

Software Systems

An introduction to Software Systems
and Concurrency in Rust

Vivian Roest
Jonathan Dönszelmann

Delft University of Technology, The Netherlands

November 16, 2024

Today's Lecture

Course Information


- Course overview
- Staff
- Deadlines and Grading
- Group work

Mutability and Concurrency

- Globals
- Interior Mutability
- Send + Sync
- Concurrency

- Continuation of Software Fundamentals^a
- Two parts
 - 1 Rust for embedded systems
 - 2 Model based software development,

^aThe homologation course for Electrical Engineers, that was focused on learning programming (in Rust). If you haven't done this course: there are notes available on the course website and at <https://cese.ewi.tudelft.nl>



<https://marabos.nl/atomics/>

Staff

- Part 1:
 - Vivian Roest
 - George Hellouin de Menibus (Head TA)
 - & TAs
- Part 2:
 - Rosilde Corvino
 - Guohao Lan
 - & TAs

Deadlines

- Wednesday Week 3 (27 November)
Assignment 1 about concurrency
- Wednesday Week 4 (4 December)
Assignment 2 about performance
- Wednesday Week 6 (18 December)
Assignment 3 about embedded development

Programs, processes and threads

Question:

What's the difference between a program, a process and a thread?

Programs, processes and threads

Question:

What's the difference between a program, a process and a thread?

Question:

Is multithreading possible on a single core system?

Programs, processes and threads

- Programs: Binaries and Scripts
- Process: An instantiation of a program
- Subprocess: A copy of a process without shared memory
- Thread: Subprocess with shared memory

Concurrency vs Parallelism

Question:

Is multithreading possible on a single core system?

- Parallelism: two things are happening at exactly the same time
- Concurrency: two things happen intertwined.
- Multithreading doesn't need to be parallel, it can be concurrent

Scheduler

- More threads/processes than cores
- Concurrency: the illusion of parallelism
- Priorities and niceness
- Generally:
 - Processes get time slices
 - Processes can yield their remaining time
 - The scheduler can cooperate with needs of processes

Scheduler

```
1  std::thread::yield_now();
2
3  // approximately the same as
4  std::thread::sleep(Duration::from_secs(0));
5
```

- Make sure another thread is temporarily allowed to execute

Blocking

Question:

What does **blocking** mean in the context of concurrent execution?

Blocking

- When waiting for a lock to unlock
- Another thread is allowed to run
- The blocking thread is restarted when it can make progress again
- More efficient than waiting in an infinite loop

Memory Safety

- Buffer Overflow
- Use After Free
- Double Free
- Race Conditions
- Null Pointer Dereference

Why does Rust work like it does?

Concurrency bugs in C

```
1  int v = 0;
2
3  void count(int* delta) {
4      for (int i = 0; i < 100000; i++) v += *delta;
5  }
6
7  int main() {
8      thrd_t t1, t2;
9      int d1 = 1, d2 = -1;
10
11     thrd_create(&t1, count, &d1);
12     thrd_create(&t2, count, &d2);
13
14     thrd_join(t1, NULL);
15     thrd_join(t2, NULL);
16
17     printf("%d\n", v);
18 }
19
```

Aliasing

- One memory location
- Accessed through multiple pointers

```
fn update(a: &mut u64) {  
    *a += 1;  
}
```

```
1  update:  
2      push    rbp  
3      mov     rbp, rsp  
4      mov     eax, DWORD PTR [rip+0x2f18]  
5      add     eax, 0x1  
6      mov     DWORD PTR [rip+0x2f0f], eax  
7      mov     eax, 0x0  
8      pop     rbp  
9      ret  
10
```


Mutexes

```
1  int v = 0, n = 100000;
2  mtx_t m;
3
4  void count(int* delta) {
5      for (int i=0; i<n; i++) {
6          mtx_lock(&m);
7          v += *delta;
8          mtx_unlock(&m);
9      }
10 }
11
12 int main() {
13     mtx_init(&m, mtx_plain);
14
15     // spawn and stop threads
16
17     printf("%d\n", v);
18 }
19
```

