

Higher-Order Leak and Deadlock Free Locks

Jules Jacobs

Radboud University

Stephanie Balzer

Carnegie Mellon University

Memory management with substructural types

```
fn min(x: u32, y: u32) → u32 {  
  let mut v = Vec::new();  
  v.push(x);  
  v.push(y);  
  v.sort();  
  return v[0];  
  // v is deallocated  
}
```

- ▶ Each heap allocation has a single owning reference
- ▶ Deallocated when owning reference disappears
- ▶ Prevents memory leaks...?

Memory leaks in Rust

`Arc<Mutex<T>>`

- ▶ Shareable mutable reference to `T`
 - ▶ Guarded by a lock
 - ▶ Reference-counted
- ▶ “Higher-order”: can store mutexes in mutexes

```
enum List { Nil, Cons(u32, Arc<Mutex<List>>) }
```

Memory leaks in Rust

`Arc<Mutex<T>>`

- ▶ Shareable mutable reference to `T`
 - ▶ Guarded by a lock
 - ▶ Reference-counted
- ▶ “Higher-order”: can store mutexes in mutexes

```
enum List { Nil, Cons(u32, Arc<Mutex<List>>) }
```

Memory leaks!

```
let x = Arc::new(Mutex::new(List::Nil)); // create list
*x.lock() = List::Cons(1, x.clone()); // create cycle

// refcount=2
drop(x);
// refcount=1 → list is leaked
```

Deadlocks in Rust

```
fn swap(x : &Mutex<u32>, y : &Mutex<u32>){  
    let mut gx = x.lock(); // acquire locks  
    let mut gy = y.lock();  
  
    let tmp = *gx; // swap contents  
    *gx = *gy;  
    *gy = tmp;  
  
    drop(gx); // release locks  
    drop(gy);  
}
```

Deadlocks!

```
spawn(||{ swap(x,y) });  
spawn(||{ swap(y,x) });
```

The paper in a nutshell

Can we get leak and deadlock freedom by type checking?

The paper in a nutshell

Can we get leak and deadlock freedom by type checking?

Language λ_{lock} with a *linearly typed lock API*

- ▶ Any lock in scope can be safely acquired
- ▶ “Higher-order locks”: can store locks in locks (recursively)
- ▶ *Sharing topology* remains acyclic by typing
- ▶ No leaks/deadlocks (✓ **in Coq**)

The paper in a nutshell

Can we get leak and deadlock freedom by type checking?

Language λ_{lock} with a *linearly typed lock API*

- ▶ Any lock in scope can be safely acquired
- ▶ “Higher-order locks”: can store locks in locks (recursively)
- ▶ *Sharing topology* remains acyclic by typing
- ▶ No leaks/deadlocks (✓ **in Coq**)

Extension $\lambda_{\text{lock++}}$ with *cyclic sharing topology*

- ▶ *Lock groups with local lock orders*
- ▶ Locks across groups can still be acquired in arbitrary order
- ▶ No leaks/deadlocks (✓ **in Coq**)

λ_{lock} 's lock type

$$\langle \tau \begin{matrix} a \\ b \end{matrix} \rangle \quad \begin{matrix} a \in \{0,1\} \\ b \in \{0,1\} \end{matrix}$$

A shareable reference to τ , similar to `Arc<Mutex<t>>` in Rust

λ_{lock} 's lock type

$$\langle \tau \begin{matrix} a \\ b \end{matrix} \rangle \quad \begin{matrix} a \in \{0,1\} \\ b \in \{0,1\} \end{matrix}$$

A shareable reference to τ , similar to `Arc<Mutex<t>>` in Rust


- ▶ **But** $\ell : \langle \tau \begin{matrix} a \\ b \end{matrix} \rangle$ is *linear*
- ▶ $a = 1$: this reference has to deallocate the lock
- ▶ $b = 1$: this reference has to release the lock

λ_{lock} 's lock API




new : $\mathbf{1} \multimap \langle \tau_1^1 \rangle$ (initially empty)

