



Self-Referential Structs in Rust

February 2025



Pete LeVasseur
Eclipse uProtocol
Maintainer





background

What's a self-referential struct in Rust?

 **It's everyone's favorite recurring topic: self-referential structs**

 help

```
struct Struct { view: &mut u8, stashed: u8 };  
let stashed = 1u8;  
let s = Struct { view: &mut stashed, stashed };
```

Attempting to store something in a struct that references something else in the same struct

The C or C++ equivalent of the above code would compile and be usable. *Why's Rust different?*

Aliasing: Why self-referential structs are hard in Rust

“*Aliasing* occurs when one pointer or reference points to a "span" of memory that overlaps with the span of another pointer or reference. A span of memory is similar to how a slice works: there's a base byte address as well as a length in bytes.” [1]

“...to emphasize, one you didn't list here and the main one in general is that aliasing something covered by a `&mut` (or owned by a `Box`) is UB. That rule is at the heart of Rust's memory and concurrency guarantees.


Alternatively put, once you "use" something, any exclusive references to it must no longer be valid.” [2]

References:

[1] Rust's Unsafe Coding Guidelines: *Glossary* <https://rust-lang.github.io/unsafe-code-guidelines/glossary.html#aliasing>

[2] Rust Programming Language Users Forum: *It's everyone's favorite recurring topic: self-referential structs*

<https://users.rust-lang.org/t/its-everyones-favorite-recurring-topic-self-referential-structs/91105/3>



motivation

Occurrence in the wild: mcap crate

We want to:

1. use the memmap crate to memory map in an .mcap file to a Mmap
2. use the MessageStream struct from the mcap crate to take a reference to that Mmap
3. so that we can keep the Mmap alive long enough for the duration that MessageStream is alive

```
1 pub struct MemmapMessageStream<'a> {  
2     // keep the Mmap here so it's alive as long as  
3     mmap: Mmap,  
4     // ... needed by the MessageStream  
5     message_stream: MessageStream<'a>,  
6 }
```

Note: All code snippets in this presentation are sample code to explain concepts!

Trying to use the MemmapMessageStream

```
1 use anyhow::Result;
2 use mcap::{McapError, Message, MessageStream};
3 use memmap::Mmap;
4 use std::{fs::File, path::Path};
5
6 pub struct MemmapMessageStream<'a> {
7     // keep the Mmap here so it's alive as long as
8     mmap: Mmap,
9     // ... needed by the MessageStream
10    message_stream: MessageStream<'a>,
11 }
12
13 impl<'a> MemmapMessageStream<'a> {
14     pub fn new<P: AsRef<Path>>(p: P) -> Result<MemmapMessageStream<'a>> {
15         let file = File::open(p)?;
16         let mmap = unsafe { Mmap::map(&file) };
17         let message_stream = MessageStream::new(&mmap)?;
18         Ok(MemmapMessageStream {
19             mmap,
20             message_stream,
21         })
22     }
23
24     pub fn next(&mut self) -> Option<std::result::Result<Message, McapError>> {
25         self.message_stream.next()
26     }
27 }
```

```
error[E0515]: cannot return value referencing local variable `mmap`
--> src/memmap_message_stream.rs:18:9
|
17 |         let message_stream = MessageStream::new(&mmap)?;
|                                             ----- `mmap` is borrowed here
18 | /         Ok(MemmapMessageStream {
19 | |             mmap,
20 | |             message_stream,
21 | |         })
| |_____^ returns a value referencing data owned by the current function
```

```
error[E0505]: cannot move out of `mmap` because it is borrowed
--> src/memmap_message_stream.rs:19:13
|
13 |     impl<'a> MemmapMessageStream<'a> {
|     -- lifetime `'a` defined here
...
16 |         let mmap = unsafe { Mmap::map(&file) };
|         ---- binding `mmap` declared here
17 |         let message_stream = MessageStream::new(&mmap)?;
|                                             ----- borrow of `mmap` occurs here
18 | /         Ok(MemmapMessageStream {
19 | |             mmap,
| |             ^^^^ move out of `mmap` occurs here
20 | |             message_stream,
21 | |         })
| |_____ - returning this value requires that `mmap` is borrowed for `'a`
```