Atomics

```
1  #include<stdatomic.h>
2
3  _Atomic int v = 0;
4  int n = 100000;
5
6  void count(int* delta) {
7      for (int i=0; i<n; i++) {
8          atomic_fetch_add_explicit(
9              &v,
10             *(int *)delta,
11             memory_order_relaxed
12         );
13     }
14 }
15
16 int main() {
17     // spawn and stop threads
18     printf("%d\n", v);
19 }
20
```

Ordering

Atomics require ordering¹

```
1  pub enum Ordering {
2      Relaxed,
3      Release,
4      Acquire,
5      AcqRel,
6      SeqCst, // Always correct
7  }
8
```

¹More info: <https://marabos.nl/atomics/memory-ordering.html>

Concurrency

```
1  use std::thread::spawn;
2  static v: i32 = 0;
3
4  fn count(delta: i32) {
5      for _ in 0..100_000 {
6          v += delta;
7      }
8  }
9
10 fn main() {
11     let t1 = spawn(|| count(1));
12     let t2 = spawn(|| count(-1));
13
14     t1.join().unwrap();
15     t2.join().unwrap();
16
17     println!("{}", v);
18 }
```

Concurrency

```
1  use std::thread::spawn;
2  static mut v: i32 = 0;
3
4  fn count(delta: i32) {
5      for _ in 0..100_000 {
6          v += delta;
7      }
8  }
9
10 fn main() {
11     let t1 = spawn(|| count(1));
12     let t2 = spawn(|| count(-1));
13
14     t1.join().unwrap();
15     t2.join().unwrap();
16
17     println!("{}", v);
18 }
```

Concurrency

```
error[E0133]: use of mutable static is unsafe and requires unsafe
↳ function or block
--> src/main.rs:7:9
   |
7  |         v += delta;
   |         ~~~~~ use of mutable static
   |
= note: mutable statics can be mutated by multiple threads:
aliasing violations or data races will cause undefined behaviour
```

Intermezzo: Unsafe

- Relaxes some of Rust's strict guarantees
- Puts the programmer in charge of creating **sound** programs.
- More in Lecture 3.

Concurrency

```
1  static mut v: i32 = 0;
2
3  fn count(delta: i32) {
4      for _ in 0..100_000 {
5          unsafe{v += delta};
6      }
7  }
8
9  fn main() {
10     let t1 = thread::spawn(|| count(1));
11     let t2 = thread::spawn(|| count(-1));
12
13     t1.join().unwrap();
14     t2.join().unwrap();
15
16     println!("{}", unsafe{v});
17 }
```

Rust Playground

Data Sharing Rules

- Sharing data immutably is okay
- Moving values to other threads is okay
- Sharing data mutably is not okay.

Notice that these rules are the same as those enforced by the borrow checker!

Mutex as a Monad

```
1  int v = 0, n = 100000;
2  mtx_t m;
3
4  void count(int* delta) {
5      for (int i=0; i<n; i++) {
6          mtx_lock(&m);
7          v += *delta;
8          mtx_unlock(&m);
9      }
10 }
11
12 int main() {
13     mtx_init(&m, mtx_plain);
14
15     // spawn and stop threads
16
17     printf("%d\n", v);
18 }
19
```

Mutex as a Monad

Mutexes work like options, they wrap the value. Providing added context to the data, guarding it.

```
1  static v: Mutex<i32> = Mutex::new(0);
2
3  fn count(delta: i32) {
4      for _ in 0..100_000 {
5          *v.lock() += delta;
6      }
7  }
```

Mutex as a Monad

What does a Mutex look like?

```
pub struct Mutex<T: ?Sized> {  
    inner: sys::MovableMutex,  
    poison: poison::Flag,  
    data: UnsafeCell<T>,  
}
```