λ_{lock} 's lock API



new : $1 \multimap \langle \tau_1^1 \rangle$ (initially empty)


λ_{lock} 's lock API



new : $\mathbf{1} \multimap \langle \tau_1^1 \rangle$ (initially empty)

fork : $\langle \tau_{b_1+b_2}^{a_1+a_2} \rangle \times (\langle \tau_{b_2}^{a_2} \rangle \multimap \mathbf{1}) \multimap \langle \tau_{b_1}^{a_1} \rangle$


λ_{lock} 's lock API



new : $\mathbf{1} \multimap \langle \tau_1^1 \rangle$ (initially empty)

fork : $\langle \tau_{b_1+b_2}^{a_1+a_2} \rangle \times (\underbrace{\langle \tau_{b_2}^{a_2} \rangle \multimap \mathbf{1}}_{\text{to child thread}}) \multimap \underbrace{\langle \tau_{b_1}^{a_1} \rangle}_{\text{to parent thread}}$


λ_{lock} 's lock API



new : $\mathbf{1} \multimap \langle \tau_1^1 \rangle$ (initially empty)

fork : $\langle \tau_{b_1+b_2}^{a_1+a_2} \rangle \times (\underbrace{\langle \tau_{b_2}^{a_2} \rangle \multimap \mathbf{1}}_{\text{to child thread}}) \multimap \underbrace{\langle \tau_{b_1}^{a_1} \rangle}_{\text{to parent thread}}$


λ_{lock} 's lock API



new : $\mathbf{1} \multimap \langle \tau_1^1 \rangle$ (initially empty)

fork : $\langle \tau_{b_1+b_2}^{a_1+a_2} \rangle \times (\underbrace{\langle \tau_{b_2}^{a_2} \rangle}_{\text{to child thread}} \multimap \mathbf{1}) \multimap \underbrace{\langle \tau_{b_1}^{a_1} \rangle}_{\text{to parent thread}}$

λ_{lock} 's lock API




new : $\mathbf{1} \multimap \langle \tau_1^1 \rangle$ (initially empty)

fork : $\langle \tau_{b_1+b_2}^{a_1+a_2} \rangle \times \underbrace{(\langle \tau_{b_2}^{a_2} \rangle \multimap \mathbf{1})}_{\text{to child thread}} \multimap \underbrace{\langle \tau_{b_1}^{a_1} \rangle}_{\text{to parent thread}}$

release : $\langle \tau_1^a \rangle \times \tau \multimap \langle \tau_0^a \rangle$

acquire : $\langle \tau_0^a \rangle \multimap \langle \tau_1^a \rangle \times \tau$

λ_{lock} 's lock API




new : $\mathbf{1} \multimap \langle \tau_1^1 \rangle$ (initially empty)

fork : $\langle \tau_{b_1+b_2}^{a_1+a_2} \rangle \times (\underbrace{\langle \tau_{b_2}^{a_2} \rangle}_{\text{to child thread}} \multimap \mathbf{1}) \multimap \underbrace{\langle \tau_{b_1}^{a_1} \rangle}_{\text{to parent thread}}$

release : $\langle \tau_1^a \rangle \times \tau \multimap \langle \tau_0^a \rangle$

acquire : $\langle \tau_0^a \rangle \multimap \langle \tau_1^a \rangle \times \tau$

λ_{lock} 's lock API




new : $\mathbf{1} \multimap \langle \tau_1^1 \rangle$ (initially empty)

fork : $\langle \tau_{b_1+b_2}^{a_1+a_2} \rangle \times (\underbrace{\langle \tau_{b_2}^{a_2} \rangle}_{\text{to child thread}} \multimap \mathbf{1}) \multimap \underbrace{\langle \tau_{b_1}^{a_1} \rangle}_{\text{to parent thread}}$

release : $\langle \tau_1^a \rangle \times \tau \multimap \langle \tau_0^a \rangle$

acquire : $\langle \tau_0^a \rangle \multimap \langle \tau_1^a \rangle \times \tau$