ouroboros

ouroboros, designed for self-referential structs

struct using ouroboros

```
1 use anyhow::Result;
2 use mcap::{McapError, Message, MessageStream};
3 use memmap::Mmap;
4 use ouroboros::self_referencing;
5 use std::{fs::File, path::Path};
6
7 #[self_referencing]
8 pub struct MemmapMessageStream {
9     // keep the Mmap here so it's alive as long as
10    mmap: Mmap,
11    // ... needed by the MessageStream
12    #[not_covariant]
13    #[borrows(mmap)]
14    message_stream: MessageStream<'this>,
15 }
```

`self_referencing` attribute macro
`#[not_covariant]`: covariance is a bit out of scope, but this says that the `'this` lifetime cannot be shortened

`'this`: the `'this` lifetime is created by the `#[self_referencing]` macro and should be used on all references marked by the `#[borrows]` macro

`#[borrows(mmap)]` informs that this struct member will borrow `mmap`

ouroboros, designed for self-referential structs *impl using ouroboros*

```
1 impl MemmapMessageStream {
2     pub fn new_stream<P: AsRef<Path>>(p: P) -> Result<MemmapMessageStream> {
3         let file = File::open(p)?;
4         let mmap = unsafe { Mmap::map(&file) }?;
5         Ok(MemmapMessageStreamBuilder {
6             mmap,
7             message_stream_builder: |m: &Mmap| {
8                 MessageStream::new(m).expect("Unable to initialize with Mmap")
9             },
10        }
11        .build())
12    }
13
14    pub fn next(&mut self) -> Option<std::result::Result<Message, McapError>> {
15        self.with_message_stream_mut(|ms| ms.next())
16    }
17 }
```

MemmapMessageStreamBuilder created by `#[self_referencing]` which allows us to pass in a closure accepting the Mmap

The `with_message_stream_mut` function created allows us mutable access to the MessageStream



macro

ouroboros, designed for self-referential structs

Expanded struct

```
cargo rustc -- -Zunpretty=expanded
```

```
1 pub struct MemmapMessageStream {
2     actual_data: ::core::mem::MaybeUninit<MemmapMessageStreamInternal>,
3 }
4 struct MemmapMessageStreamInternal {
5     #[doc(hidden)]
6     message_stream: MessageStream<'static>,
7     #[doc(hidden)]
8     mmap: ::ouroboros::macro_help::AliasableBox<Mmap>,
9 }
```


ouroboros, designed for self-referential structs

Expanded MemmapMessageStreamBuilder

message_stream_builder is an FnOnce, a closure which accepts a Mmap with lifetime 'this and returns a MessageStream of lifetime 'this

```
1 pub(super) struct MemmapMessageStreamBuilder<
2     MessageStreamBuilder_: for<'this> ::core::ops::FnOnce(&'this Mmap) -> MessageStream<'this>,
3 > {
4     pub(super) mmap: Mmap,
5     pub(super) message_stream_builder: MessageStreamBuilder_,
6 }
7 impl<
8     MessageStreamBuilder_: for<'this> ::core::ops::FnOnce(&'this Mmap) -> MessageStream<'this>,
9     > MemmapMessageStreamBuilder<MessageStreamBuilder_>
10 {
11     pub(super) fn build(self) -> MemmapMessageStream {
12         MemmapMessageStream::new(self.mmap, self.message_stream_builder)
13     }
14 }
```

