

RustBelt: Securing the Foundations of the Rust Programming Language



Ralf Jung
Jacques-Henri Jourdan
Robert Krebbers
Derek Dreyer

December 18th, 2017
Séminaire GALLIUM

Rust

Mozilla's replacement for C/C++

Systems programming language focusing on safety

- Control over memory allocation & layout
- **Sound type system** with guarantees:
 - Type and memory safety
 - Absence of data races
 - Idea: Prohibit aliased mutable state
 - Using *borrow types* with “lifetimes”
 - First-class functions, polymorphism/generics
 - *Traits* \approx Type classes + associated types



Rust

Mozilla's replacement for C/C++

Systems programming language **focusing on safety**

- Control over memory allocation & layout
- **Sound? type system** with guarantees:
 - Type and memory safety
 - Absence of data races
 - Idea: Prohibit aliased mutable state
 - Using *borrow types* with “lifetimes”
 - First-class functions, polymorphism/generics
 - *Traits* \approx Type classes + associated types



RustBelt: **prove the soundness** of Rust's type system (idealized)

The key challenge

Superficially: **no mutation through aliased pointers**

But this is not always true!

- **Many Rust libraries permit mutation through aliased pointers**
- The safety of this is highly non-obvious because these libraries make use of unsafe features!

The key challenge

Superficially: **no mutation through aliased pointers**

But this is not always true!

- **Many Rust libraries permit mutation through aliased pointers**
- The safety of this is highly non-obvious because these libraries make use of unsafe features!

So why is any of this sound?

Introduction

Overview of Rust

λ_{Rust} : a small idealized Rust

A semantic model for λ_{Rust}

Conclusion

```
let (snd, rcv) = channel();
join(
  move || { // First thread
    // Allocating [b] as Box<i32> (pointer to heap)
    let mut b = Box::new(0);
    *b = 1;

    // Transferring the ownership to the other thread...
    snd.send(b);

  },
  move || { // Second thread
    let b = rcv.recv().unwrap(); // ... that receives it
    println!("{}", *b);         // ... and uses it.
  });
```

```
let (snd, rcv) = channel();
join(
  move || { // First thread
    // Allocating [b] as Box<i32> (pointer to heap)
    let mut b = Box::new(0);
    *b = 1;

    // Transferring the ownership to the other thread...
    snd.send(b);
    *b = 2;    // Error: lost ownership of [b]
               // ==> Prevents data race
  },
  move || { // Second thread
    let b = rcv.recv().unwrap(); // ... that receives it
    println!("{}", *b);         // ... and uses it.
  });
```


Borrowing and lifetimes

```
let mut v = vec![1, 2, 3];
```

```
v[1] = 4;
```

```
v.push(6);  
println!("{:?}", v);
```

Borrowing and lifetimes

```
let mut v = vec![1, 2, 3];  
  
{ let mut inner_ptr = Vec::index_mut(&mut v, 1);  
  
    *inner_ptr = 4; }  
  
v.push(6);  
println!("{:?}", v);
```

Borrowing and lifetimes

```
let mut v = vec![1, 2, 3];

{ let mut inner_ptr = Vec::index_mut(&mut v, 1);
  // Error: can invalidate [inner_ptr]
  v.push(1);
  *inner_ptr = 4; }

v.push(6);
println!("{:?}", v);
```


Borrowing and lifetimes

```
let mut v = vec![1, 2, 3];
```

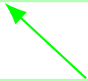
```
{ let mut inner_ptr = Vec::index_mut(&mut v, 1);  
  // Error: can invalidate [inner_ptr]  
  v.push(1);  
  *inner_ptr = 4; }
```

```
v.push(6);  
println!("{:?}", v);
```

We **temporarily** lost ownership of vector `v`



We get back the full ownership of vector `v`



Borrowing and lifetimes

```
let mut
```

```
{ let
```

```
*inn
```

```
v.push
```

```
println!
```

Type of `index_mut`:

```
fn<'a> index_mut(&'a mut Vec<i32>, usize)  
    -> &'a mut i32
```

