



Rust

For C++

Programmers



Published
with GitBook



Table of Contents

Introduction	0
Hello world!	1
Intermission - why Rust?	2
Control flow	3
Primitive types and operators	4
Unique pointers	5
Borrowed pointers	6
Rc and raw pointers	7
Data types	8
Destructuring pt 1	9
Destructuring pt 2	10
Arrays and vecs	11
Graphs and arena allocation	12

Rust for C++ Programmers

This gitbook is a collection of "Rust for C++ programmers" posts by Nick Cameron.

I found it a bit hard to read the posts on [his blog](#), hence I started copy-pasting them into a [git repo](#) to make it easier to read. Of course, all the credit for content goes to [Nick Cameron](#).

After a while, I decided to turn them into a gitbook, so here it is!

Update (June 2015): I've realized that Nick has created a repo for his posts on GitHub ([nrc/r4cxxx](#)); I've merged new content and will look into possibility of merging this repo with Nick's.

Introduction - hello world!

This is the first in a series of blog posts (none written yet) which aim to help experienced C++ programmers learn Rust. Expect updates to be sporadic at best. In this first blog post we'll just get setup and do a few super basic things. Much better resources are at the tutorial and reference manual.

First you need to install Rust. You can download a nightly build from <http://www.rust-lang.org/install.html> (I recommend the nightlies rather than 'stable' versions - the nightlies are stable in that they won't crash too much (no more than the stable versions) and you're going to have to get used to Rust evolving under you sooner or later anyway). Assuming you manage to install things properly, you should then have a `rustc` command available to you. Test it with `rustc -v`.

Now for our first program. Create a file, copy and paste the following into it and save it as `hello.rs` or something equally imaginative.

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

Compile this using `rustc hello.rs`, and then run `./hello`. It should display the expected greeting `\o/`

Two compiler options you should know are `-o ex_name` to specify the name of the executable and `-g` to output debug info; you can then debug as expected using `gdb` or `lldb`, etc. Use `-h` to show other options.

OK, back to the code. A few interesting points - we use `fn` to define a function or method. `main()` is the default entry point for our programs (we'll leave program args for later). There are no separate declarations or header files as with C++. `println!` is Rust's equivalent of `printf`. The `!` means that it is a macro, for now you can just treat it like a regular function. A subset of the standard library is available without needing to be explicitly imported/included (we'll talk about that later). The `println!` macros is included as part of that subset.

Lets change our example a little bit:

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

`let` is used to introduce a variable, `world` is the variable name and it is a string (technically the type is `&'static str`, but more on that in a later post). We don't need to specify the type, it will be inferred for us.

Using `{}` in the `println!` statement is like using `%s` in `printf`. In fact, it is a bit more general than that because Rust will try to convert the variable to a string if it is not one already¹. You can easily play around with this sort of thing - try multiple strings and using numbers (integer and float literals will work).

If you like, you can explicitly give the type of `world`:

```
let world: &'static str = "world";
```

In C++ we write `T x` to declare a variable `x` with type `T`. In Rust we write `x: T`, whether in `let` statements or function signatures, etc. Mostly we omit explicit types in `let` statements, but they are required for function arguments. Lets add another function to see it work:

```
fn foo(_x: &'static str) -> &'static str {
    "world"
}

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

The function `foo` has a single argument `_x` which is a string literal (we pass it "bar" from `main`). We don't actually use that argument in `foo`. Usually, Rust will warn us about this. By prefixing the argument name with `_` we avoid these warnings. In fact, we don't need to name the argument at all, we could just use `_`.

The return type for a function is given after `->`. If the function doesn't return anything (a void function in C++), we don't need to give a return type at all (as in `main`). If you want to be super-explicit, you can write `-> ()`, `()` is the void type in Rust. `foo` returns a string literal.

You don't need the `return` keyword in Rust, if the last expression in a function body (or any other body, we'll see more of this later) is not finished with a semicolon, then it is the return value. So `foo` will always return "world". The `return` keyword still exists so we can do early returns. You can replace `"world"` with `return "world";` and it will have the same effect.

¹

This is a programmer specified conversion which uses the `Display` trait, which works a bit like `toString` in Java. You can also use `{:?}` which gives a compiler generated representation which is sometimes useful for debugging. As with `printf`, there are many other options.

TODO check this footnote hack works

I realise that in terms of learning Rust, I had jumped straight to the 'how' and skipped the 'why'. I guess I am in enough of a Rust bubble that I can't imagine why you wouldn't want to learn it. So, I will make a bit more of an effort to explain why things are how they are. Here I will try to give a bit of an overview/motivation.

If you are using C or C++, it is probably because you have to - either you need low-level access to the system, or need every last drop of performance, or both. Rust aims to do offer the same level of abstraction around memory, the same performance, but be safer and make you more productive.

Concretely, there are many languages out there that you might prefer to use to C++: Java, Scala, Haskell, Python, and so forth, but you can't because either the level of abstraction is too high (you don't get direct access to memory, you are forced to use garbage collection, etc.), or there are performance issues (either performance is unpredictable or it's simply not fast enough). Rust does not force you to use garbage collection, and as in C++, you get raw pointers to memory to play with. Rust subscribes to the 'pay for what you use' philosophy of C++. If you don't use a feature, then you don't pay any performance overhead for its existence. Furthermore, all language features in Rust have predictable (and usually small) cost.

Whilst these constraints make Rust a (rare) viable alternative to C++, Rust also has benefits: it is memory safe - Rust's type system ensures that you don't get the kind of memory errors which are common in C++ - memory leaks, accessing un-initialised memory, dangling pointers - all are impossible in Rust. Furthermore, whenever other constraints allow, Rust strives to prevent other safety issues too - for example, all array indexing is bounds checked (of course, if you want to avoid the cost, you can (at the expense of safety) - Rust allows you to do this in unsafe blocks, along with many other unsafe things. Crucially, Rust ensures that unsafety in unsafe blocks stays in unsafe blocks and can't affect the rest of your program). Finally, Rust takes many concepts from modern programming languages and introduces them to the systems language space. Hopefully, that makes programming in Rust more productive, efficient, and enjoyable.

I would like to motivate some of the language features from part 1. Local type inference is convenient and useful without sacrificing safety or performance (it's even in modern versions of C++ now). A minor convenience is that language items are consistently denoted by keyword (`fn` , `let` , etc.), this makes scanning by eye or by tools easier, in general the syntax of Rust is simpler and more consistent than C++. The `println!` macro is safer than `printf` - the number of arguments is statically checked against the number of 'holes' in the string and the arguments are type checked. This means you can't make the `printf` mistakes of printing memory as if it had a different type or addressing memory further down the stack by mistake. These are fairly minor things, but I hope they illustrate the philosophy behind the design of Rust.

Control flow

If

The `if` statement is pretty much the same in Rust as C++. One difference is that the braces are mandatory, but brackets around the expression being tested are not. Another is that `if` is an expression, so you can use it the same way as the ternary `?` operator in C++ (remember from last time that if the last expression in a block is not terminated by a semi-colon, then it becomes the value of the block). There is no ternary `?` in Rust. So, the following two functions do the same thing:

```
fn foo(x: i32) -> &'static str {
    let mut result: &'static str;
    if x < 10 {
        result = "less than 10";
    } else {
        result = "10 or more";
    }
    return result;
}

fn bar(x: i32) -> &'static str {
    if x < 10 {
        "less than 10"
    } else {
        "10 or more"
    }
}
```

The first is a fairly literal translation of what you might write in C++. The second is better Rust style.

You can also write `let x = if ...`, etc.

Loops

Rust has while loops, again just like C++:


```
fn main() {
    let mut x = 10;
    while x > 0 {
        println!("Current value: {}", x);
        x -= 1;
    }
}
```

There is no do...while loop in Rust, but we do have the `loop` statement which just loops forever:

```
fn main() {
    loop {
        println!("Just looping");
    }
}
```

Rust has `break` and `continue` just like C++.

For loops

Rust also has `for` loops, but these are a bit different. Lets say you have a vector of integers and you want to print them all (we'll cover vectors/arrays, iterators, and generics in more detail in the future. For now, know that a `Vec<T>` is a sequence of `T`s and `iter()` returns an iterator from anything you might reasonably want to iterate over). A simple `for` loop would look like:

```
fn print_all(all: Vec<i32>) {
    for a in all.iter() {
        println!("{}", a);
    }
}
```

TODO also `&all/all` instead of `all.iter()`

If we want to index over the indices of `all` (a bit more like a standard C++ for loop over an array), you could do

```
fn print_all(all: Vec<i32>) {
    for i in ..all.len() {
        println!("{}", i, all.get(i));
    }
}
```

Hopefully, it is obvious what the `len` function does. TODO range notation

Switch/Match

Rust has a match expression which is similar to a C++ switch statement, but much more powerful. This simple version should look pretty familiar:

```
fn print_some(x: i32) {
    match x {
        0 => println!("x is zero"),
        1 => println!("x is one"),
        10 => println!("x is ten"),
        y => println!("x is something else {}", y),
    }
}
```

There are some syntactic differences - we use `=>` to go from the matched value to the expression to execute, and the match arms are separated by `,` (that last `,` is optional). There are also some semantic differences which are not so obvious: the matched patterns must be exhaustive, that is all possible values of the matched expression (`x` in the above example) must be covered. Try removing the `y => ...` line and see what happens; that is because we only have matches for 0, 1, and 10, but there are obviously lots of other integers which don't get matched. In that last arm, `y` is bound to the value being matched (`x` in this case). We could also write:

```
fn print_some(x: i32) {
    match x {
        x => println!("x is something else {}", x)
    }
}
```

Here the `x` in the match arm introduces a new variable which hides the argument `x`, just like declaring a variable in an inner scope.

If we don't want to name the variable, we can use `_` for an unnamed variable, which is like having a wildcard match. If we don't want to do anything, we can provide an empty branch:

```
fn print_some(x: i32) {
    match x {
        0 => println!("x is zero"),
        1 => println!("x is one"),
        10 => println!("x is ten"),
        _ => {}
    }
}
```

Another semantic difference is that there is no fall through from one arm to the next.

We'll see in later posts that `match` is extremely powerful. For now I want to introduce just a couple more features - the 'or' operator for values and `if` clauses on arms. Hopefully an example is self-explanatory:

```
fn print_some_more(x: i32) {
    match x {
        0 | 1 | 10 => println!("x is one of zero, one, or ten"),
        y if y < 20 => println!("x is less than 20, but not zero, one, or ten"),
        y if y == 200 => println!("x is 200 (but this is not very stylish)"),
        _ => {}
    }
}
```

Just like `if` expressions, `match` statements are actually expressions so we could re-write the last example as:

```
fn print_some_more(x: i32) {
    let msg = match x {
        0 | 1 | 10 => "one of zero, one, or ten",
        y if y < 20 => "less than 20, but not zero, one, or ten",
        y if y == 200 => "200 (but this is not very stylish)",
        _ => "something else"
    };

    println!("x is {}", msg);
}
```

Note the semi-colon after the closing brace, that is because the `let` statement is a statement and must take the form `let msg = ...;`. We fill the rhs with a `match` expression (which doesn't usually need a semi-colon), but the `let` statement does. This catches me out all the time.

Motivation: Rust `match` statements avoid the common bugs with C++ `switch` statements - you can't forget a `break` and unintentionally fall through; if you add a case to an enum (more later on) the compiler will make sure it is covered by your `match` statement.

Method call

Finally, just a quick note that methods exist in Rust, similarly to C++. They are always called via the `.` operator (no `->`, more on this in another post). We saw a few examples above (`len`, `iter`). We'll go into more detail in the future about how they are defined and called. Most assumptions you might make from C++ or Java are probably correct.

Primitive types and operators

TODO int/uint -> isize/usize

Rust has pretty much the same arithmetic and logical operators as C++. `bool` is the same in both languages (as are the `true` and `false` literals). Rust has similar concepts of integers, unsigned integers, and floats. However the syntax is a bit different. Rust uses `int` to mean an integer and `uint` to mean an unsigned integer. These types are pointer sized. E.g., on a 32 bit system, `uint` means a 32 bit unsigned integer. Rust also has explicitly sized types which are `u` or `i` followed by 8, 16, 32, or 64. So, for example, `u8` is an 8 bit unsigned integer and `i32` is a 32 bit signed integer. For floats, Rust has `f32` and `f64`.

Numeric literals can take suffixes to indicate their type (using `i` and `u` instead of `int` and `uint`). If no suffix is given, Rust tries to infer the type. If it can't infer, it uses `int` or `f64` (if there is a decimal point). Examples:

```
fn main() {
    let x: bool = true;
    let x = 34; // type int
    let x = 34u; // type uint
    let x: u8 = 34u8;
    let x = 34i64;
    let x = 34f32;
}
```

As a side note, Rust lets you redefine variables so the above code is legal - each `let` statement creates a new variable `x` and hides the previous one. This is more useful than you might expect due to variables being immutable by default.

Numeric literals can be given as binary, octal, and hexadecimal, as well as decimal. Use the `0b`, `0o`, and `0x` prefixes, respectively. You can use an underscore anywhere in a numeric literal and it will be ignored. E.g,

```
fn main() {
    let x = 12;
    let x = 0b1100;
    let x = 0o14;
    let x = 0xe;
    let y = 0b_1100_0011_1011_0001;
}
```

Rust has chars and strings, but since they are Unicode, they are a bit different from C++. I'm going to postpone talking about them until after I've introduced pointers, references, and vectors (arrays).

Rust does not implicitly coerce numeric types. In general, Rust has much less implicit coercion and subtyping than C++. Rust uses the `as` keyword for explicit coercions and casting. Any numeric value can be cast to another numeric type. `as` cannot be used to convert between booleans and numeric types. E.g.,

```
fn main() {
    let x = 34u as int;      // cast unsigned int to int
    let x = 10 as f32;      // int to float
    let x = 10.45f64 as i8; // float to int (loses precision)
    let x = 4u8 as u64;     // gains precision
    let x = 400u16 as u8;   // 144, loses precision (and thus changes the value)
    println!("`400u16 as u8` gives {}", x);
    let x = -3i8 as u8;     // 253, signed to unsigned (changes sign)
    println!("`-3i8 as u8` gives {}", x);
    //let x = 45u as bool;  // FAILS!
}
```

Rust has the following numeric operators: `+`, `-`, `*`, `/`, `%`; bitwise operators: `|`, `&`, `^`, `<<`, `>>`; comparison operators: `==`, `!=`, `>`, `<`, `>=`, `<=`; short-circuit logical operators: `||`, `&&`. All of these behave as in C++, however, Rust is a bit stricter about the types the operators can be applied to - the bitwise operators can only be applied to integers and the logical operators can only be applied to booleans. Rust has the `-` unary operator which negates a number. The `!` operator negates a boolean and inverts every bit on an integer type (equivalent to `~` in C++ in the latter case). Rust has compound assignment operators as in C++, e.g., `+=`, but does not have increment or decrement operators (e.g., `++`).

Unique pointers

Rust is a systems language and therefore must give you raw access to memory. It does this (as in C++) via pointers. Pointers are one area where Rust and C++ are very different, both in syntax and semantics. Rust enforces memory safety by type checking pointers. That is one of its major advantages over other languages. Although the type system is a bit complex, you get memory safety and bare-metal performance in return.

I had intended to cover all of Rust's pointers in one post, but I think the subject is too large. So this post will cover just one kind - unique pointers - and other kinds will be covered in follow up posts.

First, an example without pointers:

```
fn foo() {
    let x = 75;

    // ... do something with `x` ...
}
```

When we reach the end of `foo`, `x` goes out of scope (in Rust as in C++). That means the variable can no longer be accessed and the memory for the variable can be reused.

In Rust, for every type `T` we can write `Box<T>` for an owning (aka unique) pointer to `T`. We use `Box::new(...)` to allocate space on the heap and initialise that space with the supplied value. This is similar to `new` in C++. For example,

```
fn foo() {
    let x = Box::new(75);
}
```

Here `x` is a pointer to a location on the heap which contains the value `75`. `x` has type `Box<int>`; we could have written `let x: Box<int> = Box::new(75);`. This is similar to writing `int* x = new int(75);` in C++. Unlike in C++, Rust will tidy up the memory for us, so there is no need to call `free` or `delete`. Unique pointers behave similarly to values - they are deleted when the variable goes out of scope. In our example, at the end of the function `foo`, `x` can no longer be accessed and the memory pointed at by `x` can be reused.

Owning pointers are dereferenced using the `*` as in C++. E.g.,

```
fn foo() {
    let x = Box::new(75);
    println!("`x` points to {}", *x);
}
```

As with primitive types in Rust, owning pointers and the data they point to are immutable by default. Unlike C, you can't have a mutable (unique) pointer to immutable data or vice-versa. Mutability of the data follows from the pointer. E.g.,

```
fn foo() {
    let x = Box::new(75);
    let y = Box::new(42);
    // x = y;           // Not allowed, x is immutable.
    // *x = 43;        // Not allowed, *x is immutable.
    let mut x = Box::new(75);
    x = y;             // OK, x is mutable.
    *x = 43;          // OK, *x is mutable.
}
```

Owning pointers can be returned from a function and continue to live on. If they are returned, then their memory will not be freed, i.e., there are no dangling pointers in Rust. The memory will not leak. However, it will eventually go out of scope and then it will be free. E.g.,

```
fn foo() -> Box<i32> {
    let x = Box::new(75);
    x
}

fn bar() {
    let y = foo();
    // ... use y ...
}
```

Here, memory is initialised in `foo`, and returned to `bar`. `x` is returned from `foo` and stored in `y`, so it is not deleted. At the end of `bar`, `y` goes out of scope and so the memory is reclaimed.

Owning pointers are unique (also called linear) because there can be only one (owning) pointer to any piece of memory at any time. This is accomplished by move semantics. When one pointer points at a value, any previous pointer can no longer be accessed. E.g.,


```
fn foo() {
    let x = Box::new(75);
    let y = x;
    // x can no longer be accessed
    // let z = *x; // Error.
}
```

Likewise, if an owning pointer is passed to another function or stored in a field, it can no longer be accessed:

```
fn bar(y: Box<int>) {
}

fn foo() {
    let x = Box::new(75);
    bar(x);
    // x can no longer be accessed
    // let z = *x; // Error.
}
```

Rust's unique pointers are similar to C++ `std::unique_ptr`s. In Rust, as in C++, there can be only one unique pointer to a value and that value is deleted when the pointer goes out of scope. Rust does most of its checking statically rather than at runtime. So, in C++ accessing a unique pointer whose value has moved will result in a runtime error (since it will be null). In Rust this produces a compile time error and you cannot go wrong at runtime.

We'll see later that it is possible to create other pointer types which point at a unique pointer's value in Rust. This is similar to C++. However, in C++ this allows you to cause errors at runtime by holding a pointer to freed memory. That is not possible in Rust (we'll see how when we cover Rust's other pointer types).

As shown above, owning pointers must be dereferenced to use their values. However, method calls automatically dereference, so there is no need for a `->` operator or to use `*` for method calls. In this way, Rust pointers are a bit similar to both pointers and references in C++. E.g.,

```
fn bar(x: Box<Foo>, y: Box<Box<Box<Box<Box<Foo>>>>) {
    x.foo();
    y.foo();
}
```

Assuming that the type `Foo` has a method `foo()`, both these expressions are OK.

Calling `Box::new()` with an existing value does not take a reference to that value, it copies that value. So,