- UnsafeCell: Interior mutability
- Allows mutations even when not mutable
- Unsafe to use, like mutable statics
- Mutex has OS help

Removing globals

What would happen if we made the value a local variable:

```
1  fn count(delta: i32, v: &Mutex<i32>) {
2      for i in 0..100_000 {
3          *v.lock() += delta
4      }
5  }
6
7  fn main() {
8      let v = Mutex::new(0);
9
10     let t1 = thread::spawn(|| count(1, &v));
11     let t2 = thread::spawn(|| count(-1, &v));
12
13     t1.join().unwrap();
14     t2.join().unwrap();
15
16     println!("{}", v.lock());
17 }
```

Removing globals

```
error[E0373]: closure may outlive the current function, but it
  → borrows `v`, which is owned by the current function
--> src/main.rs:14:28
|
14 |     let t1 = thread::spawn(|| count(1, &v));
    |                               ^^           - `v` is borrowed here
    |                               |
    |                               may outlive borrowed value `v`
note: function requires argument type to outlive `static`
--> src/main.rs:14:14
|
14 |     let t1 = thread::spawn(|| count(1, &v));
    |                               ~~~~~~
```

Removing globals

- The spawned threads *could* run for longer than the main thread
- v is deallocated at the end of main
- the threads may need v for longer

Removing globals

- The spawned threads *could* run for longer than the main thread
- v is deallocated at the end of main
- the threads may need v for longer
- Solution: Throw v on the heap so it could live longer!

Boxing it up

```
1  fn count(delta: i32, v: Box<Mutex<i32>>) {
2      for i in 0..100_000 {
3          *v.lock() += delta
4      }
5  }
6
7  fn main() {
8      let v = Box::new(Mutex::new(0));
9
10     let t1 = thread::spawn(|| count(1, v));
11     let t2 = thread::spawn(|| count(-1, v));
12
13     t1.join().unwrap();
14     t2.join().unwrap();
15
16     println!("{}", v.lock());
17 }
18
```

Boxing it up

We try to access the same Box at different locations.

```
error[E0382]: use of moved value: `v`
12 |     let v = Box::new(Mutex::new(0));
    |         - move occurs because `v` has type `Box<_>`,
    |         which does not implement the `Copy` trait
13 |
14 |     let t1 = thread::spawn(|| count(1, v));
    |                                --          - variable moved
    |                                |          due to use in closure
    |                                |
    |                                value moved into closure here
15 |     let t2 = thread::spawn(|| count(-1, v));
    |                                ^^          - use occurs
    |                                |          due to use in closure
    |                                |
    |                                value used here after move
```

So when do we deallocate?

Reference Counting

- Keep track of the number of owners
- When it reaches 0 → deallocate
- When we clone the reference, increment the reference count
- Rc doesn't implement Copy like other references, because Copying them is *not* trivial
- Like C++ `std::shared_ptr`

Reference Counting

```
1  fn create_vectors() -> (Rc<Vec<i32>>, Rc<Vec<i32>>) {
2      // one vector on the heap. rc=1
3      let a = Rc::new(vec![1, 2, 3]);
4
5      // two *extra* references to it. rc=3
6      let ref_1 = a.clone(); // doesn't clone the vec! only the ref!
7      let ref_2 = a.clone(); // doesn't clone the vec! only the ref!
8
9      (ref_1, ref_2) // return both. rc=2
10 }
11
12 fn main() {
13     let (a, b) = create_vectors(); // Both are the same vector
14     println!("{}", a);
15     println!("{}", b);
16     // rc finally drops to 0
17 }
```

Rc is clonable even if the internal value is not.