ouroboros, designed for self-referential structs

Expanded MemmapMessageStream::new()

```
1 impl MemmapMessageStream {
2     pub(super) fn new(
3         mmap: Mmap,
4         message_stream_builder: impl for<'this> ::core::ops::FnOnce(&'this Mmap) -> MessageStream<'this>,
5     ) -> MemmapMessageStream {
6         let mmap = ::ouroboros::macro_help::aliasable_boxed(mmap);
7         let mmap_illegal_static_reference =
8             unsafe { ::ouroboros::macro_help::change_lifetime(&*mmap) };
9         let message_stream = message_stream_builder(mmap_illegal_static_reference);
10        unsafe {
11            Self {
12                actual_data: ::core::mem::MaybeUninit::new(MemmapMessageStreamInternal {
13                    mmap,
14                    message_stream,
15                }),
16            }
17        }
18    }
19
20    // ...
21 }
```

ouroboros, designed for self-referential structs

aliasable_boxed()

```
1 pub fn aliasable_boxed<T>(data: T) -> AliasableBox<T> {  
2     AliasableBox::from_unique(UniqueBox::new(data))  
3 }
```

rust-aliasable

Rust library providing basic aliasable (non `core::ptr::Unique`) types

Documentation hosted on docs.rs.

```
aliasable = "0.1"
```

Why?

Used for escaping `noalias` when multiple raw pointers may point to the same data.

ouroboros, designed for self-referential structs

change_lifetime()

```
1 /// Converts a reference to an object to a static reference This is
2 /// obviously unsafe because the compiler can no longer guarantee that the
3 /// data outlives the reference. It is up to the consumer to get rid of the
4 /// reference before the container is dropped. The + 'static ensures that
5 /// whatever we are referring to will remain valid indefinitely, that there
6 /// are no limitations on how long the pointer itself can live.
7 ///
8 /// # Safety
9 ///
10 /// The caller must ensure that the returned reference is not used after the originally passed
11 /// reference would become invalid.
12 pub unsafe fn change_lifetime<'old, 'new: 'old, T: 'new>(data: &'old T) -> &'new T {
13     &*(data as *const _)
14 }
```


ouroboros, designed for self-referential structs

Drop impl

The `MaybeUninit` held inside of `MemmapMessageStream` is dropped by the `Drop` impl, ensuring that we deallocate the `mmap_illegal_static_reference`

```
1 #[repr(transparent)]
2 pub struct MemmapMessageStream {
3     actual_data: ::core::mem::MaybeUninit<MemmapMessageStreamInternal>,
4 }
5 struct MemmapMessageStreamInternal {
6     #[doc(hidden)]
7     message_stream: MessageStream<'static>,
8     #[doc(hidden)]
9     mmap: ::ouroboros::macro_help::AliasableBox<Mmap>,
10 }
11 impl ::core::ops::Drop for MemmapMessageStream {
12     fn drop(&mut self) {
13         unsafe { self.actual_data.assume_init_drop() };
14     }
15 }
```

Rust stdlib documentation: *MaybeUninit*

https://doc.rust-lang.org/std/mem/union.MaybeUninit.html#method.assume_init_drop

ouroboros, designed for self-referential structs *with_message_stream_mut()*

Allows us to pass in an FnOnce to be able to do mutable operations on the MessageStream

```
1 #[inline(always)]
2 pub(super) fn with_message_stream_mut<'outer_borrow, ReturnType>(
3     &'outer_borrow mut self,
4     user: impl for<'this> ::core::ops::FnOnce(&'outer_borrow mut MessageStream<'this>) -> ReturnType,
5 ) -> ReturnType {
6     let field = &mut unsafe { self.actual_data.assume_init_mut() }.message_stream;
7     user(field)
8 }
```

