

Rust programming

Module C: concurrency and parallelism

Who am i?

- I'm Folkert
- work on Network Time Protocol and other systems programming things
- work on the Roc compiler (and other low-level shenanigans)

Last time...


- Cargo and dependencies
- Creating a nice API
- Testing and benchmarking
- Setting up your own project

Learning objectives

after this lecture + exercises, you can:

- parallelize a program with Rayon
- work with threads in rust
- reason about exclusive access
- implement a basic Mutex

48 Years of Microprocessor Trend Data




Original data up to the year 2010 collected and plotted by M. Horowitz, F. Labonte, O. Shacham, K. Olukotun, L. Hammond, and C. Batten
New plot and data collected for 2010-2019 by K. Rupp

Concurrency vs. Parallelism

Concurrency

Interleaves work

1 or more cores



Waiting for events

Parallelism

Parallelizes work

2 or more cores



Waiting for computation

Parallelism with Rayon

solving Pleasantly Parallel Problems

TF-IDF

An algorithm for searching in a big collection of text documents

- term frequency–inverse document frequency
- TF: "how often does a word occur in a particular document"
- IDF: "how rare is the word across all documents"

Problem:

- how do we aggregate the results?

TF-IDF in Rayon

```
1 use std::collections::HashMap;
2 use rayon::prelude::*;
3
4 fn document_frequency(documents: &[&str]) -> HashMap<&str, usize> {
5     documents
6         .par_iter()
7         .map(|document| term_occurrence(document))
8         .reduce(HashMap::default, combine_occurrences);
9 }
10
11 /// Map each word in the document to the value 1
12 fn term_occurrence(document: &str) -> HashMap<&str, usize> {
13     todo!()
14 }
15
16 /// combine the counts from maps a and b.
17 fn combine_occurrences<'a>(
18     a: HashMap<&'a str, usize>,
19     b: HashMap<&'a str, usize>,
20 ) -> HashMap<&'a str, usize> {
21     todo!()
22 }
```

Combining results

The `combine_documents` function has several useful properties

- our operation is associative $a \cdot (b \cdot c) = (a \cdot b) \cdot c$
- our operation has a neutral value `HashMap::default()`: $0 \cdot x = x \cdot 0 = x$
- therefore we can split the computation $a \cdot b \cdot c \cdot d = (0 \cdot a \cdot b) \cdot (0 \cdot c \cdot d)$
- an associative operation with a neutral value is called a "monoid"

```
1 // for each word, how often it occurs across all documents
2 documents
3     .par_iter()
4     .map(|document| count_words(document))
5     .reduce(HashMap::default, combine_documents);
```

- this idea means each thread can start accumulating values

Intermezzo: Closures

- Closures are anonymous (unnamed) functions
- they can capture ("close over") values in their scope
- they are first-class values

```
1  fn foo() -> impl Fn(i64, i64) -> i64 {
2      z = 42;
3      |x, y| x + y + z
4  }
5
6  fn bar() -> i64 {
7      // construct the closure
8      let f = foo();
9
10     // evaluate the closure
11     f(1, 2)
12 }
```

- very useful when working with iterators, `Option`` and `Result``.

```
1  let evens: Vec<_> = some_iterator.filter(|x| x % 2 == 0).collect();
```

So far

- Closures are unnamed inline functions
- Rayon makes data-parallel programming in rust extremely convenient

Fearless concurrency

thread-based concurrency in rust

Fearless concurrency

```
1  use std::thread;
2
3  fn main() {
4      thread::spawn(f);
5      thread::spawn(f);
6
7      println!("Hello from the main thread.");
8  }
9
10 fn f() {
11     println!("Hello from another thread!");
12
13     let id = thread::current().id();
14     println!("This is my thread id: {id:?}");
15 }
```

- A process can spawn multiple threads of execution. These run concurrently (and may run in parallel)
- Question: what is the output of this program?

Expected output

maybe

```
1 Hello from another thread!  
2 This is my thread id: ThreadId(411)  
3 Hello from another thread!  
4 This is my thread id: ThreadId(412)  
5 Hello from the main thread.
```

or

```
1 Hello from another thread!  
2 Hello from another thread!  
3 This is my thread id: ThreadId(411)  
4 This is my thread id: ThreadId(412)  
5 Hello from the main thread.
```

Expected output

but most likely

```
1 Hello from the main thread.
```

The process exits when the main thread is done!