λ_{lock} 's lock API




new : $\mathbf{1} \multimap \langle \tau_1^1 \rangle$ (initially empty)

fork : $\langle \tau_{b_1+b_2}^{a_1+a_2} \rangle \times \underbrace{(\langle \tau_{b_2}^{a_2} \rangle \multimap \mathbf{1})}_{\text{to child thread}} \multimap \underbrace{\langle \tau_{b_1}^{a_1} \rangle}_{\text{to parent thread}}$

release : $\langle \tau_1^a \rangle \times \tau \multimap \langle \tau_0^a \rangle$

acquire : $\langle \tau_0^a \rangle \multimap \langle \tau_1^a \rangle \times \tau$

λ_{lock} 's lock API




new : $\mathbf{1} \multimap \langle \tau_1^1 \rangle$ (initially empty)

fork : $\langle \tau_{b_1+b_2}^{a_1+a_2} \rangle \times (\underbrace{\langle \tau_{b_2}^{a_2} \rangle}_{\text{to child thread}} \multimap \mathbf{1}) \multimap \underbrace{\langle \tau_{b_1}^{a_1} \rangle}_{\text{to parent thread}}$

release : $\langle \tau_1^a \rangle \times \tau \multimap \langle \tau_0^a \rangle$

acquire : $\langle \tau_0^a \rangle \multimap \langle \tau_1^a \rangle \times \tau$

λ_{lock} 's lock API



new : $\mathbf{1} \multimap \langle \tau_1^1 \rangle$ (initially empty)

fork : $\langle \tau_{b_1+b_2}^{a_1+a_2} \rangle \times (\underbrace{\langle \tau_{b_2}^{a_2} \rangle}_{\text{to child thread}} \multimap \mathbf{1}) \multimap \underbrace{\langle \tau_{b_1}^{a_1} \rangle}_{\text{to parent thread}}$


release : $\langle \tau_1^a \rangle \times \tau \multimap \langle \tau_0^a \rangle$

acquire : $\langle \tau_0^a \rangle \multimap \langle \tau_1^a \rangle \times \tau$

drop : $\langle \tau_0^0 \rangle \multimap \mathbf{1}$ (decrements refcount)

wait : $\langle \tau_0^1 \rangle \multimap \tau$ (blocks until refcount = 1)

λ_{lock} 's lock API



new : $\mathbf{1} \multimap \langle \tau_1^1 \rangle$ (initially empty)

fork : $\langle \tau_{b_1+b_2}^{a_1+a_2} \rangle \times (\underbrace{\langle \tau_{b_2}^{a_2} \rangle \multimap \mathbf{1}}_{\text{to child thread}}) \multimap \underbrace{\langle \tau_{b_1}^{a_1} \rangle}_{\text{to parent thread}}$


release : $\langle \tau_1^a \rangle \times \tau \multimap \langle \tau_0^a \rangle$

acquire : $\langle \tau_0^a \rangle \multimap \langle \tau_1^a \rangle \times \tau$

drop : $\langle \tau_0^0 \rangle \multimap \mathbf{1}$ (decrements refcount)

wait : $\langle \tau_0^1 \rangle \multimap \tau$ (blocks until refcount = 1)

λ_{lock} 's lock API



new : $\mathbf{1} \multimap \langle \tau_1^1 \rangle$ (initially empty)

fork : $\langle \tau_{b_1+b_2}^{a_1+a_2} \rangle \times (\underbrace{\langle \tau_{b_2}^{a_2} \rangle \multimap \mathbf{1}}_{\text{to child thread}}) \multimap \underbrace{\langle \tau_{b_1}^{a_1} \rangle}_{\text{to parent thread}}$

release : $\langle \tau_1^a \rangle \times \tau \multimap \langle \tau_0^a \rangle$

acquire : $\langle \tau_0^a \rangle \multimap \langle \tau_1^a \rangle \times \tau$

drop : $\langle \tau_0^0 \rangle \multimap \mathbf{1}$ (decrements refcount)

wait : $\langle \tau_0^1 \rangle \multimap \tau$ (blocks until refcount = 1)

