

# Rust programming

Module E: Async and Rust for Web

# Last time...

- Trait objects: dynamic dispatch
- Design patterns
- Anti-patterns

*Any questions?*

# In this module

- Introduction Rust `async` programming
- The `Future` type
- The `async` and `await` keywords
- How to run futures using Tokio
- Introduction to the Axum web framework

# Learning objectives

- Understand the mechanics of `async`/`await`` from a high-level point of view
- Understand the reason of Rusts implementation of `async``
- Understand the trade-offs concerning Rust `async`` programming
- Understand the mechanics behind `async`` and `await``
- Work with `Future`s` using the Tokio runtime
- Apply knowledge on `async`` Rust in a web context

# Module E

Async and Rust for Web

# Content overview

- Rust's `async` implementation
- The `Future` trait
- `async` and `await`
- Running `Future`s with Tokio
- Rust for web with Axum

# Async in Rust


# Recap: Concurrency vs. Parallelism

## Concurrency

Interleaves work

1 or more cores

Waiting for events



## Parallelism

Parallelizes work

2 or more cores

Waiting for computation



Today, we're focusing on concurrency: *asynchronous programming*



# What's async?

- Concurrent programming model
- Very suitable for running large number of I/O bound tasks
  - like web servers!
- Look and feel\* of synchronous code through `async`/`await`` syntax

*\*Well, not perfectly. We'll get to that*

# Async vs OS Threads

	Async	OS Threads
Spawning & switching	Cheap	Expensive
Blocking is ok	No	Yes
Usage	I/O bound tasks (web servers)	CPU-bound tasks (Number crunching)
Reuse sync code	No	Yes

What Color is Your Function?

# What `async` looks like in Rust

To get an idea

```
1  /// An async function
2  async fn run() -> anyhow::Result<> {
3      /// Await loading and parsing config file
4      let config = load_config(CONFIG_PATH).await?;
5      /// Await scraping job
6      let data = scrape(&config.urls).await?;
7      data.report();
8      Ok(())
9  }
10
11 fn main() {
12     // Set up a `tokio` runtime with default configurations
13     let runtime = tokio::runtime::Runtime::new().unwrap();
14     // Run a Future to completion
15     runtime.block_on(run());
16 }
```

*Question: What stands out to you? What strikes you as odd?*

# Async in Rust

- Revolve around `Future` trait (~like JS `Promise`, C# `Task`)  
→ `async fn`'s return `Future`'s
- `Future`'s are inert
- `async` is zero-cost
- No built-in runtime
- Single- or multithreaded execution
- Can be mixed with other concurrency models
- Relatively new and lacks some features and nice diagnostics

# State of the `async`` art

What you can expect doing `async`` Rust

- Blazingly fast applications
- More interaction with advanced language features
- Compatibility issues (re: colored functions)
- Faster evolving ecosystem
- `async fn`` in Traits not stable

*But still a work in progress*

# Support of `async`

- Fundamental types and traits are in `std`
- `async`/`await` are native to the language
- Utilities/extensions in `futures` crate
- Async runtimes are third party

Example runtimes: `async-std`, `tokio`, `smol`

# The `Future` trait

Foundation of async

# A `VerySimpleFuture`

```
1 trait VerySimpleFuture {
2     type Output;
3     /// Do work and check if task is completed.
4     /// Returns [Poll::Ready], containing the
5     /// `Output` if task is ready,
6     /// [Poll::Pending] if not
7     fn poll(&mut self) -> Poll<Self::Output>;
8 }
9
10 enum Poll<T> {
11     Ready(T),
12     Pending,
13 }
```

```
1 struct VerySimpleAlarm {
2     alarm_time: Instant,
3 }
4
5 impl VerySimpleFuture for VerySimpleAlarm {
6     type Output = ();
7
8     fn poll(&mut self) -> Poll<()> {
9         if Instant::now() >= self.alarm_time {
10             Poll::Ready(())
11         } else {
12             Poll::Pending
13         }
14     }
15 }
```



# Executing

## `VerySimpleFuture``

```
1  fn main() {
2      let mut first_alarm = VerySimpleAlarm {
3          alarm_time: Instant::now()
4              + Duration::from_secs(3)
5      };
6      let mut snooze_alarm = VerySimpleAlarm {
7          alarm_time: Instant::now()
8              + Duration::from_secs(5)
9      };
10
11     loop {
12         if let Poll::Ready(_) = first_alarm.poll() {
13             println!("Beep beep beep");
14         }
15         if let Poll::Ready(_) = snooze_alarm.poll() {
16             println!("You're late for work!");
17         }
18     }
19 }
```

```
1  [pause...]
2  Beep beep beep
3  Beep beep beep
4  [... a few moments later...]
5  You're late for work!
6  Beep beep beep
7  You're late for work!
8  Beep beep beep
9  You're late for work!
10 [...ad infinitum]
```

It works!