- `.join()` can be used to block the main thread until the child is done

```
1 fn main() {  
2     let t1 = thread::spawn(f);  
3     let t2 = thread::spawn(f);  
4  
5     println!("Hello from the main thread.");  
6  
7     t1.join().unwrap();  
8     t2.join().unwrap();  
9 }
```

- `.join()` turns a panic in the thread into an `Err`

Thread lifetime

- a more typical example

```
1  let numbers = Vec::from_iter(0..=1000);
2
3  let t = thread::spawn(move || {
4      let len = numbers.len();
5      let sum = numbers.iter().sum::<usize>();
6      sum / len
7  });
8
9  let average = t.join().unwrap();
10
11 println!("average: {average}");
```

- `numbers` must be `move` d into the closure!

Thread lifetime

- otherwise `numbers`` might be dropped while the thread is still using it!

```
1 let numbers = Vec::from_iter(0..=1000);
2
3 let t = thread::spawn(|| {
4     let len = numbers.len();
5     let sum = numbers.iter().sum::<usize>();
6     sum / len
7 });
8
9 drop(numbers); // compile error: would create a dangling reference
10
11 let average = t.join().unwrap();
12
13 println!("average: {average}");
```

Thread lifetime: make it known

```
1 let numbers = Vec::from_iter(0..=1000);
2
3 let average = thread::scope(|spawner| {
4     spawner.spawn(|| {
5         let len = numbers.len();
6         let sum = numbers.iter().sum::<usize>();
7         sum / len
8     }).join().unwrap()
9 });
10
11 println!("average: {average:?}");
```

- explicitly bound the lifetime with a scope
- threads are always joined at the end of that scope
- makes immutable references just work

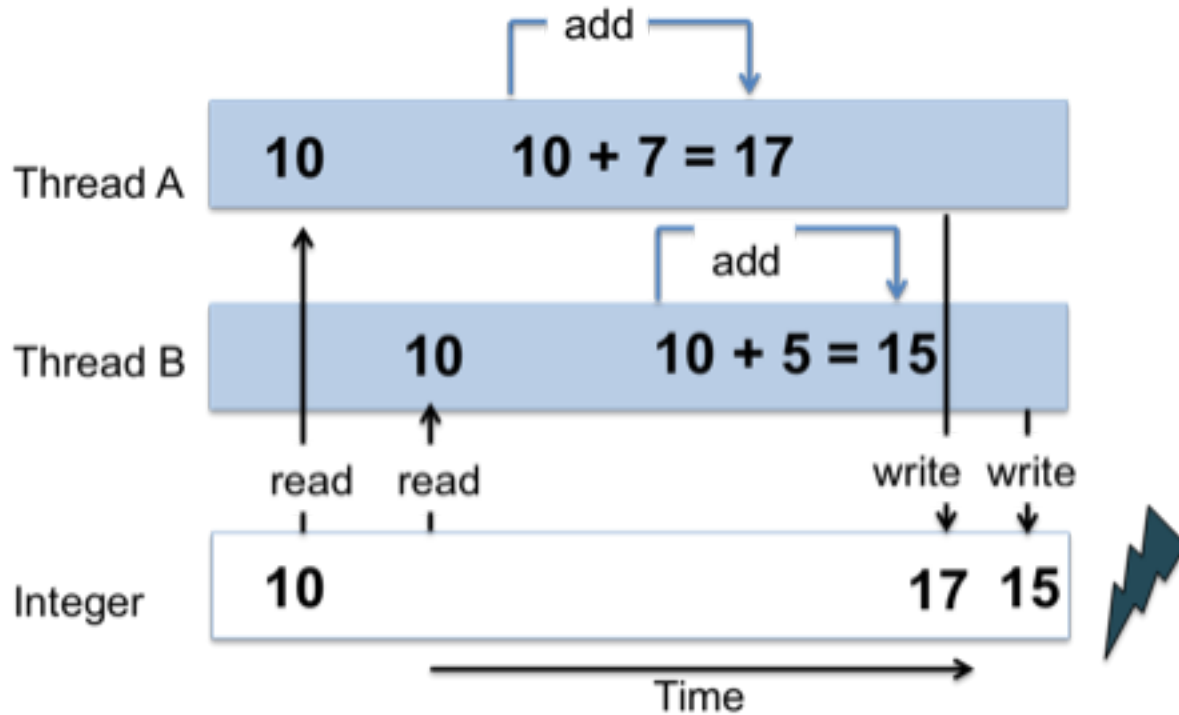
- but mutable borrowing rules still apply:

```
1 let mut count = 0;
2 let counter = &mut count;
3
4 std::thread::scope(|s| {
5     s.spawn(|| { *counter = *counter + 1; });
6     s.spawn(|| { *counter = *counter + 1; });
7 });
```


```
1 error[E0499]: cannot borrow `*counter` as mutable more than once at a time
2 6 |     thread::scope(|s| {
3   |         - has type `&'1 Scope<'1, '1>`
4 7 |         s.spawn(|| { *counter = *counter + 1; });
5   |         -----
6   |         |     |
7   |         |     | first borrow occurs due to use of `*counter` in closure
8   |         |     | first mutable borrow occurs here
9   |         |     | argument requires that `*counter` is borrowed for `1`
10 8 |         s.spawn(|| { *counter = *counter + 1; });
11   |         ^^^ ----- second borrow occurs due to use of `*counter` in closure
12   |         |
13   |         second mutable borrow occurs here
```

Race Conditions

- if multiple mutable borrows were allowed, this could happen ...