λ_{lock} 's lock API

T_1

v

new : $\mathbf{1} \multimap \langle \tau_1^1 \rangle$ (initially empty)

fork : $\langle \tau_{b_1+b_2}^{a_1+a_2} \rangle \times \underbrace{(\langle \tau_{b_2}^{a_2} \rangle \multimap \mathbf{1})}_{\text{to child thread}} \multimap \underbrace{\langle \tau_{b_1}^{a_1} \rangle}_{\text{to parent thread}}$

release : $\langle \tau_1^a \rangle \times \tau \multimap \langle \tau_0^a \rangle$

acquire : $\langle \tau_0^a \rangle \multimap \langle \tau_1^a \rangle \times \tau$

drop : $\langle \tau_0^0 \rangle \multimap \mathbf{1}$ (decrements refcount)

wait : $\langle \tau_0^1 \rangle \multimap \tau$ (blocks until refcount = 1)

Which concurrency patterns does λ_{lock} support?

Simultaneously acquire multiple locks:

```
let swap =  $\lambda(l_1 : \langle \tau_0^0 \rangle, l_2 : \langle \tau_0^0 \rangle)$ .
```

```
  let  $(l_1 : \langle \tau_1^0 \rangle, x_1 : \tau)$  = acquire( $l_1$ ) in
```

```
  let  $(l_2 : \langle \tau_1^0 \rangle, x_2 : \tau)$  = acquire( $l_2$ ) in
```

```
  let  $l_1 : \langle \tau_0^0 \rangle$  = release( $l_1, x_2$ ) in
```

```
  let  $l_2 : \langle \tau_0^0 \rangle$  = release( $l_2, x_1$ ) in
```

```
  ( $l_1, l_2$ )
```



deadlock free

Which concurrency patterns does λ_{lock} support?

Simultaneously acquire multiple locks:

```
let swap =  $\lambda(l_1 : \langle \tau_0^0 \rangle, l_2 : \langle \tau_0^0 \rangle)$ .
```

```
  let  $(l_1 : \langle \tau_1^0 \rangle, x_1 : \tau)$  = acquire( $l_1$ ) in
```

```
  let  $(l_2 : \langle \tau_1^0 \rangle, x_2 : \tau)$  = acquire( $l_2$ ) in
```

```
  let  $l_1 : \langle \tau_0^0 \rangle$  = release( $l_1, x_2$ ) in
```

```
  let  $l_2 : \langle \tau_0^0 \rangle$  = release( $l_2, x_1$ ) in
```

```
  ( $l_1, l_2$ )
```



deadlock free

Futures / promises / fork-join:

```
let  $l : \langle \tau_0^1 \rangle$  = fork(new() :  $\langle \tau_1^1 \rangle$ ,  $\lambda l : \langle \tau_1^0 \rangle$ ).
```

```
  drop(release( $l, E : \tau$ ))
```

```
) in ... wait( $l$ ) ...
```

Obligation to fulfill promise cannot be discarded

Locks in locks in λ_{lock}

Store locks in locks:

release($\ell_1 : \langle \langle \tau_b^a \rangle_1^0 \rangle$, $\ell_2 : \langle \tau_b^a \rangle$)



leak free

Another thread can **acquire**(ℓ_1) to obtain ℓ_2

Locks in locks in λ_{lock}

Store locks in locks:

release($\ell_1 : \langle \langle \tau_b^a \rangle_1^0 \rangle$, $\ell_2 : \langle \tau_b^a \rangle$)



leak free

Another thread can **acquire**(ℓ_1) to obtain ℓ_2

Recursive mutable data structures:

$\text{tree} = \langle \mathbf{1} + \tau \times \text{tree} \times \text{tree} \rangle_0^1$

Locks in locks in λ_{lock}

Store locks in locks:

release($\ell_1 : \langle \langle \tau_b^a \rangle_1^0 \rangle$, $\ell_2 : \langle \tau_b^a \rangle$)



leak free

Another thread can **acquire**(ℓ_1) to obtain ℓ_2

Recursive mutable data structures:

$\text{tree} = \langle \mathbf{1} + \tau \times \text{tree} \times \text{tree} \rangle_0^1$