*Question: How can `VerySimpleFuture`` be improved?*

# Limitation of `VerySimpleAlarm`

- Busy waiting
- How to signal the executor the future is *actually* ready to be polled?

## Introduce a Waker

General idea:

- Run some callback to notify executor
- Have executor implement some job queue

# A `SimpleFuture`

```
1  trait SimpleFuture {
2      type Output;
3
4      fn poll(&self, wake: fn()) -> Poll<Self::Output>;
5  }
6
7  pub struct SocketRead<'a> {
8      socket: &'a Socket,
9  }
10
11  impl SimpleFuture for SocketRead<'_> {
12      type Output = Vec<u8>;
13
14      fn poll(&mut self, wake: fn()) -> Poll<Self::Output> {
15          if self.socket.has_data_to_read() {           // <-- Does syscall
16              Poll::Ready(self.socket.read_buf())
17          } else {
18              self.socket.set_readable_callback(wake); // <-- Does syscall
19              Poll::Pending
20          }
21      }
22  }
```

*Adapted from [Asynchronous programming in Rust](#)*

# Joining `SimpleFuture`s

```
1 pub struct Join<FutureA, FutureB> {
2     a: Option<FutureA>,
3     b: Option<FutureB>,
4 }
5
6 impl<FutureA, FutureB> SimpleFuture
7     for Join<FutureA, FutureB>
8 where
9     FutureA: SimpleFuture<Output = ()>,
10    FutureB: SimpleFuture<Output = ()>,
11 {
12     type Output = ();
```

*Adapted from Asynchronous programming in Rust*

```
1 fn poll(&mut self, wake: fn())
2     -> Poll<Self::Output>
3 {
4     if let Some(a) = &mut self.a {
5         if let Poll::Ready(()) = a.poll(wake) {
6             self.a.take(); // Drop future A
7         }
8     }
9     if let Some(b) = &mut self.b {
10        if let Poll::Ready(()) = b.poll(wake) {
11            self.b.take(); // Drop future B
12        }
13    }
14    if self.a.is_none() && self.b.is_none() {
15        Poll::Ready(()) // Both futures dropped
16    } else {
17        Poll::Pending // A future is pending
18    }
19 }
20 }
```

# And then...

```
1  pub struct AndThenFut<FutureA, FutureB> {
2      first: Option<FutureA>,
3      second: FutureB,
4  }
5
6  impl<FutureA, FutureB> SimpleFuture for AndThenFut<FutureA, FutureB>
7  where
8      FutureA: SimpleFuture<Output = ()>,
9      FutureB: SimpleFuture<Output = ()>,
10 {
11     type Output = ();
12     fn poll(&mut self, wake: fn()) -> Poll<Self::Output> {
13         if let Some(first) = &mut self.first {
14             match first.poll(wake) {
15                 Poll::Ready(()) => self.first.take(),
16                 Poll::Pending => return Poll::Pending,
17             };
18         }
19         self.second.poll(wake)
20     }
21 }
```

*Adapted from Asynchronous programming in Rust*

# `SimpleFuture` takeaways

- Composing `SimpleFuture`'s requires no heap allocations
- Composing `SimpleFuture`'s requires no deeply nested callbacks

# The `Future` is now!

```
1 pub trait Future {
2     type Output;
3
4     fn poll(self: Pin<&mut Self>, cx: &mut Context<'_>) -> Poll<Self::Output>;
5 }
```

*Question: what stands out to you?*

- `&mut self` → `Pin<&mut Self>`: makes `Self` immovable
- `wake: fn()` → `cx: &mut Context<'_>`: contains a `Waker`

*More on `Pin<&mut Self>` in the [Rust async book](#)*

``async`` and ``await``



# Expanding `async`

"`Futures` are cool, but why didn't I see them in the web scraper example?"