Reference Counting

```
error[E0277]: `Rc<Mutex<i32>>` cannot be shared between threads safely
--> src/main.rs:15:14
   |
15 |     let t1 = thread::spawn(|| count(1, v.clone()));
   |                        ~~~~~ `Rc<Mutex<i32>>` cannot be shared
   |                        between threads safely
   |
= help: the trait `Sync` is not implemented for `Rc<Mutex<i32>>`
= note: required because of the requirements on the
       impl of `Send` for `&Rc<Mutex<i32>>`
```

```
1  pub fn spawn<F, T>(f: F) -> JoinHandle<T>
2      where
3          F: (FnOnce() -> T) + Send + 'static,
4          T: Send + 'static {}
5
```

Sharing and Reference Counting

Question:

Why can't we share an `Rc` between two threads?

Sharing and Reference Counting

Question:

Why can't we share an `Rc` between two threads?

- Keeping track of ownership means updating the reference count
- Updating the reference count needs mutability (on clone and drop)
- Within one thread that's safe

So apparently, it is not safe to send or share some types between threads.

T: Send + Sync

Send: It is safe to send T to another thread

Sync: It is safe to share a T with another thread (T is Sync iff &T is Send)

Reference Counting, Send and Sync

```
error[E0277]: `Rc<Mutex<i32>>` cannot be shared between threads safely
  --> src/main.rs:15:14
     |
15   |     let t1 = thread::spawn(|| count(1, v.clone()));
     |                        ~~~~~ `Rc<Mutex<i32>>` cannot be shared
     |                        between threads safely
     |
= help: the trait `Sync` is not implemented for `Rc<Mutex<i32>>`
= note: required because of the requirements on the
       impl of `Send` for `&Rc<Mutex<i32>>`
```

```
1  pub fn spawn<F, T>(f: F) -> JoinHandle<T>
2      where
3          F: (FnOnce() -> T) + Send + 'static,
4          T: Send + 'static {}
5
```

Reference Counting, Send and Sync

```
1  impl<T: ?Sized> !Send for Rc<T> {}  
2  impl<T: ?Sized> !Sync for Rc<T> {}
```

Atomic Reference Counter (Arc)

- We could use a Mutex again
- Atomics
- `Arc<T>`: `Send` + `Sync`
- So we can use it to share references between threads

Fixing the code with Arcs

```
1  fn count(delta: i32, v: Arc<Mutex<i32>>) {
2      for i in 0..100_000 {
3          *v.lock() += delta
4      }
5  }
6
7  fn main() {
8      let v = Arc::new(Mutex::new(0));
9      let (v1, v2) = (v.clone(), v.clone());
10
11     let t1 = thread::spawn(|| count(1, v1));
12     let t2 = thread::spawn(|| count(-1, v2));
13
14     t1.join().unwrap();
15     t2.join().unwrap();
16
17     println!("{}", v.lock());
18 }
```

Concurrency in other languages

- Python: Doesn't really have it
- Java: Every object has a “thin lock”
- C or C++: You're responsible
- Go: You're responsible
- Haskell: No mutability
- JavaScript/TypeScript: relies on async/await
- Rust: The Compiler is responsible

Break

- See you back in 15 minutes
- Crash course concurrent programming.

Communicating between threads

① Communicate by sharing memory

- Make memory mutable from multiple threads
- Needs locking to avoid data races
- Can be quite slow when there is lots of communication
- Useful for exchanging large amounts of data
- More complex logic

Communicating between threads



① Communicate by sharing memory

- Make memory mutable from multiple threads
- Needs locking to avoid data races
- Can be quite slow when there is lots of communication
- Useful for exchanging large amounts of data
- More complex logic

② Communicating by sending messages

- No need for explicit locking
- Easy to reason about
- Safe by design: No mutability problems

Communicating between threads



Channels

- Like a queue, first in, first out
- One end in one thread, other end in another thread
- Lock-free insertion (magic / atomics depending on who you ask)
- Receivers **block** and wait when there are no messages