Rust stdlib documentation: *MaybeUninit*

https://doc.rust-lang.org/std/mem/union.MaybeUninit.html#method.assume_init_drop

run it



Reading an MCAP file using MemmapMessageStream

```
$ cargo run --bin mmap_backed
```


```
    Compiling memmap-message-stream-works v0.1.0  
(/home/peter/presentations/self-referential-structs/memmap-message-stream-works)
```

```
    Finished `dev` profile [unoptimized + debuginfo] target(s) in 0.89s
```

```
    Running `target/debug/mmap_backed`
```

```
next: Ok(Message { channel: Channel { id: 0, topic: "ping", schema:  
Some(Schema { name: "Buffer", encoding: "", .. }), message_encoding:  
"", metadata: {} }, sequence: 0, log_time: 1736458452082329110,  
publish_time: 1736458452065969963, data: [0, 1, 0, 0, 4, 0, 0, 0, 32,  
0, 0, 0] })
```

Takeaway: the self-referential strategy to keep the Mmap alive worked!



efficiency

ouroboros, designed for self-referential structs

Use cargo-show-asm on with_message_stream_mut()

Want to ensure that
MemmapMessageStream::next()
compiles away the closure

Looks like in the generated assembly we
directly call Iterator::next() as
implemented on MessageStream. Great!

```
1 #[self_referencing]
2 pub struct MemmapMessageStream {
3     // keep the Mmap here so it's alive as long as
4     mmap: Mmap,
5     // ... needed by the MessageStream
6     #[not_covariant]
7     #[borrows(mmap)]
8     message_stream: MessageStream<'this>,
9 }
10
11 impl MemmapMessageStream {
12     // ...
13
14     #[inline(never)]
15     pub fn next(&mut self) -> Option<std::result::Result<Message, McapError>> {
16         self.with_message_stream_mut(|ms| ms.next())
17     }
18 }
```

```
1 .section ".text.memmap_message_stream_works::memmap_message_stream::<impl memmap_message_stream_works::memmap_message_stream::ouroboros_impl_memmap_mes
2 sage_stream::MemmapMessageStream>::next", "ax", @progbits
3     .p2align    4, 0x90
4 .type    memmap_message_stream_works::memmap_message_stream::<impl memmap_message_stream_works::memmap_message_stream::ouroboros_impl_memmap_message_str
5 eam::MemmapMessageStream>::next,@function
6 memmap_message_stream_works::memmap_message_stream::<impl memmap_message_stream_works::memmap_message_stream::ouroboros_impl_memmap_message_stream::Mem
7 mapMessageStream>::next:
8     .cfi_startproc
9     jmp qword ptr [rip + <mcap::read::MessageStream as core::iter::traits::iterator::Iterator>::next@GOTPCREL]
```

safety



Any undefined behavior (UB) using ouroboros?

Generally the Rust compiler takes responsibility for ensuring certain safe properties of code, e.g. memory-safety which could lead to undefined behavior (UB)

As we saw, in order to accomplish self-referential structs, ouroboros uses `unsafe` in key locations where the compiler may be too conservative and would have caused a compilation error.

“Unsafe Rust exists because, by nature, static analysis is conservative. When the compiler tries to determine whether or not code upholds the guarantees, it’s better for it to reject some valid programs than to accept some invalid programs.”

Key point: “`unsafe`” moves the responsibility for checking safe properties of code from compiler to the engineer, e.g. memory-safety. Code marked as `unsafe` deserves more attention during design and code review.

Running miri to check for Undefined Behavior (UB)

File-backed memory mappings unsupported

```
$MIRIFLAGS=-Zmiri-disable-isolation cargo miri run --bin mmap_backed
```

```
1 error: unsupported operation: Miri does not support file-backed memory mappings
2 --> /home/peter/.cargo/registry/src/index.crates.io-6f17d22bba15001f/memmap-0.7.0/src/unix.rs:49:23
3 |
4 49 |         let ptr = libc::mmap(
5 |             ^
6 50 |             ptr::null_mut(),
7 51 |             aligned_len as libc::size_t,
8 52 |             prot,
9 ... |
10 55 |             aligned_offset as libc::off_t,
11 56 |         );
12 |         ^ Miri does not support file-backed memory mappings
13 |
```