New pointer type: `&'a mut T`:

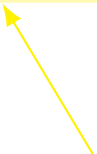
- **mutable borrowed** reference
- valid only for **lifetime** `'a`

Borrowing and lifetimes

```
let mut v = vec![1, 2, 3];
```

```
{ let mut inner_ptr = Vec::index_mut(&mut v, 1);  
  
  *inner_ptr = 4; }
```

```
v.push(6);  
println!("{:?}", v);
```



Lifetime 'a inferred by Rust

Shared borrowing

```
let mut x = 1;  
join (|| println!("{}", &x),  
      || println!("{}", &x));  
x = 2;
```

Shared borrowing

```
let mut x = 1;
join (|| println!("{}", &x),
      || println!("{}", &x));
x = 2;
```

`&x` creates a **shared borrow** of `x`

- Type: `&'a i32`
- Can be copied/shared
- Does not allow mutation

Summing up

- Rust's type system is based on **ownership**
- Three kinds of ownership:
 1. Full ownership: `Vec<T>` (vector), `Box<T>` (pointer to heap)
 2. **Mutable borrowed** reference: `&'a mut T`
 3. **Shared borrowed** reference: `&'a T`
- **Lifetimes** decide when borrows are valid
 - Remark: If `x : &'a (&'b T)`, then `&'b T` is valid during `'a`
 \implies Rust checks `'a \sqsubseteq 'b` ("**Outlives relation**")

Interior mutability

What if we want **shared mutable data structures**?

Rust standard library provides types with **interior mutability**

- Allows mutation under a shared borrow
- Written in Rust **using unsafe features**
- **Safely encapsulated**
 - The library interface restricts mutations

Mutex

Example of interior mutability

```
let m = Mutex::new(1); // m : Mutex<i32>

// We can mutate the integer
// *with a shared borrow* only
join (|| *(&m).lock().unwrap() += 1,
      || *(&m).lock().unwrap() += 1);

// Unique owner: no need to lock
println!("{}", m.into_inner().unwrap())
```

Cell

Example of interior mutability

```
fn incr_cell<'a>(c : &'a Cell<i32>) {  
    let x = c.get();  
    // Can mutate through a shared borrow only  
    c.set(x + 1)  
}  
  
fn main() {  
    let c = Cell::new(0);  
  
    incr_cell(&c);  
}
```

Cell

Example of interior mutability

```
fn incr_cell<'a>(c : &'a Cell<i32>) {  
    let x = c.get();  
    // Can mutate through a shared borrow only  
    c.set(x + 1)  
}
```

```
fn main() {  
    let c = Cell::new(0);  
    // Data race => forbidden  
    join(|| incr_cell(&c), || incr_cell(&c));  
}
```

Cell

Example of interior mutability

```
fn inc  
    le  
    //  
    c  
}  
  
fn mai  
    le  
  
    // Data race => forbidden  
    join(|| incr_cell(&c), || incr_cell(&c));  
}
```

[...]
'std::cell::Cell<i32>' cannot be shared
between threads safely
[...]
The trait 'std::marker::Sync' is not
implemented for 'std::cell::Cell<i32>'
[...]

Cell

Example of interior mutability

```
fn inc
```

```
le
```

```
//
```

```
c.
```

```
}
```

```
fn mai
```

```
le
```

```
//
```

```
join(|| incr_cell(&c), || incr_cell(&c));
```

```
}
```

T shared across threads \implies T : Sync

- Automatic for basic types (structs, i32...)
- Ex: Cell<T> **not** Sync

T moved across threads \implies T : Send

- Automatic for basic types (structs, i32...)
- Ex: Rc<T> (ref. counted smart pointer) **not** Send
- T : Sync \iff &'a T : Send

Protocols

A shared borrow **establishes a sharing protocol**:

- `&'a i32`
 - \implies **Read-only**
 - Safety: trivial

- `&'a Mutex<i32>`
 - \implies Read-write **by taking the lock**
 - Safety: ensured by proper synchronization

- `&'a Cell<i32>`
 - \implies Read-Write **via** `get()` **and** `set(...)`
 - Safety: single threaded (no `Sync`), no inner `&'a mut i32`

Introduction

Overview of Rust

λ_{Rust} : a small idealized Rust

A semantic model for λ_{Rust}

Conclusion

λ_{Rust} : Main goal

Formalize Rust's type system without its main complications.

