

Rust programming

Module F: safe and unsafe rust

Unsafe: Learning objectives

- when to reach for `unsafe` code
- reason about undefined behavior
- familiarity with raw pointers in rust
- practical experience with raw pointers, C strings and untagged unions

Content overview

- Why is unsafe needed?
- Undefined behavior and optimizations
- Break
- common types in unsafe code
- examples of unsafe usage

Rust guarantees that references are valid

for any `&T` or `&mut T`

- the address is not `NULL`
- the address is well-aligned for type `T`
- it points into memory belonging to the process

These guarantees make rust memory safe

The borrow checker

The borrow + type checker ensures that these conditions are met

- We want a 100% guarantee that when the compiler says that really means our program is correct
- An analysis that is wrong in 1 out of 100 cases is worthless

The borrow checker is conservative

- "if it is not a hell yes, it's a no"

fail-proof borrow checker

```
1 fn borrow_checker<P>(program: P) -> bool {  
2     false  
3 }
```

all programs that my borrow checker accepts are memory safe!

Hence

There are (many) correct programs that the rust borrow checker does not accept

Hence

There are (many) useful programs that the rust borrow checker does not accept

- interacting with other languages (FFI)
- interacting with the OS/hardware
- optimization

Unsafe, morally

In rust "unsafe" means "I, the programmer, am responsible for checking the correctness of this code"

The type and borrow checker are still in fully effect. But we can use types like raw pointers on which the conditions that the type/borrow checker places are less strict.

Unsafe in code

- unsafe blocks: "programmer must check the rules"

```
1 // BAD!
2 let reference: &u8 = unsafe {
3     let ptr = 0usize as *const u8;
4
5     &*ptr
6 };
```

- unsafe functions: "programmer must check the preconditions"

```
1 unsafe fn foobar() {
2     ...
3 }
```

- unsafe impl: "programmer must check impl is valid"

```
1 unsafe impl Send for MyType {}
```

Undefined Behavior & Optimizations

```
1 // std::hint::unreachable_unchecked
2 pub const unsafe fn unreachable_unchecked() -> !
```

- `unsafe fn`: to call this function, the programmer has to check the preconditions
- returns "never", the type of an infinite loop (diverging computation)

Undefined Behavior & Optimizations

```
1  if expensive_pure_computation() == 0 {  
2      println!("hello there");  
3      unsafe { std::hint::unreachable_unchecked() }  
4  } else {  
5      different_computation()  
6  }
```

- that print is unreachable if the rest of the branch is unreachable

Undefined Behavior & Optimizations

```
1  if expensive_pure_computation() == 0 {  
2      unsafe { std::hint::unreachable_unchecked() }  
3  } else {  
4      different_computation()  
5  }
```

- actually the whole branch is unreachable

Undefined Behavior & Optimizations

```
1 expensive_pure_computation() == 0;  
2 different_computation()
```

- actually that whole condition does not need to be computed

Undefined Behavior & Optimizations

```
1  if expensive_pure_computation() == 0 {  
2      println!("hello there");  
3      unsafe { std::hint::unreachable_unchecked() }  
4  } else {  
5      different_computation()  
6  }
```

becomes just

```
1  different_computation()
```

- but if the condition turns out to be reachable, behavior is confusing

Undefined Behavior & Optimizations

- misusing `unreachable_unchecked` is very explicit. There are many more subtle ways to introduce UB

```
1 // BAD!
2 let reference: &u8 = unsafe {
3     let ptr = 0usize as *const u8;
4
5     &*ptr
6 };
```

- the rust compiler assumes that references are valid, so this snippet contains UB!
- LLVM encodes and exploits assumptions like `nonnull` or `noalias`

```
1 define internal fastcc void @str.RocStr.reallocate(
2     %str.RocStr* noalias nocapture nonnull %arg,
3     %str.RocStr* nocapture nonnull readonly align 8 %arg1,
4     i64 %arg2
5 )
```

`transmute`

e.g. bitcast a `i64` into a `f64`:

```
1 std::mem::transmute::<i64, f64>(42i64)
```

there are still checks! `transmute` errors when the size does not correspond

```
1 error[E0512]: cannot transmute between types of different sizes, or dependently-sized types
2   --> src/main.rs:2:12
3     |
4 2 |   unsafe { std::mem::transmute::<i64, f32>(42i64) };
5     |           ^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^
6     |
7     = note: source type: `i64` (64 bits)
8     = note: target type: `f32` (32 bits)
```

`transmute`

Only some bit patterns are valid for a type. Creating an invalid bit pattern is UB!