Channels

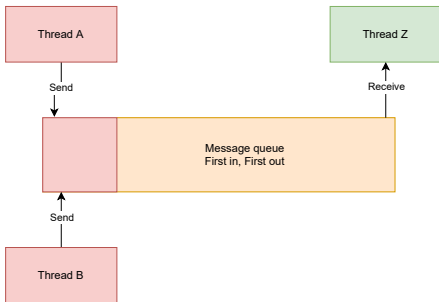
A simple example

```
1 use std::sync::mpsc::channel;
2
3 fn main() {
4     // tx: Sender<i32>, rx: Receiver<i32>
5     let (tx, rx) = channel();
6
7     spawn(move || {
8         while let Ok(i) = rx.recv() {
9             println!("hello, {i}")
10        }
11    });
12
13    tx.send(1).unwrap();
14    tx.send(2).unwrap();
15    tx.send(3).unwrap();
16 }
```

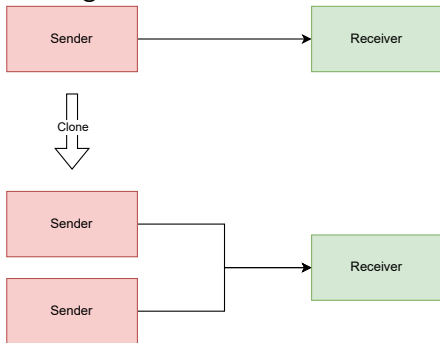
Channels: mpsc

- Multi Producer Single Consumer
- Senders can be cloned, receivers can not

Multiple Producers



Cloning



Channels: Cloning Senders

```
1  fn main() {
2      let (tx, rx) = channel();
3
4      spawn(move || {
5          while let Ok(i) = rx.recv() {
6              println!("hello, {i}")
7          }
8      });
9
10     // clone senders
11     let tx1 = tx.clone();
12     let tx2 = tx.clone();
13
14     spawn(move || for i in 0..5 {tx1.send(i).unwrap()});
15     spawn(move || for i in 5..10 {tx2.send(i).unwrap()});
16 }
```

Channels: multiple receivers?


What are the semantics?

- Send messages to all receivers
- Send messages to one receiver
- Various libraries implement both



Bounded Channels

- By default, channels are unbounded
- Problem: senders can get ahead of receivers
- Solution: Buffer with a certain size:
Bounded channel / sync channel
- Senders can block



Bounded Channels Example

```
1  fn main() {
2      let (tx, rx) = sync_channel(3);
3
4      spawn(move || {
5          while let Ok(i) = rx.recv() {
6              println!("hello, {i}")
7          }
8      });
9
10     // clone senders 40 times
11     for i in 0..40 {
12         let tx_clone = tx.clone();
13         spawn(move || for j in (i*10)..(i*10+10) {
14             tx_clone.send(j).unwrap()
15         });
16     }
17 }
```


Bounded Channels: 0 size

Question:

What happens when the size is 0?

Bounded Channels: 0 size

- Sender blocks until a receiver is ready

```
1  fn main() {
2      let (tx, rx) = sync_channel(0);
3
4      spawn(move || {
5          while let Ok(recv_msg) = rx.recv() {
6              println!("hello, {i}")
7          }
8      });
9
10     // clone senders 40 times
11     for thread_count in 0..40 {
12         let tx_clone = tx.clone();
13         spawn(move || for i in
14     → (thread_count*10)..(thread_count*10+10) {
15             tx_clone.send(i).unwrap()
16         });
17     }
```

Thread Management

The lifecycle of a thread:

```
1  let handle = spawn(|| {...});  
2  
3  ...  
4  
5  handle.join();  
6
```

Thread Management

The lifecycle of multiple threads:

```
1  let mut handles = Vec::new();
2
3  for ... {
4      handles.push(spawn(|| {...}));
5  }
6
7  ...
8
9  for handle in handles {
10     handle.join();
11 }
12
```

Question:

How many threads should we spawn?