Session typed channels as a library:

$s ::= !\tau.s \mid ?\tau.s \mid s \& s \mid s \oplus s \mid \text{End}_! \mid \text{End}_? \mid \mu x.s \mid x$

Locks in locks in λ_{lock}

Store locks in locks:

release($\ell_1 : \langle \langle \tau_b^a \rangle_1^0 \rangle, \ell_2 : \langle \tau_b^a \rangle$)



leak free

Another thread can **acquire**(ℓ_1) to obtain ℓ_2

Recursive mutable data structures:

$\text{tree} = \langle \mathbf{1} + \tau \times \text{tree} \times \text{tree} \rangle_0^1$

Session typed channels as a library:

$s ::= !\tau.s \mid ?\tau.s \mid s \& s \mid s \oplus s \mid \text{End}_! \mid \text{End}_? \mid \mu x.s \mid x$

Shared sessions:

$\langle \mu x. !\tau. ?\tau.s \oplus \text{End}_! \rangle_0^0$

From λ_{lock} to $\lambda_{\text{lock}++}$

In λ_{lock} , we can duplicate *one* lock on **fork**.

Is it sound to allow duplicating *two*?

let $(\ell_1, \ell_2) = \mathbf{fork}((\ell_1, \ell_2), \lambda(\ell_1, \ell_2). \dots)$ **in** \dots

From λ_{lock} to $\lambda_{\text{lock}++}$

In λ_{lock} , we can duplicate *one* lock on **fork**.
Is it sound to allow duplicating *two*?

let $(l_1, l_2) = \mathbf{fork}((l_1, l_2), \lambda(l_1, l_2). \dots)$ **in** \dots

No, because

- ▶ Deadlock: acquiring $(l_1$ then $l_2)$ in parallel with $(l_2$ then $l_1)$
- ▶ Leak: storing $(l_1$ into $l_2)$ in parallel with $(l_2$ into $l_1)$

From λ_{lock} to $\lambda_{\text{lock}++}$

In λ_{lock} , we can duplicate *one* lock on **fork**.
Is it sound to allow duplicating *two*?

let $(\ell_1, \ell_2) = \mathbf{fork}((\ell_1, \ell_2), \lambda(\ell_1, \ell_2). \dots)$ **in** \dots

No, because

- ▶ Deadlock: acquiring $(\ell_1$ then $\ell_2)$ in parallel with $(\ell_2$ then $\ell_1)$
- ▶ Leak: storing $(\ell_1$ into $\ell_2)$ in parallel with $(\ell_2$ into $\ell_1)$

Yes, if

- ▶ We make **acquire** and **wait** follow a lock order
- ▶ We prevent storing those locks inside each other

$\lambda_{\text{lock}++}$'s lock group type

$$\langle \tau_1^{a_1}_{b_1}, \dots, \tau_n^{a_n}_{b_n} \rangle$$

- ▶ We can only **acquire** and **wait** in the given order
- ▶ We can add and remove locks dynamically
- ▶ The type level list is a *local view* into a complete order.

$\lambda_{\text{lock++}}$'s lock group API

newgroup : $\mathbf{1} \multimap \langle \rangle$

dropgroup : $\langle \rangle \multimap \mathbf{1}$

new[k] : $\langle A, B \rangle \multimap \langle A, \tau_1^1, B \rangle$ (length(A) = k)

drop[k] : $\langle A, \tau_0^0, B \rangle \multimap \langle A, B \rangle$

release[k] : $\langle A, \tau_1^a, B \rangle \times \tau \multimap \langle A, \tau_0^a, B \rangle$

acquire[k] : $\langle A, \tau_0^a, B_0 \rangle \multimap \langle A, \tau_1^a, B_0 \rangle \times \tau$

wait[k] : $\langle A_0, \tau_0^1, B_0^1 \rangle \multimap \langle A_0, B_0^1 \rangle \times \tau$

fork : $\langle A \rangle \times (\langle B \rangle \multimap \mathbf{1}) \multimap \langle C \rangle$