Fearless concurrency



- borrowing rules prevent data races & deadlocks
- but also any shared mutable state between threads
- many correct, useful programs are disallowed!

Re-defining references

- `&T`: (possibly) shared reference
- `&mut T`: exclusive reference

for safe mutation, we need exclusive *access*, which we can get in multiple ways:

- we have an exclusive reference to the value
- we own the value (we can exclusively borrow from ourselves)
- access is inherently exclusive (atomic operations)

Atomics

- atomic operations are indivisible, but relatively expensive

```
1 use std::sync::atomic::{AtomicU32, Ordering};
2
3 let foo = AtomicU32::new(0);
4 assert_eq!(foo.fetch_add(10, Ordering::SeqCst), 0);
5 assert_eq!(foo.load(Ordering::SeqCst), 10);
```

- no risk of a race condition: another thread cannot read the value while an atomic operation is ongoing

```
1 pub fn fetch_add(&self, val: u32, order: Ordering) -> u32
```


Mutual Exclusion

- `Mutex` allows mutation of a `T` through a shared `&Mutex<T>` reference

```
1 use std::sync::Mutex;
2 use std::thread;
3
4 fn main() {
5     let n = Mutex::new(String::from("foo"));
6     thread::scope(|s| {
7
8         s.spawn(|| { n.lock().unwrap().push_str("bar"); });
9
10        s.spawn(|| { n.lock().unwrap().push_str("baz"); });
11
12    });
13
14    println!("{}", n.into_inner().unwrap());
15 }
```

- threads lock the mutex, but there is no `unlock` ?!

Sharing ownership between threads

```
1 impl<T> Mutex<T> {
2     pub fn lock<'a>(&'a self) -> LockResult<MutexGuard<'a, T>> {
3         ...
4     }
5 }
```

- Acquires a mutex, blocking the current thread until it is able to do so
- Returns a `PoisonError` if a thread panicked while holding the lock
- Returns a `MutexGuard`, proof to the type checker that we hold the lock
- `MutexGuard<'a, T>` implements `DerefMut<Target = T>`, so we can use it like a mutable reference

```
1 impl<'a, T> DerefMut for MutexGuard<'a, T> {
2     fn deref_mut(&mut self) -> &mut T {
3         // ...
4     }
5 }
```

- dropping the `MutexGuard` unlocks the mutex

Moving ownership between threads

- Some values should never be shared or moved between threads

The `Send` and `Sync` marker traits enforce this:

```
1 pub unsafe auto trait Send { /* no method */ }
2 pub unsafe auto trait Sync { /* no method */ }
```

- `Send`: A type is `Send` if it can be sent to another thread. In other words, if ownership of a value of that type can be transferred to another thread
- `Sync`: A type is `Sync` if it can be shared with another thread. In other words, a type `T` is `Sync` if and only if a shared reference to that type `&T` is `Send`

`Send`

- A type is `Send` if it can be sent to another thread. In other words, if ownership of a value of that type can be transferred to another thread

```
1 impl<T: ?Sized> !Send for MutexGuard<'_, T>  
2 impl<T: ?Sized + Sync> Sync for MutexGuard<'_, T>
```

- On certain OS's, only the thread that locked a mutex may unlock it again!


MPSC: many producer single consumer

```
1 fn main() {
2     let (tx, rx) = std::sync::mpsc::channel();
3
4     std::thread::scope(|s| {
5         for (i, tx) in std::iter::repeat(tx).take(10).enumerate() {
6             s.spawn(move || { tx.send(i).unwrap(); });
7         }
8
9         s.spawn(move || {
10            while let Ok(msg) = rx.recv() {
11                println!("{msg}");
12            }
13        });
14    });
15 }
```

where the `Receiver` is:

```
1 impl<T: Send> Send for Receiver<T>
2 impl<T> !Sync for Receiver<T>
```

Further reading



- read for free at <https://marabos.nl/atomics/>

Summary

- Rayon makes parallel computation easy
- Scoped threads allow borrowing into threads
- Mutation requires exclusive access
- Some data structures guarantee exclusive access (even through a shared reference)
- The borrow checker, ``Send`` and ``Sync`` prevent many common problems