```
fn foo() {  
    let x = 3;  
    let mut y = Box::new(x);  
    *y = 45;  
    println!("x is still {}", x);  
}
```

In general, Rust has move rather than copy semantics (as seen above with unique pointers). Primitive types have copy semantics, so in the above example the value `3` is copied, but for more complex values it would be moved. We'll cover this in more detail later.

Sometimes when programming, however, we need more than one reference to a value. For that, Rust has borrowed pointers. I'll cover those in the next post.

Borrowed pointers

In the last post I introduced unique pointers. This time I will talk about another kind of pointer which is much more common in most Rust programs: borrowed pointers (aka borrowed references, or just references).

If we want to have a reference to an existing value (as opposed to creating a new value on the heap and pointing to it, as with unique pointers), we must use `&`, a borrowed reference. These are probably the most common kind of pointer in Rust, and if you want something to fill in for a C++ pointer or reference (e.g., for passing a parameter to a function by reference), this is probably it.

We use the `&` operator to create a borrowed reference and to indicate reference types, and `*` to dereference them. The same rules about automatic dereferencing apply as for unique pointers. For example,

```
fn foo() {
    let x = &3; // type: &i32
    let y = *x; // 3, type: i32
    bar(x, *x);
    bar(&y, y);
}

fn bar(z: &i32, i: i32) {
    // ...
}
```

The `&` operator does not allocate memory (we can only create a borrowed reference to an existing value) and if a borrowed reference goes out of scope, no memory gets deleted.

Borrowed references are not unique - you can have multiple borrowed references pointing to the same value. E.g.,

```
fn foo() {
    let x = 5; // type: i32
    let y = &x; // type: &i32
    let z = y; // type: &i32
    let w = y; // type: &i32
    println!("These should all 5: {} {} {}", *w, *y, *z);
}
```

Like values, borrowed references are immutable by default. You can also use `&mut` to take a mutable reference, or to denote mutable reference types. Mutable borrowed references are unique (you can only take a single mutable reference to a value, and you can only have a mutable reference if there are no immutable references). You can use mutable reference where an immutable one is wanted, but not vice versa. Putting all that together in an example:

```
fn bar(x: &i32) { ... }
fn bar_mut(x: &mut i32) { ... } // &mut i32 is a reference to an i32 which
                                // can be mutated

fn foo() {
    let x = 5;
    //let xr = &mut x;           // Error - can't make a mutable reference to an
                                // immutable variable

    let xr = &x;                // Ok (creates an immutable ref)
    bar(xr);
    //bar_mut(xr);              // Error - expects a mutable ref

    let mut x = 5;
    let xr = &x;                // Ok (creates an immutable ref)
    //*xr = 4;                  // Error - mutating immutable ref
    //let xr = &mut x;          // Error - there is already an immutable ref, so we
                                // can't make a mutable one

    let mut x = 5;
    let xr = &mut x;            // Ok (creates a mutable ref)
    *xr = 4;                    // Ok
    //let xr = &x;              // Error - there is already a mutable ref, so we
                                // can't make an immutable one

    //let xr = &mut x;          // Error - can only have one mutable ref at a time
    bar(xr);                    // Ok
    bar_mut(xr);                // Ok
}
```

Note that the reference may be mutable (or not) independently of the mutableness of the variable holding the reference. This is similar to C++ where pointers can be `const` (or not) independently of the data they point to. This is in contrast to unique pointers, where the mutableness of the pointer is linked to the mutableness of the data. For example,

```
fn foo() {
    let mut x = 5;
    let mut y = 6;
    let xr = &mut x;
    //xr = &mut y;          // Error xr is immutable

    let mut x = 5;
    let mut y = 6;
    let mut xr = &mut x;
    xr = &mut y;          // Ok

    let mut x = 5;
    let mut y = 6;
    let mut xr = &x;
    xr = &y;              // Ok - xr is mut, even though the referenced data is not
}
```

If a mutable value is borrowed, it becomes immutable for the duration of the borrow. Once the borrowed pointer goes out of scope, the value can be mutated again. This is in contrast to unique pointers, which once moved can never be used again. For example,

```
fn foo() {
    let mut x = 5;          // type: i32
    {
        let y = &x;        // type: &i32
        //x = 4;          // Error - x has been borrowed
        println!("{}", x); // Ok - x can be read
    }
    x = 4;                 // OK - y no longer exists
}
```

The same thing happens if we take a mutable reference to a value - the value still cannot be modified. In general in Rust, data can only ever be modified via one variable or pointer. Furthermore, since we have a mutable reference, we can't take an immutable reference. That limits how we can use the underlying value:

```
fn foo() {
    let mut x = 5;          // type: i32
    {
        let y = &mut x;    // type: &mut i32
        //x = 4;          // Error - x has been borrowed
        //println!("{}", x); // Error - requires borrowing x
    }
    x = 4;                 // OK - y no longer exists
}
```

Unlike C++, Rust won't automatically reference a value for you. So if a function takes a parameter by reference, the caller must reference the actual parameter. However, pointer types will automatically be converted to a reference:

```
fn foo(x: &i32) { ... }

fn bar(x: i32, y: Box<i32>) {
    foo(&x);
    // foo(x); // Error - expected &i32, found i32
    foo(y);    // Ok
    foo(&*y);  // Also ok, and more explicit, but not good style
}
```

mut VS const

At this stage it is probably worth comparing `mut` in Rust to `const` in C++. Superficially they are opposites. Values are immutable by default in Rust and can be made mutable by using `mut`. Values are mutable by default in C++, but can be made constant by using `const`. The subtler and more important difference is that C++ const-ness applies only to the current use of a value, whereas Rust's immutability applies to all uses of a value. So in C++ if I have a `const` variable, someone else could have a non-const reference to it and it could change without me knowing. In Rust if you have an immutable variable, you are guaranteed it won't change.

As we mentioned above, all mutable variables are unique. So if you have a mutable value, you know it is not going to change unless you change it. Furthermore, you can change it freely since you know that no one else is relying on it not changing.

Borrowing and lifetimes

One of the primary safety goals of Rust is to avoid dangling pointers (where a pointer outlives the memory it points to). In Rust, it is impossible to have a dangling borrowed reference. It is only legal to create a borrowed reference to memory which will be alive longer than the reference (well, at least as long as the reference). In other words, the lifetime of the reference must be shorter than the lifetime of the referenced value.

That has been accomplished in all the examples in this post. Scopes introduced by `{ }` or functions are bounds on lifetimes - when a variable goes out of scope its lifetime ends. If we try to take a reference to a shorter lifetime, such as in a narrower scope, the compiler will give us an error. For example,

```
fn foo() {  
    let x = 5;  
    let mut xr = &x; // Ok - x and xr have the same lifetime  
    {  
        let y = 6;  
        //xr = &y // Error - xr will outlive y  
    } // y is released here  
} // x and xr are released here
```

In the above example, `x` and `xr` don't have the same lifetime because `xr` starts later than `x`, but it's the end of lifetimes which is more interesting, since you can't reference a variable before it exists in any case - something else which Rust enforces and which makes it safer than C++.

Explicit lifetimes

After playing with borrowed pointers for a while, you'll probably come across borrowed pointers with an explicit lifetime. These have the syntax `&'a T` (cf `&T`). They're kind of a big topic since I need to cover lifetime-polymorphism at the same time so I'll leave it for another post (there are a few more less common pointer types to cover first though). For now, I just want to say that `&T` is a shorthand for `&'a T` where `a` is the current scope, that is the scope in which the type is declared.

Reference counted and raw pointers

TODO remove all references to Gc pointers

TODO add discussion of custom pointers and Deref trait (maybe later, not here)

So far we've covered unique and borrowed pointers. Unique pointers are very similar to the new `std::unique_ptr` in C++ and borrowed references are the 'default' pointer you usually reach for if you would use a pointer or reference in C++. Rust has a few more, rarer pointers either in the libraries or built in to the language. These are mostly similar to various kinds of smart pointers you might be used to in C++.

This post took a while to write and I still don't like it. There are a lot of loose ends here, both in my write up and in Rust itself. I hope some will get better with later posts and some will get better as the language develops. If you are learning Rust, you might even want to skip this stuff for now, hopefully you won't need it. Its really here just for completeness after the posts on other pointer types.

It might feel like Rust has a lot of pointer types, but it is pretty similar to C++ once you think about the various kinds of smart pointers available in libraries. In Rust, however, you are more likely to meet them when you first start learning the language. Because Rust pointers have compiler support, you are also much less likely to make errors when using them.

I'm not going to cover these in as much detail as unique and borrowed references because, frankly, they are not as important. I might come back to them in more detail later on.

Rc

Reference counted pointers come as part of the rust standard library. They are in the `std::rc` module (we'll cover modules soon-ish. The modules are the reason for the `use` incantations in the examples). A reference counted pointer to an object of type `T` has type `Rc<T>`. You create reference counted pointers using a static method (which for now you can think of like C++'s, but we'll see later they are a bit different) - `Rc::new(...)` which takes a value to create the pointer to. This constructor method follows Rust's usual move/copy semantics (like we discussed for unique pointers) - in either case, after calling `Rc::new`, you will only be able to access the value via the pointer.

As with the other pointer types, the `.` operator does all the dereferencing you need it to. You can use `*` to manually dereference.

To pass a ref-counted pointer you need to use the `clone` method. This kinda sucks, and hopefully we'll fix that, but that is not for sure (sadly). You can take a (borrowed) reference to the pointed at value, so hopefully you don't need to clone too often. Rust's type system ensures that the ref-counted variable will not be deleted before any references expire. Taking a reference has the added advantage that it doesn't need to increment or decrement the ref count, and so will give better performance (although, that difference is probably marginal since `Rc` objects are limited to a single thread and so the ref count operations don't have to be atomic). As in C++, you can also take a reference to the `Rc` pointer.

An `Rc` example:

```
use std::rc::Rc;

fn bar(x: Rc<int>) { }
fn baz(x: &int) { }

fn foo() {
    let x = Rc::new(45);
    bar(x.clone()); // Increments the ref-count
    baz(&*x);       // Does not increment
    println!("{}", 100 - *x);
} // Once this scope closes, all Rc pointers are gone, so ref-count == 0
// and the memory will be deleted.
```

Ref counted pointers are always immutable. If you want a mutable ref-counted object you need to use a `RefCell` (or `Cell`) wrapped in an `Rc`.

Cell and RefCell

`Cell` and `RefCell` are structs which allow you to 'cheat' the mutability rules. This is kind of hard to explain without first covering Rust data structures and how they work with mutability, so I'm going to come back to these slightly tricky objects later. For now, you should know that if you want a mutable, ref counted object you need a `Cell` or `RefCell` wrapped in an `Rc`. As a first approximation, you probably want `Cell` for primitive data and `RefCell` for objects with move semantics. So, for a mutable, ref-counted int you would use `Rc<Cell<int>>`.

*T - unsafe pointers

TODO raw pointers must be const or mut now

Finally Rust has unsafe pointers. These are denoted `*T` and are created using `&` (you might need to specify a type to get a `*T` rather than a `&T` since the `&` operator can create either a borrowed reference or an unsafe pointer). These are like C pointers, just a pointer to memory with no restrictions on how they are used (you can't do pointer arithmetic without casting, but you can do it that way if you must). Unsafe pointers are the only pointer type in Rust which can be null. There is no automatic dereferencing of unsafe pointers (so to call a method you have to write `(*x).foo()`) and no automatic referencing. The most important restriction is that they can't be dereferenced (and thus can't be used) outside of an unsafe block. In regular Rust code you can only pass them around.

So, what is unsafe code? Rust has strong safety guarantees, and (rarely) they prevent you doing something you need to do. Since Rust aims to be a systems language, it has to be able to do anything that is possible and sometimes that means doing things the compiler can't verify is safe. To accomplish that, Rust has the concept of unsafe blocks, marked by the `unsafe` keyword. In unsafe code you can do unsafe things - dereference an unsafe pointer, index into an array without bounds checking, call code written in another language via the FFI, or cast variables. Obviously, you have to be much more careful writing unsafe code than writing regular Rust code. In fact, you should only very rarely write unsafe code. Mostly it is used in very small chunks in libraries, rather than in client code. In unsafe code you must do all the things you normally do in C++ to ensure safety. Furthermore, you must ensure that by the time the unsafe block finishes, you have re-established all of the invariants that the Rust compiler would usually enforce, otherwise you risk causing bugs in safe code too.

An example of using an unsafe pointer:

```
fn foo() {
    let x = 5;
    let xp: *int = &5;
    println!("x+5={}", add_5(xp));
}

fn add_5(p: *int) -> int {
    unsafe {
        if !p.is_null() { // Note that *-pointers do not auto-deref, so this is
                          // a method implemented on *int, not int.
            *p + 5
        } else {
            -1 // Not a recommended error handling strategy.
        }
    }
}
```

As with borrowed references, unsafe pointers are immutable by default and can be made mutable using the `mut` keyword, for example `*mut int`.

And that concludes our tour of Rust's pointers. Next time we'll take a break from pointers and look at Rust's data structures. We'll come back to borrowed references again in a later post though.

Data types

In this post I'll discuss Rust's data types. These are roughly equivalent to classes, structs, and enums in C++. One difference with Rust is that data and behaviour are much more strictly separated in Rust than C++ (or Java, or other OO languages). Behaviour is defined by functions and those can be defined in traits and `impl s` (implementations), but traits cannot contain data, they are similar to Java's interfaces in that respect. I'll cover traits and `impl s` in a later post, this one is all about data.

Structs

A rust struct is similar to a C struct or a C++ struct without methods. Simply a list of named fields. The syntax is best seen with an example:

```
struct S {
    field1: int,
    field2: SomeOtherStruct
}
```

Here we define a struct called `s` with two fields. The fields are comma separated; if you like, you can comma-terminate the last field too.

Structs introduce a type. In the example, we could use `s` as a type. `SomeOtherStruct` is assumed to be another struct (used as a type in the example), and (like C++) it is included by value, that is, there is no pointer to another struct object in memory.

Fields in structs are accessed using the `.` operator and their name. An example of struct use:

```
fn foo(s1: S, s2: &S) {
    let f = s1.field1;
    if f == s2.field1 {
        println!("field1 matches!");
    }
}
```

Here `s1` is struct object passed by value and `s2` is a struct object passed by reference. As with method call, we use the same `.` to access fields in both, no need for `->`.

Structs are initialised using struct literals. These are the name of the struct and values for each field. For example,

```
fn foo(sos: SomeOtherStruct) {
    let x = S { field1: 45, field2: sos }; // initialise x with a struct literal
    println!("x.field1 = {}", x.field1);
}
```

Structs cannot be recursive, that is you can't have cycles of struct names involving definitions and field types. This is because of the value semantics of structs. So for example, `struct R { r: Option<R> }` is illegal and will cause a compiler error (see below for more about `Option`). If you need such a structure then you should use some kind of pointer; cycles with pointers are allowed:

```
struct R {
    r: Option<Box<R>>
}
```

If we didn't have the `Option` in the above struct, there would be no way to instantiate the struct and Rust would signal an error.

Structs with no fields do not use braces in either their definition or literal use. Definitions do need a terminating semi-colon though, presumably just to facilitate parsing.

```
struct Empty;

fn foo() {
    let e = Empty;
}
```

Tuples

Tuples are anonymous, heterogeneous sequences of data. As a type, they are declared as a sequence of types in parentheses. Since there is no name, they are identified by structure. For example, the type `(int, int)` is a pair of integers and `(i32, f32, S)` is a triple. Tuple values are initialised in the same way as tuple types are declared, but with values instead of types for the components, e.g., `(4, 5)`. An example:

```
// foo takes a struct and returns a tuple
fn foo(x: SomeOtherStruct) -> (i32, f32, S) {
    (23, 45.82, S { field1: 54, field2: x })
}
```

Tuples can be used by destructuring using a `let` expression, e.g.,

```
fn bar(x: (int, int)) {
    let (a, b) = x;
    println!("x was ({} , {})", a, b);
}
```

We'll talk more about destructuring next time.

Tuple structs

Tuple structs are named tuples, or alternatively, structs with unnamed fields. They are declared using the `struct` keyword, a list of types in parentheses, and a semicolon. Such a declaration introduces their name as a type. Their fields must be accessed by destructuring (like a tuple), rather than by name. Tuple structs are not very common.

```
struct IntPoint (int, int);

fn foo(x: IntPoint) {
    let IntPoint(a, b) = x; // Note that we need the name of the tuple
                           // struct to destructure.
    println!("x was ({} , {})", a, b);
}
```

Enums

Enums are types like C++ enums or unions, in that they are types which can take multiple values. The simplest kind of enum is just like a C++ enum:

```
enum E1 {
    Var1,
    Var2,
    Var3
}

fn foo() {
    let x: E1 = Var2;
    match x {
        Var2 => println!("var2"),
        _ => {}
    }
}
```

However, Rust enums are much more powerful than that. Each variant can contain data. Like tuples, these are defined by a list of types. In this case they are more like unions than enums in C++. Rust enums are tagged unions rather untagged (as in C++), that means you can't mistake one variant of an enum for another at runtime. An example:

```
enum Expr {
    Add(int, int),
    Or(bool, bool),
    Lit(int)
}

fn foo() {
    let x = Or(true, false); // x has type Expr
}
```

Many simple cases of object-oriented polymorphism are better handled in Rust using enums.

To use enums we usually use a match expression. Remember that these are similar to C++ switch statements. I'll go into more depth on these and other ways to destructure data next time. Here's an example:

```
fn bar(e: Expr) {
    match e {
        Add(x, y) => println!("An `Add` variant: {} + {}", x, y),
        Or(..) => println!("An `Or` variant"),
        _ => println!("Something else (in this case, a `Lit`)",)
    }
}
```

Each arm of the match expression matches a variant of `Expr`. All variants must be covered. The last case (`_`) covers all remaining variants, although in the example there is only `Lit`. Any data in a variant can be bound to a variable. In the `Add` arm we are binding the two ints in an `Add` to `x` and `y`. If we don't care about the data, we can use `..` to match any data, as we do for `or`.

Option

One particularly common enum in Rust is `option`. This has two variants - `Some` and `None`. `None` has no data and `Some` has a single field with type `T` (`option` is a generic enum, which we will cover later, but hopefully the general idea is clear from C++). Options are used to indicate a value might be there or might not. Any place you use a null pointer in C++ to indicate a value which is in some way undefined, uninitialised, or false, you should probably

use an `Option` in Rust. Using `Option` is safer because you must always check it before use; there is no way to do the equivalent of dereferencing a null pointer. They are also more general, you can use them with values as well as pointers. An example:

```
use std::rc::Rc;

struct Node {
    parent: Option<Rc<Node>>,
    value: int
}

fn is_root(node: Node) -> bool {
    match node.parent {
        Some(_) => false,
        None => true
    }
}
```

Here, the `parent` field could be either a `None` or a `Some` containing an `Rc<Node>`. In the example, we never actually use that payload, but in real life you usually would.

There are also convenience methods on `Option`, so you could write the body of `is_root` as `node.is_none()` or `!node.is_some()`.

Inherited mutability and Cell/RefCell

Local variables in Rust are immutable by default and can be marked mutable using `mut`. We don't mark fields in structs or enums as mutable, their mutability is inherited. This means that a field in a struct object is mutable or immutable depending on whether the object itself is mutable or immutable. Example:


```

struct S1 {
    field1: int,
    field2: S2
}
struct S2 {
    field: int
}

fn main() {
    let s = S1 { field1: 45, field2: S2 { field: 23 } };
    // s is deeply immutable, the following mutations are forbidden
    // s.field1 = 46;
    // s.field2.field = 24;

    let mut s = S1 { field1: 45, field2: S2 { field: 23 } };
    // s is mutable, these are OK
    s.field1 = 46;
    s.field2.field = 24;
}

```

Inherited mutability in Rust stops at references. This is similar to C++ where you can modify a non-const object via a pointer from a const object. If you want a reference field to be mutable, you have to use `&mut` on the field type:

```

struct S1 {
    f: int
}
struct S2<'a> {
    f: &'a mut S1 // mutable reference field
}
struct S3<'a> {
    f: &'a S1 // immutable reference field
}

fn main() {
    let mut s1 = S1{f:56};
    let s2 = S2 { f: &mut s1};
    s2.f.f = 45; // legal even though s2 is immutable
    // s2.f = &mut s1; // illegal - s2 is not mutable
    let s1 = S1{f:56};
    let mut s3 = S3 { f: &s1};
    s3.f = &s1; // legal - s3 is mutable
    // s3.f.f = 45; // illegal - s3.f is immutable
}

```

(The `'a` parameter on `s2` and `s3` is a lifetime parameter, we'll cover those soon).

Sometimes whilst an object is logically immutable, it has parts which need to be internally mutable. Think of various kinds of caching or a reference count (which would not give true logical immutability since the effect of changing the ref count can be observed via destructors). In C++, you would use the `mutable` keyword to allow such mutation even when the object is `const`. In Rust we have the `Cell` and `RefCell` structs. These allow parts of immutable objects to be mutated. Whilst that is useful, it means you need to be aware that when you see an immutable object in Rust, it is possible that some parts may actually be mutable.

`RefCell` and `Cell` let you get around Rust's strict rules on mutation and aliasability. They are safe to use because they ensure that Rust's invariants are respected dynamically, even though the compiler cannot ensure that those invariants hold statically. `Cell` and `RefCell` are both single threaded objects.

Use `Cell` for types which have copy semantics (pretty much just primitive types). `Cell` has `get` and `set` methods for changing the stored value, and a `new` method to initialise the cell with a value. `Cell` is a very simple object - it doesn't need to do anything smart since objects with copy semantics can't keep references elsewhere (in Rust) and they can't be shared across threads, so there is not much to go wrong.

Use `RefCell` for types which have move semantics, that means nearly everything in Rust, struct objects are a common example. `RefCell` is also created using `new` and has a `set` method. To get the value in a `RefCell`, you must borrow it using the borrow methods (`borrow` , `borrow_mut` , `try_borrow` , `try_borrow_mut`) these will give you a borrowed reference to the object in the `RefCell`. These methods follow the same rules as static borrowing - you can only have one mutable borrow, and can't borrow mutably and immutably at the same time. However, rather than a compile error you get a runtime failure. The `try_` variants return an `Option` - you get `Some(val)` if the value can be borrowed and `None` if it can't. If a value is borrowed, calling `set` will fail too.

Here's an example using a ref-counted pointer to a `RefCell` (a common use-case):

```
use std::rc::Rc;
use std::cell::RefCell;

Struct S {
    field: int
}

fn foo(x: Rc<RefCell<S>>) {
    {
        let s = x.borrow();
        println!("the field, twice {} {}", s.f, x.borrow().field);
        // let s = x.borrow_mut(); // Error - we've already borrowed the contents of x
    }

    let s = x.borrow_mut(); // O, the earlier borrows are out of scope
    s.f = 45;
    // println!("The field {}", x.borrow().field); // Error - can't mut and immut borrow
    println!("The field {}", s.f);
}
```

If you're using Cell/RefCell, you should try to put them on the smallest object you can. That is, prefer to put them on a few fields of a struct, rather than the whole struct. Think of them like single threaded locks, finer grained locking is better since you are more likely to avoid colliding on a lock.

Destructuring

Last time we looked at Rust's data types. Once you have some data structure, you will want to get that data out. For structs, Rust has field access, just like C++. For tuples, tuple structs, and enums you must use destructuring (there are various convenience functions in the library, but they use destructuring internally). Destructuring of data structures doesn't happen in C++, but it might be familiar from languages such as Python or various functional languages. The idea is that just as you can create a data structure by filling out its fields with data from a bunch of local variables, you can fill out a bunch of local variables with data from a data structure. From this simple beginning, destructuring has become one of Rust's most powerful features. To put it another way, destructuring combines pattern matching with assignment into local variables.

Destructuring is done primarily through the `let` and `match` statements. The `match` statement is used when the structure being destructured can have different variants (such as an enum). A `let` expression pulls the variables out into the current scope, whereas `match` introduces a new scope. To compare:

```
fn foo(pair: (int, int)) {
    let (x, y) = pair;
    // we can now use x and y anywhere in foo

    match pair {
        (x, y) => {
            // x and y can only be used in this scope
        }
    }
}
```

The syntax for patterns (used after `let` and before `=>` in the above example) in both cases is (pretty much) the same. You can also use these patterns in argument position in function declarations:

```
fn foo((x, y): (int, int)) {
}
```

(Which is more useful for structs or tuple-structs than tuples).

Most initialisation expressions can appear in a destructuring pattern and they can be arbitrarily complex. That can include references and primitive literals as well as data structures. For example,

```

struct St {
    f1: int,
    f2: f32
}

enum En {
    Var1,
    Var2,
    Var3(int),
    Var4(int, St, int)
}

fn foo(x: &En) {
    match x {
        &Var1 => println!("first variant"),
        &Var3(5) => println!("third variant with number 5"),
        &Var3(x) => println!("third variant with number {} (not 5)", x),
        &Var4(3, St { f1: 3, f2: x }, 45) => {
            println!("destructuring an embedded struct, found {} in f2", x)
        }
        &Var4(_, x, _) => {
            println!("Some other Var4 with {} in f1 and {} in f2", x.f1, x.f2)
        }
        _ => println!("other (Var2)")
    }
}

```

Note how we destructure through a reference by using `&` in the patterns and how we use a mix of literals (`5`, `3`, `st { ... }`), wildcards (`_`), and variables (`x`).

You can use `_` wherever a variable is expected if you want to ignore a single item in a pattern, so we could have used `&Var3(_)` if we didn't care about the integer. In the first `Var4` arm we destructure the embedded struct (a nested pattern) and in the second `Var4` arm we bind the whole struct to a variable. You can also use `..` to stand in for all fields of a tuple or struct. So if you wanted to do something for each enum variant but don't care about the content of the variants, you could write:

```

fn foo(x: En) {
    match x {
        Var1 => println!("first variant"),
        Var2 => println!("second variant"),
        Var3(..) => println!("third variant"),
        Var4(..) => println!("fourth variant")
    }
}

```

When destructuring structs, the fields don't need to be in order and you can use `..` to elide the remaining fields. E.g.,

```
struct Big {
    field1: int,
    field2: int,
    field3: int,
    field4: int,
    field5: int,
    field6: int,
    field7: int,
    field8: int,
    field9: int,
}

fn foo(b: Big) {
    let Big { field6: x, field3: y, .. } = b;
    println!("pulled out {} and {}", x, y);
}
```

As a shorthand with structs you can use just the field name which creates a local variable with that name. The `let` statement in the above example created two new local variables `x` and `y`. Alternatively, you could write

```
fn foo(b: Big) {
    let Big { field6, field3, .. } = b;
    println!("pulled out {} and {}", field3, field6);
}
```

Now we create local variables with the same names as the fields, in this case `field3` and `field6`.

There are a few more tricks to Rust's destructuring. Lets say you want a reference to a variable in a pattern. You can't use `&` because that matches a reference, rather than creates one (and thus has the effect of dereferencing the object). For example,

```
struct Foo {
    field: &'static int
}

fn foo(x: Foo) {
    let Foo { field: &y } = x;
}
```

Here, `y` has type `int` and is a copy of the field in `x`.

To create a reference to something in a pattern, you use the `ref` keyword. For example,

```
fn foo(b: Big) {  
    let Big { field3: ref x, ref field6, .. } = b;  
    println!("pulled out {} and {}", *x, *field6);  
}
```

Here, `x` and `field6` both have type `&int` and are references to the fields in `b`.

One last trick when destructuring is that if you are destructuring a complex object, you might want to name intermediate objects as well as individual fields. Going back to an earlier example, we had the pattern `&Var4(3, St{ f1: 3, f2: x }, 45)`. In that pattern we named one field of the struct, but you might also want to name the whole struct object. You could write `&Var4(3, s, 45)` which would bind the struct object to `s`, but then you would have to use field access for the fields, or if you wanted to only match with a specific value in a field you would have to use a nested match. That is not fun. Rust lets you name parts of a pattern using `@` syntax. For example `&Var4(3, s @ St{ f1: 3, f2: x }, 45)` lets us name both a field (`x`, for `f2`) and the whole struct (`s`).

That just about covers your options with Rust pattern matching. There are a few features I haven't covered, such as matching vectors, but hopefully you know how to use `match` and `let` and have seen some of the powerful things you can do. Next time I'll cover some of the subtle interactions between `match` and borrowing which tripped me up a fair bit when learning Rust.

Destructuring pt2 - match and borrowing

When destructuring there are some surprises in store where borrowing is concerned. Hopefully, nothing surprising once you understand borrowed references really well, but worth discussing (it took me a while to figure out, that's for sure. Longer than I realised, in fact, since I screwed up the first version of this blog post).

Imagine you have some `&Enum` variable `x` (where `Enum` is some enum type). You have two choices: you can match `*x` and list all the variants (`Variant1 => ...`, etc.) or you can match `x` and list reference to variant patterns (`&Variant1 => ...`, etc.). (As a matter of style, prefer the first form where possible since there is less syntactic noise). `x` is a borrowed reference and there are strict rules for how a borrowed reference can be dereferenced, these interact with match expressions in surprising ways (at least surprising to me), especially when you a modifying an existing enum in a seemingly innocuous way and then the compiler explodes on a match somewhere.

Before we get into the details of the match expression, lets recap Rust's rules for value passing. In C++, when assigning a value into a variable or passing it to a function there are two choices - pass-by-value and pass-by-reference. The former is the default case and means a value is copied either using a copy constructor or a bitwise copy. If you annotate the destination of the parameter pass or assignment with `&`, then the value is passed by reference - only a pointer to the value is copied and when you operate on the new variable, you are also operating on the old value.

Rust has the pass-by-reference option, although in Rust the source as well as the destination must be annotated with `&`. For pass-by-value in Rust, there are two further choices - copy or move. A copy is the same as C++'s semantics (except that there are no copy constructors in Rust). A move copies the value but destroys the old value - Rust's type system ensures you can no longer access the old value. As examples, `int` has copy semantics and `Box<int>` has move semantics:

```
fn foo() {
    let x = 7i;
    let y = x;           // x is copied
    println!("x is {}", x); // OK

    let x = box 7i;
    let y = x;           // x is moved
    //println!("x is {}", x); // error: use of moved value: `x`
}
```


Rust determines if an object has move or copy semantics by looking for destructors. Destructors probably need a post of their own, but for now, an object in Rust has a destructor if it implements the `Drop` trait. Just like C++, the destructor is executed just before an object is destroyed. If an object has a destructor then it has move semantics. If it does not, then all of its fields are examined and if any of those do then the whole object has move semantics. And so on down the object structure. If no destructors are found anywhere in an object, then it has copy semantics.

Now, it is important that a borrowed object is not moved, otherwise you would have a reference to the old object which is no longer valid. This is equivalent to holding a reference to an object which has been destroyed after going out of scope - it is a kind of dangling pointer. If you have a pointer to an object, there could be other references to it. So if an object has move semantics and you have a pointer to it, it is unsafe to dereference that pointer. (If the object has copy semantics, dereferencing creates a copy and the old object will still exist, so other references will be fine).

OK, back to match expressions. As I said earlier, if you want to match some `x` with type `&T` you can dereference once in the match clause or match the reference in every arm of the match expression. Example:

```
enum Enum1 {
    Var1,
    Var2,
    Var3
}

fn foo(x: &Enum1) {
    match *x { // Option 1: deref here.
        Var1 => {}
        Var2 => {}
        Var3 => {}
    }

    match x {
        // Option 2: 'deref' in every arm.
        &Var1 => {}
        &Var2 => {}
        &Var3 => {}
    }
}
```

In this case you can take either approach because `Enum1` has copy semantics. Let's take a closer look at each approach: in the first approach we dereference `x` to a temporary variable with type `Enum1` (which copies the value in `x`) and then do a match against the three variants of `Enum1`. This is a 'one level' match because we don't go deep into the value's type. In the second approach there is no dereferencing. We match a value with type

`&Enum1` against a reference to each variant. This match goes two levels deep - it matches the type (always a reference) and looks inside the type to match the referred type (which is `Enum1`).

Either way, we must ensure that we (that is, the compiler) must ensure we respect Rust's invariants around moves and references - we must not move any part of an object if it is referenced. If the value being matched has copy semantics, that is trivial. If it has move semantics then we must make sure that moves don't happen in any match arm. This is accomplished either by ignoring data which would move, or making references to it (so we get by-reference passing rather than by-move).

```
enum Enum2 {
    // Box has a destructor so Enum2 has move semantics.
    Var1(Box<int>),
    Var2,
    Var3
}

fn foo(x: &Enum2) {
    match *x {
        // We're ignoring nested data, so this is OK
        Var1(..) => {}
        // No change to the other arms.
        Var2 => {}
        Var3 => {}
    }

    match x {
        // We're ignoring nested data, so this is OK
        &Var1(..) => {}
        // No change to the other arms.
        &Var2 => {}
        &Var3 => {}
    }
}
```

In either approach we don't refer to any of the nested data, so none of it is moved. In the first approach, even though `x` is referenced, we don't touch its innards in the scope of the dereference (i.e., the match expression) so nothing can escape. We also don't bind the whole value (i.e., bind `*x` to a variable), so we can't move the whole object either.

We can take a reference to any variant in the second match, but not in the dereferenced version. So, in the second approach replacing the second arm with `a @ &Var2 => {}` is OK (`a` is a reference), but under the first approach we couldn't write `a @ Var2 => {}` since that would mean moving `*x` into `a`. We could write `ref a @ Var2 => {}` (in which `a` is also a reference), although it's not a construct you see very often.

But what about if we want to use the data nested inside `var1` ? We can't write:

```
match *x {
  Var1(y) => {}
  _ => {}
}
```

or

```
match x {
  &Var1(y) => {}
  _ => {}
}
```

because in both cases it means moving part of `x` into `y` . We can use the 'ref' keyword to get a reference to the data in `var1 : &Var1(ref y) => {}` . That is OK, because now we are not dereferencing anywhere and thus not moving any part of `x` . Instead we are creating a pointer which points into the interior of `x` .

Alternatively, we could destructure the Box (this match is going three levels deep):

`&Var1(box y) => {}` . This is OK because `int` has copy semantics and `y` is a copy of the `int` inside the `Box` inside `var1` (which is 'inside' a borrowed reference). Since `int` has copy semantics, we don't need to move any part of `x` . We could also create a reference to the int rather than copy it: `&Var1(box ref y) => {}` . Again, this is OK, because we don't do any dereferencing and thus don't need to move any part of `x` . If the contents of the Box had move semantics, then we could not write `&Var1(box y) => {}` , we would be forced to use the reference version. We could also use similar techniques with the first approach to matching, which look the same but without the first `&` . For example, `Var1(box ref y) => {}` .

Now lets get more complex. Lets say you want to match against a pair of reference-to-enum values. Now we can't use the first approach at all:

```
fn bar(x: &Enum2, y: &Enum2) {
    // Error: x and y are being moved.
    // match (*x, *y) {
    //     (Var2, _) => {}
    //     _ => {}
    // }

    // OK.
    match (x, y) {
        (&Var2, _) => {}
        _ => {}
    }
}
```

The first approach is illegal because the value being matched is created by dereferencing `x` and `y` and then moving them both into a new tuple object. So in this circumstance, only the second approach works. And of course, you still have to follow the rules above for avoiding moving parts of `x` and `y`.

If you do end up only being able to get a reference to some data and you need the value itself, you have no option except to copy that data. Usually that means using `clone()`. If the data doesn't implement `clone`, you're going to have to further destructure to make a manual copy or implement `clone` yourself.

What if we don't have a reference to a value with move semantics, but the value itself. Now moves are OK, because we know no one else has a reference to the value (the compiler ensures that if they do, we can't use the value). For example,

```
fn baz(x: Enum2) {
    match x {
        Var1(y) => {}
        _ => {}
    }
}
```

There are still a few things to be aware of. Firstly, you can only move to one place. In the above example we are moving part of `x` into `y` and we'll forget about the rest. If we wrote `a @ Var1(y) => {}` we would be attempting to move all of `x` into `a` and part of `x` into `y`. That is not allowed, an arm like that is illegal. Making one of `a` or `y` a reference (using `ref a`, etc.) is not an option either, then we'd have the problem described above where we move whilst holding a reference. We can make both `a` and `y` references and then we're OK - neither is moving, so `x` remains in tact and we have pointers to the whole and a part of it.

Similarly (and more common), if we have a variant with multiple pieces of nested data, we can't take a reference to one datum and move another. For example if we had a `Var4` declared as `Var4(Box<int>, Box<int>)` we can have a match arm which references both (`Var4(ref y, ref z) => {}`) or a match arm which moves both (`Var4(y, z) => {}`) but you cannot have a match arm which moves one and references the other (`Var4(ref y, z) => {}`). This is because a partial move still destroys the whole object, so the reference would be invalid.

Arrays and Vectors

Rust arrays are pretty different from C arrays. For starters they come in statically and dynamically sized flavours. These are more commonly known as fixed length arrays and slices. As we'll see, the former is kind of a bad name since both kinds of array have fixed (as opposed to growable) length. For a growable 'array', Rust provides the `Vec` collection.

Fixed length arrays

The length of a fixed length array is known statically and features in it's type. E.g., `[i32; 4]` is the type of an array of `i32` s with length four.

Array literal and array access syntax is the same as C:

```
let a: [i32; 4] = [1, 2, 3, 4];    // As usual, the type annotation is optional.
println!("The second element is {}", a[1]);
```

You'll notice that array indexing is zero-based, just like C.

However, unlike C/C++, array indexing is bounds checked. In fact all access to arrays is bounds checked, which is another way Rust is a safer language.

If you try to do `a[4]`, then you will get a runtime panic. Unfortunately, the Rust compiler is not clever enough to give you a compile time error, even when it is obvious (as in this example).

If you like to live dangerously, or just need to get every last ounce of performance out of your program, you can still get unchecked access to arrays. To do this, use the `get_unchecked` method on an array. Unchecked array accesses must be inside an `unsafe` block. You should only need to do this in the rarest circumstances.

Just like other data structures in Rust, arrays are immutable by default and mutability is inherited. Mutation is also done via the indexing syntax:

```
let mut a = [1, 2, 3, 4];
a[3] = 5;
println!("{:?}", a);
```

And just like other data, you can borrow an array by taking a reference to it:

```
fn foo(a: &[i32; 4]) {
    println!("First: {}; last: {}", a[0], a[3]);
}

fn main() {
    foo(&[1, 2, 3, 4]);
}
```

Notice that indexing still works on a borrowed array.

This is a good time to talk about the most interesting aspect of Rust arrays for C++ programmers - their representation. Rust arrays are value types: they are allocated on the stack like other values and an array object is a sequence of values, not a pointer to those values (as in C). So from our examples above, `let a = [1_i32, 2, 3, 4];` will allocate 16 bytes on the stack and executing `let b = a;` will copy 16 bytes. If you want a C-like array, you have to explicitly make a pointer to the array, this will give you a pointer to the first element.

A final point of difference between arrays in Rust and C++ is that Rust arrays can implement traits, and thus have methods. To find the length of an array, for example, you use `a.len()`.

Slices

A slice in Rust is just an array whose length is not known at compile time. The syntax of the type is just like a fixed length array, except there is no length: e.g., `[i32]` is a slice of 32 bit integers (with no statically known length).

There is a catch with slices: since the compiler must know the size of all objects in Rust, and it can't know the size of a slice, then we can never have a value with slice type. If you try and write `fn foo(x: [i32])`, for example, the compiler will give you an error.

So, you must always have pointers to slices (there are some very technical exceptions to this rule so that you can implement your own smart pointers, but you can safely ignore them for now). You must write `fn foo(x: &[i32])` (a borrowed reference to a slice) or `fn foo(x: *mut [i32])` (a mutable raw pointer to a slice), etc.

The simplest way to create a slice is by coercion. There are far fewer implicit coercions in Rust than there are in C++. One of them is the coercion from fixed length arrays to slices. Since slices must be pointer values, this is effectively a coercion between pointers. For example, we can coerce `&[i32; 4]` to `&[i32]`, e.g.,

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

Here the right hand side is a fixed length array of length four, allocated on the stack. We then take a reference to it (type `&[i32; 4]`). That reference is coerced to type `&[i32]` and given the name `a` by the `let` statement.

Again, access is just like C (using `[...]`), and access is bounds checked. You can also check the length yourself by using `len()`. So clearly the length of the array is known somewhere. In fact all arrays of any kind in Rust have known length, since this is essential for bounds checking, which is an integral part of memory safety. The size is known dynamically (as opposed to statically in the case of fixed length arrays), and we say that slice types are dynamically sized types (DSTs, there are other kinds of dynamically sized types too, they'll be covered elsewhere).

Since a slice is just a sequence of values, the size cannot be stored as part of the slice. Instead it is stored as part of the pointer (remember that slices must always exist as pointer types). A pointer to a slice (like all pointers to DSTs) is a fat pointer - it is two words wide, rather than one, and contains the pointer to the data plus a payload. In the case of slices, the payload is the length of the slice.

So in the example above, the pointer `a` will be 128 bits wide (on a 64 bit system). The first 64 bits will store the address of the `1` in the sequence `[1, 2, 3, 4]`, and the second 64 bits will contain `4`. Usually, as a Rust programmer, these fat pointers can just be treated as regular pointers. But it is good to know about (it can affect the things you can do with casts, for example).

Slicing notation and ranges

A slice can be thought of as a (borrowed) view of an array. So far we have only seen a slice of the whole array, but we can also take a slice of part of an array. There is a special notation for this which is like the indexing syntax, but takes a range instead of a single integer. E.g., `a[0..4]`, which takes a slice of the first four elements of `a`. Note that the range is exclusive at the top and inclusive at the bottom. Examples:

```
let a: [i32; 4] = [1, 2, 3, 4];
let b: &[i32] = &a; // Slice of the whole array.
let c = &a[0..4]; // Another slice of the whole array, also has type &[i32].
let c = &a[1..3]; // The middle two elements, &[i32].
let c = &a[1..]; // The last three elements.
let c = &a[..3]; // The first three element.
let c = &a[..]; // The whole array, again.
let c = &b[1..3]; // We can also slice a slice.
```


Note that in the last example, we still need to borrow the result of slicing. The slicing syntax produces an unborrowed slice (type: `[i32]`) which we must then borrow (to give a `&[i32]`), even if we are slicing a borrowed slice.

Range syntax can also be used outside of slicing syntax. `a..b` produces an iterator which runs from `a` to `b-1`. This can be combined with other iterators in the usual way, or can be used in `for` loops:

```
// Print all numbers from 1 to 10.
for i in 1..11 {
    println!("{}", i);
}
```

Vecs

A vector is heap allocated and is an owning reference. Therefore (and like `Box<_>`), it has move semantics. We can think of a fixed length array analogously to a value, a slice to a borrowed reference. Similarly, a vector in Rust is analogous to a `Box<_>` pointer.

It helps to think of `Vec<_>` as a kind of smart pointer, just like `Box<_>`, rather than as a value itself. Similarly to a slice, the length is stored in the 'pointer', in this case the 'pointer' is the `Vec` value.

A vector of `i32`s has type `Vec<i32>`. There are no vector literals, but we can get the same effect by using the `vec!` macro. We can also create an empty vector using `Vec::new()`:

```
let v = vec![1, 2, 3, 4]; // A Vec<i32> with length 4.
let v: Vec<i32> = Vec::new(); // An empty vector of i32s.
```

In the second case above, the type annotation is necessary so the compiler can know what the vector is a vector of. If we were to use the vector, the type annotation would probably not be necessary.

Just like arrays and slices, we can use indexing notation to get a value from the vector (e.g., `v[2]`). Again, these are bounds checked. We can also use slicing notation to take a slice of a vector (e.g., `&v[1..3]`).

The extra feature of vectors is that their size can change - they can get longer or shorter as needed. For example, `v.push(5)` would add the element `5` to the end of the vector (this would require that `v` is mutable). Note that growing a vector can cause reallocation, which for large vectors can mean a lot of copying. To guard against this you can pre-allocate space in a vector using `with_capacity`, see the [Vec docs](#) for more details.

The Index traits

Note for readers: there is a lot of material in this section that I haven't covered properly yet. If you're following the tutorial, you can skip this section, it is a somewhat advanced topic in any case.

The same indexing syntax used for arrays and vectors is also used for other collections, such as `HashMap`s. And you can use it yourself for your own collections. You opt-in to using the indexing (and slicing) syntax by implementing the `Index` trait. This is a good example of how Rust makes available nice syntax to user types, as well as built-ins (`Deref` for dereferencing smart pointers, as well as `Add` and various other traits, work in a similar way).

The `Index` trait looks like

```
pub trait Index<Idx: ?Sized> {
    type Output: ?Sized;

    fn index(&self, index: Idx) -> &Self::Output;
}
```

`Idx` is the type used for indexing. For most uses of indexing this is `usize`. For slicing this is one of the `std::ops::Range` types. `Output` is the type returned by indexing, this will be different for each collection. For slicing it will be a slice, rather than the type of a single element. `index` is a method which does the work of getting the element(s) out of the collection. Note that the collection is taken by reference and the method returns a reference to the element with the same lifetime.

Let's look at the implementation for `Vec` to see how what an implementation looks like:

```
impl<T> Index<usize> for Vec<T> {
    type Output = T;

    fn index(&self, index: usize) -> &T {
        &(**self)[index]
    }
}
```

As we said above, indexing is done using `usize`. For a `Vec<T>`, indexing will return a single element of type `T`, thus the value of `Output`. The implementation of `index` is a bit weird - `(**self)` gets a view of the whole vec as a slice, then we use indexing on slices to get the element, and finally take a reference to it.

If you have your own collections, you can implement `Index` in a similar way to get indexing and slicing syntax for your collection.

Initialiser syntax

As with all data in Rust, arrays and vectors must be properly initialised. Often you just want an array full of zeros to start with and using the array literal syntax is a pain. So Rust gives you a little syntactic sugar to initialise an array full of a given value: `[value; len]`. So for example to create an array with length 100 full of zeros, we'd use `[0; 100]`.

Similarly for vectors, `vec![42; 100]` would give you a vector with 100 elements, each with the value 42.

The initial value is not limited to integers, it can be any expression. For array initialisers, the length must be an integer constant expression. For `vec!`, it can be any expression with type `usize`.

Graphs and arena allocation

(Note you can run the examples in this chapter by downloading this directory and running `cargo run`).

Graphs are a bit awkward to construct in Rust because of Rust's stringent lifetime and mutability requirements. Graphs of objects are very common in OO programming. In this tutorial I'm going to go over a few different approaches to implementation. My preferred approach uses arena allocation and makes slightly advanced use of explicit lifetimes. I'll finish up by discussing a few potential Rust features which would make using such an approach easier.

A [graph](#) is a collection of nodes with edges between some of those nodes. Graphs are a generalisation of lists and trees. Each node can have multiple children and multiple parents (we usually talk about edges into and out of a node, rather than parents/children). Graphs can be represented by adjacency lists or adjacency matrices. The former is basically a node object for each node in the graph, where each node object keeps a list of its adjacent nodes. An adjacency matrix is a matrix of booleans indicating whether there is an edge from the row node to the column node. We'll only cover the adjacency list representation, adjacency matrices have very different issues which are less Rust-specific.

There are essentially two orthogonal problems: how to handle the lifetime of the graph and how to handle its mutability.

The first problem essentially boils down to what kind of pointer to use to point to other nodes in the graph. Since graph-like data structures are recursive (the types are recursive, even if the data is not) we are forced to use pointers of some kind rather than have a totally value-based structure. Since graphs can be cyclic, and ownership in Rust cannot be cyclic, we cannot use `Box<Node>` as our pointer type (as we might do for tree-like data structures or linked lists).

No graph is truly immutable. Because there may be cycles, the graph cannot be created in a single statement. Thus, at the very least, the graph must be mutable during its initialisation phase. The usual invariant in Rust is that all pointers must either be unique or immutable. Graph edges must be mutable (at least during initialisation) and there can be more than one edge into any node, thus no edges are guaranteed to be unique. So we're going to have to do something a little bit advanced to handle mutability.

One solution is to use mutable raw pointers (`*mut Node`). This is the most flexible approach, but also the most dangerous. You must handle all the lifetime management yourself without any help from the type system. You can make very flexible and efficient data structures this

way, but you must be very careful. This approach handles both the lifetime and mutability issues in one fell swoop. But it handles them by essentially ignoring all the benefits of Rust - you will get no help from the compiler here (it's also not particularly ergonomic since raw pointers don't automatically (de-)reference). Since a graph using raw pointers is not much different from a graph in C++, I'm not going to cover that option here.

The options you have for lifetime management are reference counting (shared ownership, using `Rc<...>`) or arena allocation (all nodes have the same lifetime, managed by an arena; using borrowed references `&...`). The former is more flexible (you can have references from outside the graph to individual nodes with any lifetime), the latter is better in every other way.

For managing mutability, you can either use `RefCell`, i.e., make use of Rust's facility for dynamic, interior mutability, or you can manage the mutability yourself (in this case you have to use `UnsafeCell` to communicate the interior mutability to the compiler). The former is safer, the latter is more efficient. Neither is particularly ergonomic.

Note that if your graph might have cycles, then if you use `Rc`, further action is required to break the cycles and not leak memory. Since Rust has no cycle collection of `Rc` pointers, if there is a cycle in your graph, the ref counts will never fall to zero, and the graph will never be deallocated. You can solve this by using `Weak` pointers in your graph or by manually breaking cycles when you know the graph should be destroyed. The former is more reliable. We don't cover either here, in our examples we just leak memory. The approach using borrowed references and arena allocation does not have this issue and is thus superior in that respect.

To compare the different approaches I'll use a pretty simple example. We'll just have a `Node` object to represent a node in the graph, this will hold some string data (representative of some more complex data payload) and a `Vec` of adjacent nodes (`edges`). We'll have an `init` function to create a simple graph of nodes, and a `traverse` function which does a pre-order, depth-first traversal of the graph. We'll use this to print the payload of each node in the graph. Finally, we'll have a `Node::first` method which returns a reference to the first adjacent node to the `self` node and a function `foo` which prints the payload of an individual node. These functions stand in for more complex operations involving manipulation of a node interior to the graph.

To try and be as informative as possible without boring you, I'll cover two combinations of possibilities: ref counting and `RefCell`, and arena allocation and `UnsafeCell`. I'll leave the other two combinations as an exercise.