Running miri to check for Undefined Behavior (UB)

Use a Vec-backed approximation

```
$cargo miri run --bin vec_backed
```

```
error: unsupported operation: can't call foreign function `ZSTD_createDCtx` on OS `linux`  
--> /home/peter/.cargo/registry/src/index.crates.io-6f17d22bba15001f/zstd-safe-5.0.2+zstd.1.5.2/src/lib.rs:707:35  
707 |         NonNull::new(unsafe { zstd_sys::ZSTD_createDCtx() }?),  
      |         ^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^ can't call foreign function `ZSTD_createDCtx` on OS `linux`  
  
= help: if this is a basic API commonly used on this target, please report an issue with Miri  
= help: however, note that Miri does not aim to support every FFI function out there; for instance, we will not support APIs for things such as GUI  
s, scripting languages, or databases
```

Running miri to check for Undefined Behavior (UB)

Just check the ouroboros tests!

No exact test for this scenario, hmmm

Running miri to check for Undefined Behavior (UB)

Write a test for ouroboros then! - test setup

```
1 struct PhraseRef<'a> {
2     data: &'a mut String,
3 }
4
5 impl<'a> PhraseRef<'a> {
6     fn change_phrase(&mut self) {
7         *self.data = self.data.replace("Hello", "Goodbye");
8     }
9 }
10
11 #[self_referencing]
12 struct DataAndCustomRef {
13     data: String,
14     #[borrows(mut data)]
15     #[not_covariant]
16     phrase: PhraseRef<'this>,
17 }
```

Running miri to check for Undefined Behavior (UB)

Write a test for ouroboros then! - test

```
1 #[test]
2 fn custom_ref() {
3     let mut instance = DataAndCustomRefBuilder {
4         data: "Hello world!".to_owned(),
5         phrase_builder: |data| PhraseRef { data },
6     }
7     .build();
8     instance.with_phrase_mut(|phrase| phrase.change_phrase());
9     let modified_data = instance.into_heads().data;
10    assert_eq!(modified_data, "Goodbye world!");
11 }
```

```
$ cargo miri test --features="miri"
```

```
...
```

```
test ok_tests::custom_ref ... ok
```


Running miri to check for Undefined Behavior (UB)

Write a test for ouroboros then! - merged upstream

Add test for struct holding original data and a struct which holds a reference to original data. #130

Merged someguynamedj... merged 2 commits into someguynamedjosh:main from PLeVasseur:feature/test-struct-holding-ref 2 days ago

Conversation 0 Commits 2 Checks 6 Files changed 1



PLeVasseur commented 3 days ago

Contributor ...

closes [#129](#)



Reviewers


No reviews

Assignees

No one assigned

Test merged upstream:

<https://github.com/someguynamedjosh/ouroboros/pull/130>



std::pin::Pin

Question: Can we use `std::pin::Pin` instead?

Yes! Here's the code

```
1 use anyhow::Result;
2 use mcap::{McapError, Message, MessageStream};
3 use mmap::Mmap;
4 use std::{fs::File, marker::PhantomPinned, path::Path, pin::Pin};
5
6 pub struct MemmapMessageStream {
7     // No lifetime parameter needed anymore
8     mmap: Mmap,
9     message_stream: MessageStream<'static>, // We'll transmute the lifetime
10     _pin: PhantomPinned, // Prevent Unpin implementation
11 }
```

We use `PhantomPinned` to ensure that this struct's contents cannot be moved in memory

Using `std::mem::transmute` here to extend the lifetime of `mmap`. As the `stdlib` docs note:


"This is advanced, very unsafe Rust!"

```
1 impl MemmapMessageStream {
2     pub fn new_stream<P: AsRef<Path>>(p: P) -> Result<Pin<Box<Self>>> {
3         let file = File::open(p)?;
4         let mmap = unsafe { Mmap::map(&file) };
5
6         // Create a struct instance without the message_stream first
7         let mut boxed = Box::new(Self {
8             mmap,
9             // Create with dummy message stream, we'll initialize it properly below
10            message_stream: unsafe { std::mem::transmute(MessageStream::new(&[])?) },
11            _pin: PhantomPinned,
12        });
13
14        // Now we can create the message stream with a reference to the pinned mmap
15        let message_stream = unsafe {
16            // This transmute is safe because the mmap will live as long as the struct
17            // and we ensure the struct can't be moved via PhantomPinned
18            std::mem::transmute(MessageStream::new(&boxed.mmap)?)
19        };
20        boxed.message_stream = message_stream;
21
22        // Pin the box
23        Ok(Box::into_pin(boxed))
24    }
25
26    #[inline(never)]
27    pub fn next(self: &mut Pin<Box<Self>>) -> Option<std::result::Result<Message, McapError>> {
28        // Safety: we only access message_stream which is okay to mutate
29        unsafe { self.as_mut().get_unchecked_mut() }
30            .message_stream
31            .next()
32    }
33 }
```

Question: Can we use `std::pin::Pin` instead?

Miri passes without finding anything unsound

```
$cargo miri run --bin vec_backed
```

A terminal window with a dark background and a blue border. It shows the output of a Miri run. The output consists of six lines of text, with line numbers 1 through 6 on the left. Line 1: Finished `dev` profile [unoptimized + debuginfo] target(s) in 0.11s. Line 2: Running `/home/peter/.rustup/toolchains/nightly-x86_64-unknown-linux-gnu/bin/cargo-miri runner target/miri/x86_64-unknown-linux-gnu/debug/vec_back`. Line 3: ed`. Line 4: next: Ok(Message { channel: Channel { id: 0, topic: "foo", schema: None, message_encoding: "", metadata: {} }, sequence: 0, log_time: 0, publish_time: 0, data: [0] }). Line 5: 0, data: [0] }). Line 6: Hello, world!

```
1 Finished `dev` profile [unoptimized + debuginfo] target(s) in 0.11s
2 Running `/home/peter/.rustup/toolchains/nightly-x86_64-unknown-linux-gnu/bin/cargo-miri runner target/miri/x86_64-unknown-linux-gnu/debug/vec_back
3 ed`
4 next: Ok(Message { channel: Channel { id: 0, topic: "foo", schema: None, message_encoding: "", metadata: {} }, sequence: 0, log_time: 0, publish_time:
5 0, data: [0] })
6 Hello, world!
```

Question: Can we use `std::pin::Pin` instead?

Are there differences in generated assembly of `next()`?

Using ouroboros

```
1 .section ".text.memmap_message_stream_works::memmap_message_stream::<impl memmap_message_stream_works::memmap_message_stream::ouroboros_impl_memmap_mes
2 sage_stream::MemmapMessageStream>::next","ax",@progbits
3     .p2align        4, 0x90
4 .type    memmap_message_stream_works::memmap_message_stream::<impl memmap_message_stream_works::memmap_message_stream::ouroboros_impl_memmap_message_str
5 eam::MemmapMessageStream>::next,@function
6 memmap_message_stream_works::memmap_message_stream::<impl memmap_message_stream_works::memmap_message_stream::ouroboros_impl_memmap_message_stream::Mem
7 mapMessageStream>::next:
8     .cfi_startproc
9     jmp qword ptr [rip + <mcap::read::MessageStream as core::iter::traits::iterator::Iterator>::next@GOTPCREL]
```

No! They are identical

Using Pin

```
1 .section .text.mmap_backed::memmap_message_stream::MemmapMessageStream::next,"ax",@progbits
2     .p2align        4, 0x90
3 .type    mmap_backed::memmap_message_stream::MemmapMessageStream::next,@function
4 mmap_backed::memmap_message_stream::MemmapMessageStream::next:
5     .cfi_startproc
6     jmp qword ptr [rip + <mcap::read::MessageStream as core::iter::traits::iterator::Iterator>::next@GOTPCREL]
```




lessons

When does it make sense to use each?