`async fn`'s and `async` blocks are syntactic sugar generating `Future`'s

```
1  async fn foo() -> u8 { 5 }
```

is equivalent to:

```
1  fn foo() -> impl Future<Output=u8> {  
2      async {  
3          5  
4      }  
5  }
```

which is equivalent to:

```
1  fn foo() -> impl Future<Output=u8> {  
2      /// Create a future that is immediately ready  
with a value.  
3      futures::future::ready(5)  
4  }
```

# Expanding `async` and `await`

```
1 let fut_one = /* ... */;
2 let fut_two = /* ... */;
3 async move {           // <-- generated Future takes ownership of referenced variables
4     fut_one.await;
5     fut_two.await;
6 }
```

Generates an opaque type implementing `Future`:

```
1 struct AsyncFuture {
2     fut_one: FutOne,
3     fut_two: FutTwo,
4     state: State,
5 }
6 enum State {
7     AwaitingFutOne,
8     AwaitingFutTwo,
9     Done,
10 }
```

*This and the following is not the actually generated code, but it's a good mental model*

## Expanding `async` and `await` (2)

```
1  impl Future for AsyncFuture {
2      type Output = ();
3
4      fn poll(mut self: Pin<&mut Self>, cx: &mut Context<'_>) -> Poll<()> {
5          loop {
6              match self.state {
7                  State::AwaitingFutOne => match self.fut_one.poll(/* - snip - */) {
8                      Poll::Ready(()) => self.state = State::AwaitingFutTwo,
9                      Poll::Pending => return Poll::Pending,
10                 }
11                 State::AwaitingFutTwo => match self.fut_two.poll(/* - snip - */) {
12                     Poll::Ready(()) => self.state = State::Done,
13                     Poll::Pending => return Poll::Pending,
14                 }
15                 State::Done => return Poll::Ready(()),
16             }
17         }
18     }
19 }
```

Kind of looks like `AndThenFut`!

*Adapted from Asynchronous programming in Rust*

# `async` / `await` expansion takeaways

- Rust generates state machines out of `async` blocks that implement `Future`
- You can `await` `Future`'s
- Every `await` point introduces a new state
- Generated code may become very complex, but original is easy to follow

Running `Future`s

# What's an `async` Runtime do?

- Spawn `Future`'s
- Keep track of pending `Future`'s
- Call `Future::poll` on each `Future` that can make progress
- Poll `Future`'s on `Waker::wake` calls

Nice to have:

- Poll `Future`'s on multiple threads
- Abstract over I/O

*Crates depending on different runtime I/O abstractions be incompatible!*

# Many Runtime flavors

- `smol`: Small
- `async-std`: API resembles `std`
- `tokio`: Bery commonly used
- `embassy`: Embedded
- Create your own?

*Note: crates may depend on a specific runtime!*

# Showcase: Tokio

```
1  /// Set up a tokio Runtime and spawn the Future returned by `main`
2  #[tokio::main]
3  async fn main() {
4      do_stuff.await();
5  }
```

It does stuff!



# A simple TCP server

```
1 use tokio::net::{TcpListener, TcpStream};
2
3 /// Read a line, and reply with that line!
4 async fn handle_connection(socket: TcpStream) -> anyhow::Result<> {
5     let mut stream = BufReader::new(socket);
6     let mut name = String::new();
7     stream.read_line(&mut name).await?;
8
9     stream.write_all(format!("Hello, {name}!").as_bytes()).await?;
10    Ok(())
11 }
12
13 #[tokio::main]
14 async fn main() -> Result<> {
15     let listener = TcpListener::bind("127.0.0.1:6379").await.unwrap();
16
17     loop {
18         // The second item contains the IP and port of the new connection.
19         let (socket, _) = listener.accept().await.unwrap();
20         handle_connection(socket).await?;
21     }
22 }
```

# It works!

```
1 $ echo -e Ferris | nc localhost 6379
2 Hello Ferris![]
```

*Question: But does it scale?*

Nope! Only one request at a time!