- Close to MIR (`rustc` intermediate language)
- Omitted: traits, polymorphism, panics, weak memory*, IO, destructors...

*Work in progress

λ_{Rust} : Syntax and operational semantics

Lambda-calculus with extensions:

- Integers and Boolean
- Heap (manual (de)allocation) and pointers (block ID + offset)
- Concurrency (`fork`, atomic memory accesses, CAS)

Operational semantics:

- Small-step style
- Stuck:
 - Type errors
 - Incorrect memory accesses (incl. double free, ...)
 - Data races

λ_{Rust} : Type system

$E; L \mid K; T \vdash F$

Programs written in **continuation-passing style**

- MIR programs: control flow graphs
- \implies No output type

λ_{Rust} : Type system

$E; L \mid K; \mathbf{T} \vdash F$

Type context

Each variable: integer or pointer.

Associated with **type + ownership**. Examples:

- $x \triangleleft \mathbf{int}$
- $p \triangleleft \mathbf{box}(\mathbf{int} \times \mathbf{int} + ())$
- $p \triangleleft \&_{\text{mut}}^{\alpha}(\mathbf{int} \times \mathbf{int} + ()); p \triangleleft \dagger^{\alpha} \mathbf{box}(\mathbf{int} \times \mathbf{int} + ())$
- $p \triangleleft \&_{\text{shr}}^{\alpha} \mathbf{int}$

λ_{Rust} : Type system

E; L | **K; T** \vdash F

Lifetime contexts

Contain information for:

- Lifetime currently alive: **E; L** \vdash α alive
- Lifetime inclusions: **E; L** \vdash $\alpha \sqsubseteq \beta$
 - We know *in advance* that α is shorter than β
- Allowing to end a lifetime
 - Check: no promise to end another lifetime earlier.

λ_{Rust} : Type system

$E; L \mid \mathbf{K}; T \vdash F$



Continuation context

- Calling a continuation \Leftrightarrow Jumping to another block in CFG
- May require constraints on typing context and lifetime context
- Example:

$k \triangleleft \mathbf{cont}(\text{"}\alpha \text{ alive" } ; r \triangleleft \mathbf{box}(() + \&_{\text{mut}}^{\alpha} \mathbf{int}), x \triangleleft \mathbf{box} \mathbf{int})$

Introduction

Overview of Rust

λ_{Rust} : a small idealized Rust

A semantic model for λ_{Rust}

Conclusion

Challenge #1: the right approach

- One can write **unsafe code** in a **safely encapsulated** manner
 - New types that are safe for **new reasons**
- Our goal: prove that these library are safe
 - \implies **Syntactic approaches will not work**

We build a **logical relation** for Rust's type system

Choosing the right logic

Rust type system: **Ownership**, complex **sharing protocols**, in a **concurrent setting**

- **Iris** is a concurrent separation logic framework that we have been developing since 2014 [POPL'15, ICFP'16, ESOP'17, POPL'17, ECOOP'17]
- Iris has built-in support for these features and furthermore supports deriving new custom logics and mechanizing proofs in Coq
- \implies Iris is the right tool for modeling Rust!