Generally ouroboros has less sharp edges, recommended

Using `std::pin::Pin`

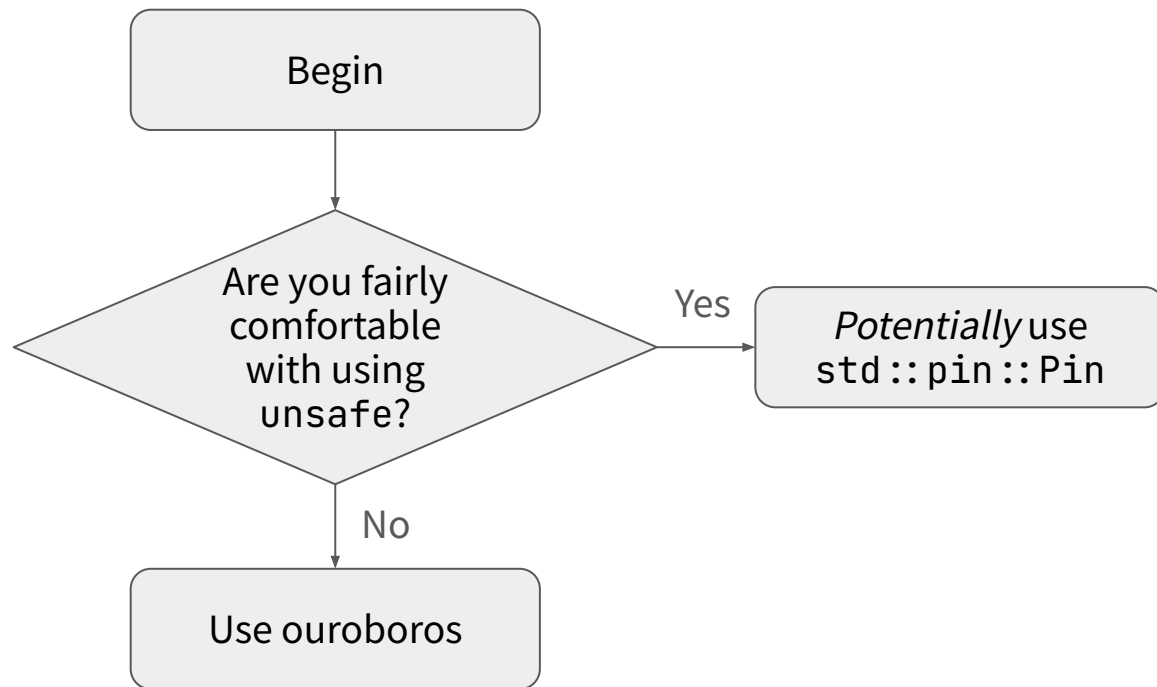
Using `ouroboros`

Pros	Cons
Built into the standard library	Requires unsafe code to implement correctly
Well-documented and understood (used heavily in <code>async/await</code>)	Easy to make mistakes with the unsafe code that could lead to undefined behavior
More flexible if we need to do unsafe (somewhat of a con too) or have specific memory layout requirements	More verbose implementation & have to work with <code>Pin</code> 's API which can be awkward

Pros	Cons
No unsafe code needed in implementation & <code>ouroboros</code> handles all the unsafe details	Additional dependency
More ergonomic API for common use cases	Procedural macros can make compilation slower
Less boilerplate code required	Less flexible for unusual cases