# Spawning tasks is cheap!

```
1  #[tokio::main]
2  async fn main() -> Result<()> {
3      let listener = TcpListener::bind("127.0.0.1:6379").await.unwrap();
4
5      loop {
6          // The second item contains the IP and port of the new connection.
7          let (socket, _) = listener.accept().await.unwrap();
8          handle_connection(socket).await?;
9      }
10 }
```

becomes:

```
1  async fn main() -> Result<()> {
2      loop {
3          // The second item contains the IP and port of the new connection.
4          let (socket, _) = listener.accept().await.unwrap();
5          tokio::task::spawn(async {
6              handle_connection(socket).await?;
7              Ok::<_, anyhow::Error>(( ))
8          });
9      }
10 }
```

Rust for web

# Are we web yet?

---

- "Yes! And it's freaking fast!"
- Several web frameworks exist
  - `rocket``
  - `actix-web``
  - `warp``
  - `axum``
  - ...lots more
- Several DB drivers and ORMs
- Much more!

*Tip: have a look if you want to do web stuff in your final project*

# Axum demo: setting up server

```
1 use axum::{
2     extract::{Path, State},
3     response::Html,
4     routing::get,
5     Router,
6 };
7 use std::net::SocketAddr;
8
9 #[tokio::main]
10 async fn main() {
11     // set up shared, mutable state.
12     let app_state = Arc::new(Mutex::new(Vec::new()));
13     // build our application with a route
14     let app = Router::new()
15         .route("/:name", get(handler))
16         .with_state(app_state);
17     // run it
18     let addr = SocketAddr::from(([127, 0, 0, 1], 3000));
19     println!("listening on {}", addr);
20     axum::Server::bind(&addr)
21         .serve(app.into_make_service())
22         .await
23         .unwrap();
24 }
```

# Axum demo: request handler

```
1  /// A very long type name warrants a type alias
2  type AppState = State<Arc<Mutex<Vec<String>>>>;
3
4  async fn handler(
5      Path(name): Path<String>,
6      State(past_names): State<AppState>,
7  ) -> Html<String> {
8      let mut response = format!("<h1>Hello, {name}</h1>");
9
10     // Of course, locking here is not very fast
11     let mut past_names = past_names.lock().await;
12
13     if !past_names.is_empty() {
14         response += "<h2>Names we saw earlier:</h2>";
15         past_names
16             .iter()
17             .for_each(|name| response += &format!("<p>{name}</p>"))
18     }
19
20     past_names.push(name);
21
22     Html(response)
23 }
```

# Summary

- Async in Rust
- The `Future` trait
- `async` and `await` expansion
- Running `Future`s
- Rust for web



# Tutorial time!

- Exercises E in 101-rs.tweede.golf
- If you haven't done so yet: project proposal deadline is today!
- Reach out if you want to discuss your proposal or need ideas!
- We'll provide feedback on proposals in coming days

*Don't forget to `git pull`!*



Bonus section: ``Pin``

# `async` and lifetime elision

`async fn`s which accept references, return a `Future` bound by argument lifetime:

```
1  async fn foo(x: &u8) -> u8 { *x }
```

is equivalent to:

```
1  fn foo_expanded<'a>(x: &'a u8) -> impl Future<Output = u8> + 'a {  
2      async move { *x }  
3  }
```

- `async move` takes ownership of any variables it references
  - i.e. `x`, which itself is a *reference*
- The returned `impl Future` internally holds the references
- The returned `impl Future` must be `await`ed within `'a`

# Self-referential structs

Consider:

```
1  async {
2      let mut x = [0; 128];           // <--
3      let read_into_buf_fut = read_into_buf(&mut x); // <-- Create future
4      read_into_buf_fut.await;       // <-- `await` future
5      println!("{:?}", x);           //
6  }
```

which becomes

```
1  struct ReadIntoBuf<'a> {
2      buf: &'a mut [u8], // <-- reference to `Async::Future.x` below
3  }
4
5  struct AsyncFuture {
6      x: [u8; 128],       // <-- referent
7      read_into_buf_fut: ReadIntoBuf<'what_lifetime?>,
8  }
```

\*Question: what happens when `AsyncFuture` is moved?

## `Pin<T>`

- Wraps pointer types
- Guarantees values can't be moved (unless `T: Unpin`) using type system

More in [Asynchronous Programming in Rust](#), and [docs on `Pin`](#)