Only ``0b0000_0000`` and ``0b0000_0001`` are valid ``bool`` bit patterns. This code has UB:

```
1  std::mem::transmute::<u8, bool>(2u8)
```

An ``if`` statement might be compiled into a jump table

```
1  const JMP_TABLE: *const u8 = [ 0x1000, 0x1100 ];  
2  
3  // this will fail horribly if `bool_value >= 2`  
4  jmp JMP_TABLE[bool_value as usize];
```

The memory representation of rust values is explicitly undefined! Bitcasting is therefore very unsafe!

So Rust is just as bad as C?

- if memory safety can be broken, how is rust any better than C?

So far

- rust is a systems language: it must provide unrestricted access
 - call code in different languages
 - exploit all capabilities of the OS/hardware
 - optimize
- rust's goals mean restricting access
- unsafe is an escape hatch: great power, but the risk of introducing UB

Part II: common types and examples

Raw Pointers

```
1 let mut x = 0;
2 let y = &mut x as *mut i32;
3 let z = 12;
4
5 unsafe {
6     std::ptr::write(y, z);
7     assert_eq!(std::ptr::read(y), 12);
8 }
```

- raw (mut or const) pointers can alias each other!

NonNull

A `*mut T` that is guaranteed to not be NULL

```
1 use std::ptr::NonNull;
2
3 let mut x = 0u32;
4 let ptr = unsafe { NonNull::new_unchecked(&mut x as *mut _) };
5
6 // NEVER DO THIS!!! This is undefined behavior.
7 let ptr = unsafe { NonNull::<u32>::new_unchecked(std::ptr::null_mut()) };
```


MaybeUninit

Working with uninitialized memory

```
1 use std::mem::MaybeUninit;
2
3 let b: bool = unsafe { MaybeUninit::uninit().assume_init() }; // undefined behavior! ☹
```

- Useful when working with pointers (which may point to uninitialized data)

```
1 pub const unsafe fn swap<T>(x: *mut T, y: *mut T) {
2     // Give ourselves some scratch space to work with.
3     // We do not have to worry about drops: `MaybeUninit` does nothing when dropped.
4     let mut tmp = MaybeUninit::<T>::uninit();
5
6     // Perform the swap
7     // SAFETY: the caller must guarantee that `x` and `y` are
8     // valid for writes and properly aligned. `tmp` cannot be
9     // overlapping either `x` or `y` because `tmp` was just allocated
10    // on the stack as a separate allocated object.
11    unsafe {
12        std::ptr::copy_nonoverlapping(x, tmp.as_mut_ptr(), 1);
13        std::ptr::copy(y, x, 1); // `x` and `y` may overlap
14        std::ptr::copy_nonoverlapping(tmp.as_ptr(), y, 1);
15    }
16 }
```

CString

A null-terminated string type

```
1 use std::ffi::CString;
2 use libc::strlen;
3
4 fn main() {
5     let cstring = CString::new("Hello, world!").expect("no NULL bytes");
6
7     // pub unsafe extern "C" fn strlen(cs: *const c_char) -> size_t
8     println!("{}", unsafe { strlen(cstring.as_ptr())});
9 }
```

Examples

- interacting with other languages (FFI)
- interacting with the OS/hardware
- optimization

Using libc functions

```
1 // pub unsafe extern "C" fn getpid() -> pid_t
2
3 use libc;
4
5 println!("My pid is {}", unsafe { libc::getpid() });
```

Using libc functions

```
1 // pub fn id() -> u32
2
3 use std::process;
4
5 println!("My pid is {}", process::id());
```

Interacting with the OS

```
1  unsafe fn execve(&self, argv: &[*const c_char], envp: &[*const c_char]) -> c_int {
2      match self {
3          // ...
4
5          #[cfg(target_family = "windows")]
6          ExecutableFile::OnDisk(_, path) => {
7              let path_cstring = CString::new(path.to_str().unwrap()).unwrap();
8
9              libc::execve(path_cstring.as_ptr().cast(), argv.as_ptr(), envp.as_ptr())
10         }
11     }
12 }
```