Thread Management

The lifecycle of multiple threads:

```
1  let mut handles = Vec::new();
2
3  for ... {
4      handles.push(spawn(|| {...}));
5  }
6
7  ...
8
9  for handle in handles {
10     handle.join();
11 }
12
```

Insight

It is not always trivial to determine the amount of threads to spawn.

Thread Management

```
1  fn task_runner(rx: Receiver<Task>) {
2      while let Some(task) = rx.recv() { task() }
3  }
4
5  let mut txs = Vec::new();
6  for _ in num_cores {
7      let (rx, tx) = channel();
8      spawn(|| task_runner(rx));
9      txs.push(tx);
10 }
11
12 // Execute task(s) on a random thread
13 txs.choose(thread_rng()).send(some_task)
14
```

Note: Task could for example be a function pointer

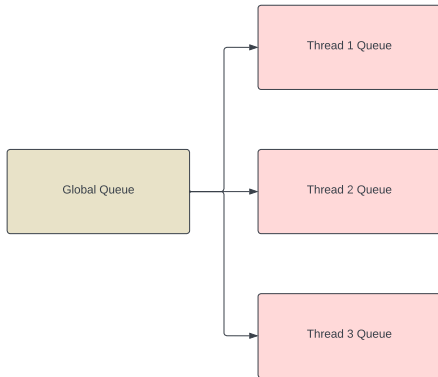
Thread Management

```
1  fn task_runner(rx: Receiver<Task>) {
2      while let Some(task) = rx.recv() { task() }
3  }
4
5  let mut txs = Vec::new();
6  for _ in num_cores {
7      let (rx, tx) = channel();
8      spawn(|| task_runner(rx));
9      txs.push(tx);
10 }
11
12 // Execute task(s) on a random thread
13 txs.choose(thread_rng()).send(some_task);
14
```

Question:

Does this ensure an equal distribution of work?

Thread Management: work stealing



Thread 1 queue and global queue empty? Steal from another thread!

Thread Management: Work Stealing

```
1 use rayon::ParallelIterator;
2
3 let tasks: [Task; 3] = [taska, taskb, taskc];
4 tasks
5     .into_par_iter()
6     .for_each(|x| x())
```

Waiting for IO

Situation:

- ① A webserver is handling many connections.
- ② Each connection has an associated TCP socket
- ③ Every time we receive a packet, we have to process it for 0.1ms
- ④ Every time we send a packet, it takes about 10 seconds to get a response

Question:

How many threads do we need for maximum utilization?

Waiting for IO

But the model of 'tasks' is really useful!

```
1  fn task() {
2      let conn = start_connection();
3      while let Some(packet) = conn.recv() {
4          conn.send(process(packet));
5      }
6  }
7
8  for i in 0..=many {
9      spawn(task);
10 }
11
```

Take scheduling into our own hands!

Don't let the OS schedule threads

- Make our own lightweight 'task'
- Make our own scheduler
- Tasks tell the scheduler what event they're waiting for
- The scheduler wakes up a task when the OS says the event has happened

Examples:

- Java, Elixir: Green Threads
- Go: Goroutines
- Python, Javascript, C++, Rust: Async/IO

Take scheduling into our own hands!

Problem:

How do we preempt a task?

- 'Easy' for interpreted languages
- Go inserts extra runtime checks

Take scheduling into our own hands!

Find the places where we could start doing something else.

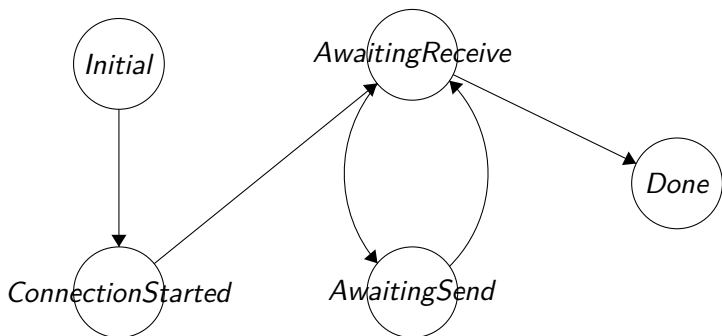
Modelling I/O as a state machine.

```
1  let conn = start_connection();
2  while let Some(packet) = conn.recv() {
3      conn.send(process(packet));
4  }
```

States:

- Initial
- ConnectionStarted{conn}
- AwaitingReceive{conn}
- AwaitingSend{conn, packet}
- Done

Take scheduling into our own hands!