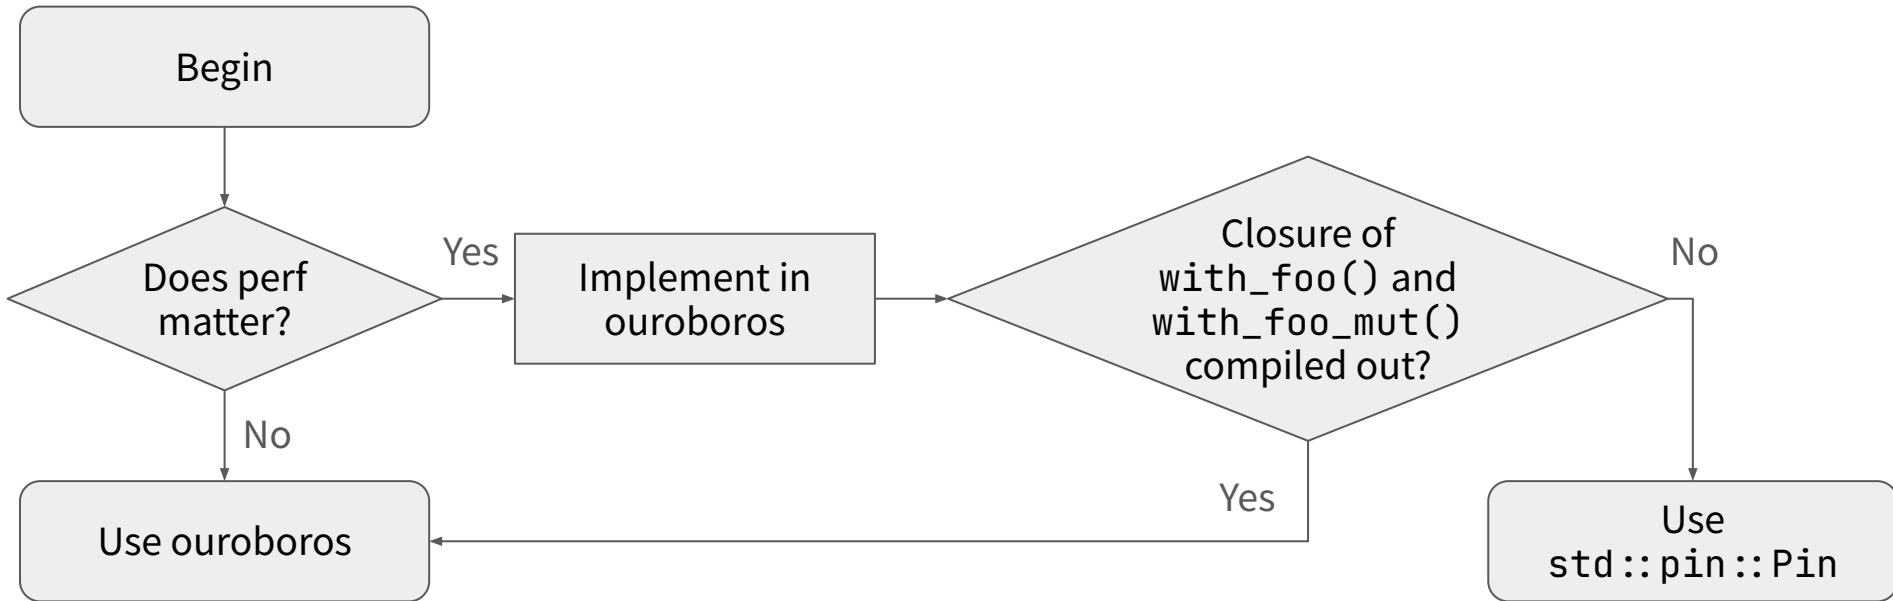
Thinking about ouroboros vs `std::pin::Pin`

Rust expertise



Thinking about ouroboros vs `std::pin::Pin`

Performance





Thank You!



Eclipse uProtocol Website:
<https://uprotocol.org/>

JOIN US ON GITHUB!

Eclipse uProtocol GitHub Project:

<https://github.com/eclipse-uprotocol>

Eclipse SDV Blueprint: Service-to-Signal:

<https://github.com/eclipse-sdv-blueprints/service-to-signal>