Proof method

- We define a **logical relation in Iris**:

$$\mathbf{E}; \mathbf{L} \mid \mathbf{K}; \mathbf{T} \models F \quad \triangleq \quad \{[\mathbf{E}] * [\mathbf{L}] * [\mathbf{K}] * [\mathbf{T}]\} F \ \{\text{True}\}$$

- The relation is **compatible** with type-checking rules:

$$\mathbf{E}; \mathbf{L} \mid \mathbf{K}; \mathbf{T} \vdash F \quad \Longrightarrow \quad \mathbf{E}; \mathbf{L} \mid \mathbf{K}; \mathbf{T} \models F$$

- The relation is **adequate**:

$$\mathbf{E}; \mathbf{L} \mid \mathbf{K}; \mathbf{T} \models F \quad \Longrightarrow \quad F \text{ is safe}$$

- Conclusion: **well-typed programs can't go wrong**
 - No data race, no memory error, ...

Logical interpretation of types

Example 1: `int`

What values are integers?

$$\llbracket \text{int} \rrbracket.\text{own}(v) \triangleq \exists n \in \mathbb{Z}. v = n$$

Logical interpretation of types

Example 2: **box** τ

We must state **ownership** of memory and inner type:

$$\llbracket \mathbf{box} \tau \rrbracket.\text{own}(v) \triangleq \exists l \in \mathcal{L}. v = l * \exists v'. l \mapsto v' * \llbracket \tau \rrbracket.\text{own}(v')$$

Logical interpretation of types

Example 3: $\tau_1 \times \tau_2$

Actually, types refer to **list of values**

$$\llbracket \tau_1 \times \tau_2 \rrbracket.\text{own}(\bar{v}) \triangleq \exists \bar{v}_1 \bar{v}_2. \bar{v} = \bar{v}_1 ++ \bar{v}_2 * \llbracket \tau_1 \rrbracket.\text{own}(\bar{v}_1) * \llbracket \tau_2 \rrbracket.\text{own}(\bar{v}_2)$$

This also ensures **no aliasing** between \bar{v}_1 and \bar{v}_2

Challenge #2: interpreting borrows

Mutable borrows $\&_{\text{mut}}^{\alpha} \tau$

Pointer with **temporary ownership**

Recall:

$$\begin{aligned} \llbracket \mathbf{box} \tau \rrbracket.\text{own}(\bar{v}) &\triangleq \\ \exists \ell \in \mathcal{L}. \bar{v} = [\ell] * \exists \bar{w}. \ell \mapsto \bar{w} * \llbracket \tau \rrbracket.\text{own}(\bar{w}) \end{aligned}$$

Challenge #2: interpreting borrows

Mutable borrows $\&_{\text{mut}}^{\alpha} \tau$

Pointer with **temporary ownership**

For mutable borrows, we use the full borrow assertion $\&_{\text{full}}^{\kappa} -$:

$$\begin{aligned} \llbracket \&_{\text{mut}}^{\alpha} \tau \rrbracket.\text{own}(\bar{v}) &\triangleq \\ \exists \ell \in \mathcal{L}. \quad \bar{v} = [\ell] * \&_{\text{full}}^{\llbracket \alpha \rrbracket} (\exists \bar{w}. \ell \mapsto \bar{w} * \llbracket \tau \rrbracket.\text{own}(\bar{w})) \end{aligned}$$

Challenge #2: interpreting borrows

Mutable borrows $\&_{\text{mut}}^{\alpha} \tau$

Pointer

$\&_{\text{full}}^{\kappa} P$ is a separation logic assertion

- Intuition: ownership of P only if κ is *alive*
- Part of the **lifetime logic** (logical library defined in Iris)

Challenge #2: interpreting borrows

Shared borrows $\&_{\text{shr}}^{\alpha} \tau$

$$\llbracket \&_{\text{shr}}^{\alpha} \tau \rrbracket.\text{own}(\bar{v}) \triangleq \exists l \in \mathcal{L}. \bar{v} = [l] * \text{ ???}$$

Challenge #2: interpreting borrows

Shared borrows $\&_{\text{shr}}^{\alpha} \tau$

$$\llbracket \&_{\text{shr}}^{\alpha} \tau \rrbracket.\text{own}(\bar{v}) \triangleq \exists l \in \mathcal{L}. \bar{v} = [l] * \text{ ???}$$

A type chooses its sharing protocol.