Reading a file

```
1  fn read_file(file: &Path) -> Vec<u8> {
2      let f = File::open(path);
3      let buffer = Vec::new();
4
5      loop {
6          let mut read_buffer = [0; 2048];
7          let num_read = f.read(&mut read_buffer);
8
9          if num_read == 0 { break; }
10         buffer.extend(&read_buffer[0..num_read]);
11     }
12     buffer
13 }
```

Question:


What are the states?

Reading a file




Reading two files simultaneously

Task 1



Task 2



Try to make progress in both: Polling

Asynchronous IO

- <https://man7.org/linux/man-pages/man7/aio.7.html>
- `aio_read` like `read`, but returns immediately
- `aio_error` see the status (error, ok, or in progress) of a read request
- `aio_return` get the return value of `read`, how many bytes were read?

Note: this is *one* option under linux. There is also `io_uring` for Linux, and similar APIs on Windows and MacOS.

Asynchronous IO

```
1  pub async fn read_file(path: &Path) -> Result<Vec<u8>, io::Error> {
2      let mut file = File::open(path).await?;
3      let mut buffer = Vec::new();
4
5      loop {
6          let mut read_buffer = Vec::new();
7          let num_bytes = file.read(&mut read_buffer).await?;
8          if num_bytes == 0 {
9              break;
10         }
11
12         buffer.extend(&read_buffer[0..num_bytes]);
13     }
14
15     Ok(buffer)
16 }
```

Asynchronous IO

In Rust `.await`

- Sets up an (asynchronous) request to the operating system
- Goes to the next state in the state machine
- Yields to poll another state machine
- When we get back to polling this state machine:
 - ask the OS if we made progress
 - continue executing code

Summary

- Achieving fearless concurrency in Rust
- Using channels to your advantage
- Be conscious of your thread count
- Not all cases call for threads
- Async I/O internals and its state machines

Rounding up

- First lab next thursday
- First assignment online today:
 - creating a concurrent `grep` clone

Lifetimes

- Every value has a lifetime
- After the end of a lifetime, a value is dropped
- Some types have a lifetime associated with them

```
1 fn main() {
2     let i = 3; // Lifetime for `i` starts.
3
4     {
5         let borrow1 = &i; // `borrow1` lifetime starts.
6
7         println!("borrow1: {}", borrow1); //
8     } // `borrow1` ends.
9
10
11    {
12        let borrow2 = &i; // `borrow2` lifetime starts.
13
14        println!("borrow2: {}", borrow2); //
15    } // `borrow2` ends.
16
17 } // Lifetime ends.
18
```


Lifetimes

In functions, you can be explicit about lifetimes

```
1 fn print_ref<'a>(x: &'a i32) {  
2     println!("print_ref: x is {}", x);  
3 }  
4
```

In structs, you **have** to be explicit about lifetimes

```
1 struct Example<'a> {  
2     field: &'a i32,  
3 }  
4
```

Lifetimes

Sometimes, logic requires two lifetimes to be the same

```
1  fn longest(a: &str, b: &str) -> &str {
2      if a.len() > b.len() {
3          a
4      } else {
5          b
6      }
7  }
8
```

Lifetimes

Sometimes, logic requires two lifetimes to be the same

```
1  fn longest<'a>(a: &'a str, b: &'a str) -> &'a str {
2      if a.len() > b.len() {
3          a
4      } else {
5          b
6      }
7  }
```