Interacting with the Hardware

Using a SIMD intrinsic

```
1  #[target_feature(enable = "avx")]
2  unsafe fn vperilps(mut current: __m128, mask: (i32, i32, i32, i32)) -> __m128 {
3      let mask = _mm_set_epi32(mask.3, mask.2, mask.1, mask.0);
4
5      std::arch::asm!(
6          "vpermilps {a:y}, {a:y}, {m:y}",
7          a = inout(ymm_reg) current,
8          m = in(ymm_reg) mask,
9
10     );
11
12     current
13 }
```

Example: Memory consumption of linked lists

```
1  enum LinkedList {
2      Nil,
3      Cons(u64, Box<LinkedList>),
4  }
5
6  use LinkedList::*;
7
8  impl LinkedList {
9      fn range(range: Range<u64>) -> Self {
10         let mut list = Nil;
11         for value in range.rev() {
12             list = Cons(value, Box::new(list));
13         }
14
15         list
16     }
17
18     fn sum(&self) -> u64 {
19         match self {
20             Nil => 0,
21             Cons(first, rest) => first + rest.sum(),
22         }
23     }
24 }
```


Example: Memory consumption of linked lists

```
1  enum LinkedList {
2      Nil,
3      Cons(u64, Box<LinkedList>),
4  }
5
6  // could be represented as
7
8  struct LinkedList {
9      tag: LinkedListTag,
10     payload: LinkedListUnion,
11 }
12
13 enum LinkedListTag {
14     Nil = 0,
15     Cons = 1,
16 }
17
18 union LinkedListUnion {
19     nil: (),
20     cons: (u64, std::mem::ManuallyDrop<Box<LinkedList>>),
21 }
```

Example: Memory consumption of linked lists

- what is the memory layout of this type?

```
1  struct LinkedList {
2      tag: LinkedListTag,
3      payload: LinkedListUnion,
4  }
5
6  enum LinkedListTag {
7      Nil = 0,
8      Cons = 1,
9  }
10
11 union LinkedListUnion {
12     nil: (),
13     cons: (u64, std::mem::ManuallyDrop<Box<LinkedList>>),
14 }
```

- field order
- alignment
- size

Example: Memory consumption of linked lists

```
1  struct LinkedList(*const Node);
2
3  struct Node {
4      first: u64,
5      rest: LinkedList,
6  }
7
8  impl LinkedList {
9      fn range(range: Range<u64>) -> Self {
10         let mut list = LinkedList(std::ptr::null());
11         for value in range.rev() {
12             let node = Node { first: value, rest: list };
13             list = LinkedList(Box::into_raw(Box::new(node)));
14         }
15
16         list
17     }
18
19     fn sum(&self) -> u64 {
20         if self.is_null() { 0 } else {
21             let node = unsafe { std::ptr::read(self) };
22             node.first + node.rest.sum()
23         }
24     }
25 }
```

Example: Memory consumption of linked lists

```
1 struct LinkedList(Option<Box<Node>>);
2
3 struct Node {
4     first: u64,
5     rest: LinkedList,
6 }
7
8 impl LinkedList {
9     fn range(range: Range<u64>) -> Self {
10         todo!()
11     }
12
13     fn sum(&self) -> u64 {
14         todo!()
15     }
16 }
```

Question: what is the memory layout of LinkedList. What is the size?

Exercises

- Implement functions for the pointer-based `LinkedList`
- Implement a process forwarding program using `execve`
- Implement a custom `Result` variant that matches a specific memory layout

Summary

- rust is a systems language: it must provide unrestricted access
 - call code in different languages
 - exploit all capabilities of the OS/hardware
 - optimize
- rust's goals mean restricting access
- unsafe is an escape hatch: great power, but the risk of introducing UB
- common types in unsafe code: `*const T`, `*mut T`, `CString`, `MaybeUninit`
- examples of unsafe
 - using the `execve` syscall wrapper
 - using custom simd instructions
 - optimizing a linked list with pointer trickery