Examples:

- Read-only for $\&'a \text{ i32}$,
- With a lock for $\&'a \text{ Mutex}\langle \text{i32} \rangle$
- Thread-local + accessors for $\&'a \text{ Cell}\langle \text{i32} \rangle$

Challenge #2: interpreting borrows

Shared borrows $\&_{\text{shr}}^{\alpha} \tau$

$$\llbracket \&_{\text{shr}}^{\alpha} \tau \rrbracket.\text{own}(\bar{v}) \triangleq \exists \ell \in \mathcal{L}. \bar{v} = [\ell] * \text{ ???}$$

A type chooses its sharing protocol.

In our model, types have **two interpretations**

- *Ownership predicate* $\llbracket \tau \rrbracket.\text{own}(\bar{v})$: when exclusively owned
- *Sharing predicate* $\llbracket \tau \rrbracket.\text{shr}(\kappa, \ell)$: when shared (under a shared borrow)

Challenge #2: interpreting borrows

Shared borrows $\&_{\text{shr}}^{\alpha} \tau$

$$\llbracket \&_{\text{shr}}^{\alpha} \tau \rrbracket.\text{own}(\bar{v}) \triangleq \exists \ell \in \mathcal{L}. \bar{v} = [\ell] * \llbracket \tau \rrbracket.\text{shr}(\llbracket \alpha \rrbracket, \ell)$$

A type chooses its sharing protocol.

In our model, types have **two interpretations**

- *Ownership predicate* $\llbracket \tau \rrbracket.\text{own}(\bar{v})$: when exclusively owned
- *Sharing predicate* $\llbracket \tau \rrbracket.\text{shr}(\kappa, \ell)$: when shared (under a shared borrow)

Challenge #2: interpreting borrows

Shared borrow α

$[[\tau]].shr(\kappa, \ell)$ has to be **persistent**.

- I.e., duplicable, does not contain resources
- So that $\&_{shr}^{\alpha} \tau$ can be shared

Lifetime logic provides tools for defining $\&_{shr}^{\alpha} \tau$

- $\&_{frac}^{\kappa} \Phi$: access to only a fraction of Φ
 - Read-only protocol
- $\&_{at}^{\kappa} P$: access to P only for atomic operations
 - Mutex...
- ...

In our

■ Own

■ Share

borrow)

Challenge #3: thread safety

Reminder, in Rust:

- T : **Send** \iff T can be **moved to another thread**
- T : **Sync** \iff T can be **shared with another thread**

How to model this?

Challenge #3: thread safety

Reminder, in Rust:

- T : **Send** \iff T can be **moved to another thread**
- T : **Sync** \iff T can be **shared with another thread**

$\llbracket \tau \rrbracket$.own and $\llbracket \tau \rrbracket$.shr **can depend on thread ID:**

$$\llbracket \tau \rrbracket.\text{own}(t, \bar{v}) = \dots \quad \llbracket \tau \rrbracket.\text{shr}(\kappa, t, \ell) = \dots$$

- τ is **Send** : $\llbracket \tau \rrbracket$.own(t, \bar{v}) independent of t
- τ is **Sync** : $\llbracket \tau \rrbracket$.shr(κ, t, ℓ) independent of t
- In the lifetime logic: $\&_{\text{na}}^{\kappa/t} P$
 - Shared P while κ is alive, limited to thread t

Introduction

Overview of Rust

λ_{Rust} : a small idealized Rust

A semantic model for λ_{Rust}

Conclusion

Conclusion

And also...

- **Lifetime logic**: a library in Iris
 - Logical model of lifetimes and borrows
 - Main idea: **split ownership over time** (instead of “over space”)
- Model of most of Rust's types with **interior mutability**
 - `Cell<T>`, `RefCell<T>`, `Rc<T>`, `Arc<T>`, `Mutex<T>`, `RwLock<T>`
- **Subtyping**/lifetime inclusions

`http://plv.mpi-sws.org/rustbelt/`

Borrows and inheritance

Usually: we split ownership with respect to space

Let's allow **splitting ownership over time**:

$$\triangleright P \Rightarrow \&_{\text{full}}^{\alpha} P * ([\dagger\alpha] \Rightarrow \triangleright P)$$



$\triangleright P$ can be transformed into...

