *l);
```

Causes *one* of the threads that are currently waiting inside a call to `cond_wait` for the condition variable `c` to wake up. Nothing happens if no threads are currently waiting for `c`. Note that it is not well-defined *which* thread wakes up. Assumes that `l` is held.

```
void cond_broadcast(struct condition *c, struct lock *l);
```

Causes *all* threads that are currently waiting inside a call to `cond_wait` for the condition variable `c` to wake up. Nothing happens if no threads are currently waiting for `c`. Assumes that `l` is held.

```
void cond_destroy(struct condition *c);
```

Deallocates any resources allocated by `cond_init`. This is not crucial to proper operation of the program, but leads to resource leaks on some platforms. Progviz does not consider it essential to deallocate semaphores, so this book considers it to be optional to destroy semaphores.

A.3 Atomic Operations

The atomic operations below are generic in the sense that they are able to operate on any integer-like type (i.e., integers and pointers), which is denoted as `T` below. The functions are listed alongside with their semantics, expressed as a regular function. Do remember, however, that the actual implementation uses *compiler intrinsics* to let the compiler and hardware know that the operations should be atomic.

```
#include "atomics.h"
```

These functions and types are defined in the file `atomics.h`.

```
T atomic_read(T *value) {
    return *value;
}
```

Read the value pointed to by `value` atomically.

```
void atomic_write(T *value, T write) {
    *value = write;
}
```

Write the value passed to `write` to the location pointed to by `value` atomically.

```
T test_and_set(T *value) {
    T old = *value;
    *value = 1;
    return old;
}
```

Read the value pointed to by `value` and set it to 1. Return the old value pointed to by `value`. Note that `test_and_set` is unique in that it is not possible to use for pointer variables (making a pointer refer to the address 1 is rarely useful).

```
T atomic_swap(T *value, T replace) {
    T old = *value;
    *value = replace;
    return old;
}
```

Read the value pointed to by `value` and swap it the value passed as `swap`. Return the old value pointed to by `value`. Note that only the contents of `value` is accessed atomically.

```
T compare_and_swap(T *value, T compare, T replace) {
    T old = *value;
    if (old == compare)
        *value = replace;
    return old;
}
```

Read the value pointed to by `value`. Compare it to the value passed as `compare`. If they were equal, replace the value pointed to by `value` with the value passed as `swap`. Regardless of the check, return the old value pointed to by `value`. Note that only the contents of `value` is accessed atomically.

```
T atomic_add(T *value, T add) {
    T old = *value;
    *value += add;
    return old;
}
```

Read the value pointed to by `value`, add one to it, and return the original value.

```
T atomic_sub(T *value, T sub) {
    T old = *value;
    *value -= sub;
    return old;
}
```

Read the value pointed to by `value`, subtract one from it, and return the original value.

```
void intentional_busy_wait();
```

Function used to mark loops that use atomic operations to wait for something. By default, the option *Run* \Rightarrow *Look for errors...* for places in the code where busy-wait may occur. Since there is no way to wait “properly” using only atomic operations, we need to tell Progviz that we are aware that some particular loop contains a busy wait, and that we do not think it is a problem. By calling `intentional_busy_wait` inside such loops we inform Progviz about this, and thereby it will ignore that particular busy wait.

Limitations in Progvis

Progvis uses its internal C compiler to compile programs. This means that it does not need a separate C compiler to be installed on your system for Progvis to work. You can just download the archive and unpack it somewhere, and Progvis will work. However, there are a number of differences between the C compiler used by Progvis and the ones you might be familiar with.

There are many reasons for this. Some differences are to make it more convenient to load programs into Progvis. Some are to make the language more memory-safe, both so that Progvis can visualize the memory properly, but also to make it less likely that the program crashes Progvis.

Below is a list of notable differences between Progvis' C compiler and most other C compilers, with short explanations of why.

- No support for floating-point types.

Progvis does not currently support the types `float` or `double`. This is not a technical limitation, they have simply not been needed so far.

- No support for type casts.

The main reason for disallowing type casts is to disallow casting between different pointer types (e.g., `(int *)`), which undermines much of the type safety. Casting between pointers and integers is not easily possible, both to retain some memory and type safety, but also since Progvis internally represents pointers as a pointer to the start of the allocation and an offset into the allocation.

Pointers between different integer types could be supported, but has not been needed so far.

- No support for `void *`.

Without pointer casts (neither explicit nor implicit), `void *` is not very useful. As such, it is not available.

Because of this, the threading library has been designed to not need `void *`. For example, `thread_new` is implemented as a language construct in Progvis to make it type safe.

- `#include` statements are ignored.

Progvis does not follow the compilation model of C, and does not attempt to parse any `#include` statements. Instead, it contains an implementation of the threading library (Chapter A) and a minimal C standard library that is always available to programs.

Note that this also means that it is not possible to include custom `.h` files. It is, however, possible to load multiple `.c` files as one program. This is done by selecting multiple files in the *File* → *Open...* dialog in Progvis (usually Ctrl+click). In contrast to C, the `.c` files are treated as a single translation unit, so no header files are necessary.

- It is necessary to `return` in `main`.

The C (and C++) standard both allows and defines what happens if `main` does not `return`. This special case is not implemented in Progvis, so you need to add `return 0;` at the end of your program. On the other hand, Progvis also allows `main` to return `void`, but then the program is not possible to compile with a normal C compiler.

- It is not possible to accept parameters to `main`.

The `main` function must not accept command line parameters (typically named `argc` and `argv`).

- No initializers for `structs`.

While Progvis supports initializing arrays using the initializer syntax (i.e., `{0, 1, ...}`), it is not supported for `structs` yet. Similarly, default values for variables in `structs` are not supported.