```
Rc<RefCell<Node>>
```

See [full example](#).

This is the safer option because there is no unsafe code. It is also the least efficient and least ergonomic option. It is pretty flexible though, nodes of the graph can be easily reused outside the graph since they are ref-counted. I would recommend this approach if you need a fully mutable graph, or need your nodes to exist independently of the graph.

The node structure looks like

```
struct Node {
    datum: &'static str,
    edges: Vec<Rc<RefCell<Node>>>,
}
```

Creating a new node is not too bad: `Rc::new(RefCell::new(Node { ... }))`. To add an edge during initialisation, you have to borrow the start node as mutable, and clone the end node into the Vec of edges (this clones the pointer, incrementing the reference count, not the actual node). E.g.,

```
let mut mut_root = root.borrow_mut();
mut_root.edges.push(b.clone());
```

The `RefCell` dynamically ensures that we are not already reading or writing the node when we write it.

Whenever you access a node, you have to use `.borrow()` to borrow the `RefCell`. Our `first` method has to return a ref-counted pointer, rather than a borrowed reference, so callers of `first` also have to borrow:

```
fn first(&self) -> Rc<RefCell<Node>> {
    self.edges[0].clone()
}

pub fn main() {
    let g = ...;
    let f = g.first();
    foo(&*f.borrow());
}
```

`&Node` and `UnsafeCell`

See [full example](#).

In this approach we use borrowed references as edges. This is nice and ergonomic and lets us use our nodes with 'regular' Rust libraries which primarily operate with borrowed references (note that one nice thing about ref counted objects in Rust is that they play nicely with the lifetime system. We can create a borrowed reference into the `Rc` to directly (and safely) reference the data. In the previous example, the `RefCell` prevents us doing this, but an `Rc / UnsafeCell` approach should allow it).

Destruction is correctly handled too - the only constraint is that all the nodes must be destroyed at the same time. Destruction and allocation of nodes is handled using an arena.

On the other hand, we do need to use quite a few explicit lifetimes. Unfortunately we don't benefit from lifetime elision here. At the end of the section I'll discuss some future directions for the language which could make things better.

During construction we will mutate our nodes which might be multiply referenced. This is not possible in safe Rust code, so we must initialise inside an `unsafe` block. Since our nodes are mutable and multiply referenced, we must use an `UnsafeCell` to communicate to the Rust compiler that it cannot rely on its usual invariants.

When is this approach feasible? The graph must only be mutated during initialisation. In addition, we require that all nodes in the graph have the same lifetime (we could relax these constraints somewhat to allow adding nodes later as long as they can all be destroyed at the same time). Similarly, we could rely on more complicated invariants for when the nodes can be mutated, but it pays to keep things simple, since the programmer is responsible for safety in those respects.

Arena allocation is a memory management technique where a set of objects have the same lifetime and can be deallocated at the same time. An arena is an object responsible for allocating and deallocating the memory. Since large chunks of memory are allocated and deallocated at once (rather than allocating individual objects), arena allocation is very efficient. Usually, all the objects are allocated from a contiguous chunk of memory, that improves cache coherency when you are traversing the graph.

In Rust, arena allocation is supported by the [libarena](#) crate and is used throughout the compiler. There are two kinds of arenas - typed and untyped. The former is more efficient and easier to use, but can only allocate objects of a single type. The latter is more flexible and can allocate any object. Arena allocated objects all have the same lifetime, which is a parameter of the arena object. The type system ensures references to arena allocated objects cannot live longer than the arena itself.

Our node struct must now include the lifetime of the graph, `'a`. We wrap our `Vec` of adjacent nodes in an `UnsafeCell` to indicate that we will mutate it even when it should be immutable:

```
struct Node<'a> {
    datum: &'static str,
    edges: UnsafeCell<Vec<&'a Node<'a>>>,
}
```

Our new function must also use this lifetime and must take as an argument the arena which will do the allocation:

```
fn new<'a>(datum: &'static str, arena: &'a TypedArena<Node<'a>>) -> &'a Node<'a> {
    arena.alloc(Node {
        datum: datum,
        edges: UnsafeCell::new(Vec::new()),
    })
}
```

We use the arena to allocate the node. The lifetime of the graph is derived from the lifetime of the reference to the arena, so the arena must be passed in from the scope which covers the graph's lifetime. For our examples, that means we pass it into the `init` method. (One could imagine an extension to the type system which allows creating values at scopes outside their lexical scope, but there are no plans to add such a thing any time soon). When the arena goes out of scope, the whole graph is destroyed (Rust's type system ensures that we can't keep references to the graph beyond that point).

Adding an edge is a bit different looking:

```
(*root.edges.get()).push(b);
```

We're essentially doing the obvious `root.edges.push(b)` to push a node (`b`) on to the list of edges. However, since `edges` is wrapped in an `UnsafeCell`, we have to call `get()` on it. That gives us a mutable raw pointer to edges (`*mut Vec<&Node>`), which allows us to mutate `edges`. However, it also requires us to manually dereference the pointer (raw pointers do not auto-deref), thus the `(*...)` construction. Finally, dereferencing a raw pointer is unsafe, so the whole lot has to be wrapped up in an unsafe block.

The interesting part of `traverse` is:

```
for n in &(*self.edges.get()) {
    n.traverse(f, seen);
}
```


We follow the previous pattern for getting at the edges list, which requires an unsafe block. In this case we know it is in fact safe because we must be post-initialisation and thus there will be no mutation.

Again, the `first` method follows the same pattern for getting at the `edges` list. And again must be in an unsafe block. However, in contrast to the graph using `Rc<RefCell<_>>`, we can return a straightforward borrowed reference to the node. That is very convenient. We can reason that the unsafe block is safe because we do no mutation and we are post-initialisation.

```
fn first(&'a self) -> &'a Node<'a> {
    unsafe {
        (*self.edges.get())[0]
    }
}
```

Future language improvements for this approach

I believe that arena allocation and using borrowed references are an important pattern in Rust. We should do more in the language to make these patterns safer and easier to use. I hope use of arenas becomes more ergonomic with the ongoing work on [allocators](#). There are three other improvements I see:

Safe initialisation

There has been lots of research in the OO world on mechanisms for ensuring mutability only during initialisation. How exactly this would work in Rust is an open research question, but it seems that we need to represent a pointer which is mutable and not unique, but restricted in scope. Outside that scope any existing pointers would become normal borrowed references, i.e., immutable *or* unique.

The advantage of such a scheme is that we have a way to represent the common pattern of mutable during initialisation, then immutable. It also relies on the invariant that, while individual objects are multiply owned, the aggregate (in this case a graph) is uniquely owned. We should then be able to adopt the reference and `UnsafeCell` approach, without the `UnsafeCell`s and the unsafe blocks, making that approach more ergonomic and more safer.

Alex Summers and Julian Viereck at ETH Zurich are investigating this further.

Generic modules

The 'lifetime of the graph' is constant for any particular graph. Repeating the lifetime is just boilerplate. One way to make this more ergonomic would be to allow the graph module to be parameterised by the lifetime, so it would not need to be added to every struct, impl, and function. The lifetime of the graph would still need to be specified from outside the module, but hopefully inference would take care of most uses (as it does today for function calls).

See [ref_graph_generic_mod.rs](#) for how that might look. (We should also be able to use safe initialisation (proposed above) to remove the unsafe code).

See also this [RFC issue](#).

This feature would vastly reduce the syntactic overhead of the reference and `UnsafeCell` approach.

Lifetime elision

We currently allow the programmer to elide some lifetimes in function signatures to improve ergonomics. One reason the `&Node` approach to graphs is a bit ugly is because it doesn't benefit from any of the lifetime elision rules.

A common pattern in Rust is data structures with a common lifetime. References into such data structures give rise to types like `&'a Foo<'a>`, for example `&'a Node<'a>` in the graph example. It would be nice to have an elision rule that helps in this case. I'm not really sure how it should work though.

Looking at the example with generic modules, it doesn't look like we need to extend the lifetime elision rules very much (I'm not actually sure if `Node::new` would work without the given lifetimes, but it seems like a fairly trivial extension to make it work if it doesn't). We might want to add some new rule to allow elision of module-generic lifetimes if they are the only ones in scope (other than `'static`), but I'm not sure how that would work with multiple in- scope lifetimes (see the `foo` and `init` functions, for example).

If we don't add generic modules, we might still be able to add an elision rule specifically to target `&'a Node<'a>`, not sure how though.