Borrows and inheritance

Usually: we split ownership with respect to space

Let's allow **splitting ownership over time**:

$$\triangleright P \quad \Rightarrow \quad \&_{\text{full}}^{\alpha} P \quad * \quad ([\dagger\alpha] \Rightarrow \triangleright P)$$



A *borrowed* part:

- access of P when α is ongoing
- P must be *preserved* when α ends

Borrows and inheritance

Usually: we split ownership with respect to space

Let's allow **splitting ownership over time**:

$$\triangleright P \quad \Rightarrow \quad \&_{\text{full}}^{\alpha} P \quad * \quad ([\dagger\alpha] \Rightarrow \triangleright P)$$

An *inheritance* part, that gives back P when α is finished.

Lifetime tokens

How to witness that α is alive?

We use a **lifetime token** $[\alpha]$

- **Left in deposit** when opening a borrow:

$$\&_{\text{full}}^{\alpha} P * [\alpha] \quad \Rightarrow \quad \triangleright P * (\triangleright P \Rightarrow \&_{\text{full}}^{\alpha} P * [\alpha])$$

- Needed to **terminate** α :

$$[\alpha] \Rightarrow [\dagger\alpha]$$

```
fn f<'a>(pair : &'a mut (i32, i32)) {  
    let fst : &'a mut i32 = &mut pair.0;  
    let snd : &'a mut i32 = &mut pair.1;  
    join(|| *fst *= 2,  
        || *snd += 1);  
}
```

```
fn f<'a>(pair : &'a mut (i32, i32)) {  
    let fst : &'a mut i32 = &mut pair.0;  
    let snd : &'a mut i32 = &mut pair.1;  
    join(|| *fst *= 2,  
        || *snd += 1);  
}
```

We need to **split borrows**:

$$\&_{\text{full}}^{\alpha}(P * Q) \iff \&_{\text{full}}^{\alpha} P * \&_{\text{full}}^{\alpha} Q$$


```
fn f<'a>(pair : &'a mut (i32, i32)) {  
    let fst : &'a mut i32 = &mut pair.0;  
    let snd : &'a mut i32 = &mut pair.1;  
    join(|| *fst *= 2,  
        || *snd += 1);  
}
```

Both threads witness that the lifetime is alive.
⇒ We make the lifetime token **fractional**:

$$[\alpha]_{q+q'} \Leftrightarrow [\alpha]_q * [\alpha]_{q'}$$

Fractional lifetime tokens

How to witness that α is alive?

We use **lifetime tokens** $[\alpha]_q$

- Fractional: $[\alpha]_{q+q'} \Leftrightarrow [\alpha]_q * [\alpha]_{q'}$
- **Full token** needed to **terminate** α :

$$[\alpha]_1 \Rightarrow [\dagger\alpha]$$

- **Fraction** left in deposit when opening a borrow:

$$\&_{\text{full}}^{\alpha} P * [\alpha]_q \quad \Rightarrow \quad \triangleright P * (\triangleright P \Rightarrow \&_{\text{full}}^{\alpha} P * [\alpha]_q)$$

Sharing protocols

$$\llbracket \&'a T \rrbracket.\text{own}(\llbracket l \rrbracket) \triangleq \llbracket T \rrbracket.\text{shr}(\llbracket 'a \rrbracket, l) \triangleq ?$$

Depends on T. Common idea:

- **Share a borrow** using an invariant:

$$\llbracket T \rrbracket.\text{shr}(\llbracket 'a \rrbracket, l) \triangleq \boxed{\&_{\text{full}}^{\alpha} [\text{Protocol}]}^{\mathcal{N}}$$

- Technical problems with step-indexing
- Specific construction: **persistent borrows**: $\&_{\text{at}}^{\alpha} [\text{Protocol}]$
 - Behave like **cancellable invariant**