(where $A = B \oplus C$)

Swap within a lock group

$\text{swap} : \langle \text{int}_0^0, \text{int}_0^0 \rangle \multimap \langle \text{int}_0^0, \text{int}_0^0 \rangle$

$\text{swap}(\ell) :=$

let $\ell, x = \text{acquire}[0](\ell)$ **in**

let $\ell, y = \text{acquire}[1](\ell)$ **in**

let $\ell = \text{release}[0](\ell, y)$ **in**


let $\ell = \text{release}[1](\ell, x)$ **in** ℓ

- ▶ Type system enforces an order *within* a group
- ▶ No restrictions between two groups
(Partial lock orders don't allow this!)

Dijkstra's dining philosophers

Lock groups allow $\lambda_{\text{lock++}}$ to have cyclic connectivity

- ▶ Example: *Dijkstra's Dining Philosophers*
- ▶ Every thread (*Philosopher*) has access to 2 locks (*forks*): $\langle \text{fork}_b^a, \text{fork}_{b'}^{a'} \rangle$
- ▶ Can grow the dining table dynamically (fractally growing example in the paper)



Leak and deadlock freedom theorem

Small-step semantics on config $\rho = \{X_1, \dots, X_n\}$ of threads & locks

$X \in \rho$ waits for $Y \in \rho$ if

- ▶ X is a thread attempting an operation on lock Y , or
- ▶ Y has a reference to lock X

Leak and deadlock freedom theorem

Small-step semantics on config $\rho = \{X_1, \dots, X_n\}$ of threads & locks

$X \in \rho$ waits for $Y \in \rho$ if

- ▶ X is a thread attempting an operation on lock Y , or
- ▶ Y has a reference to lock X

$X \in \rho$ is *reachable* if it transitively waits for $Y \in \rho$ that can step

$S \subseteq \rho$ is a *deadlock* if no $X \in S$ can step or waits for $Y \notin S$

Leak and deadlock freedom theorem

Small-step semantics on config $\rho = \{X_1, \dots, X_n\}$ of threads & locks

$X \in \rho$ waits for $Y \in \rho$ if

- ▶ X is a thread attempting an operation on lock Y , or
- ▶ Y has a reference to lock X

$X \in \rho$ is *reachable* if it transitively waits for $Y \in \rho$ that can step

$S \subseteq \rho$ is a *deadlock* if no $X \in S$ can step or waits for $Y \notin S$

Theorem: all $X \in \rho$ reachable \iff no deadlocks $\emptyset \subset S \subseteq \rho$

Theorem: well-typed programs are leak and deadlock free

Corollary: global progress: $\rho \neq \emptyset \rightarrow \rho$ steps

✓ in Coq ($\approx 13k$ loc)

Insight: leak and deadlock freedom are related

Related and future work (non-exhaustive)

Related work

- ▶ CLASS – Rocha and Caires (ICFP'21, ESOP'23)
- ▶ Client-server sessions – Qian, Kavvos, Birkedal (ICFP'21)
- ▶ Usages/obligations – Kobayashi et al. (see paper)
- ▶ Priorities – Padovani, Dharda et al. (see paper)
- ▶ Manifest sharing – Balzer et al. (ICFP'17,ESOP'19)
- ▶ Session types – (see paper)

Future work

- ▶ DAG-shaped mutable data structures / `Rc<RefCell<T>>`
- ▶ Integration with Rust features (borrowing & `unsafe`)

Rust is a *practical* memory-safe language without GC ✓

Can a *practical* language be leak and deadlock free?

I hope to have convinced you that:

- ▶ This might be possible & is worth trying
- ▶ Promising direction: single ownership → sharing topology

Rust is a *practical* memory-safe language without GC ✓

Can a *practical* language be leak and deadlock free?

I hope to have convinced you that:

- ▶ This might be possible & is worth trying
- ▶ Promising direction: single ownership → sharing topology

“The authors didn’t even have to hide a bunch of more complicated rules in an appendix.” – Reviewer A

(P.S. I’m looking for a postdoc position)