Verifying Safe Rust Clients of Internally-Unsafe Libraries

federico.poli@inf.ethz.ch

Federico Poli Peter Müller Alexander J. Summers

ETH zürich

UBC THE UNIVERSITY OF BRITISH COLUMBIA

Rust Verification Workshop 2023

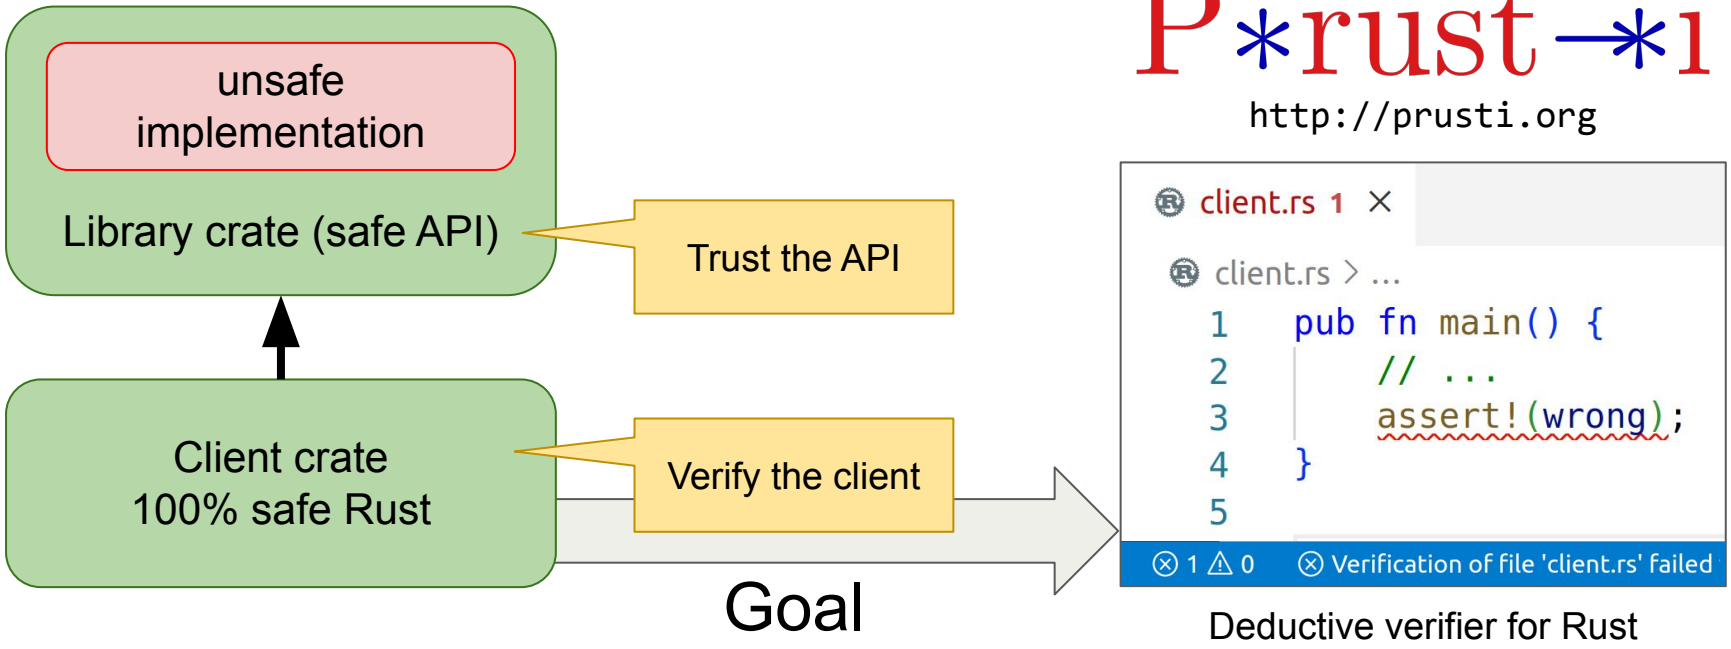
unsafe implementation
Library crate (safe API)
Trust the API
Client crate 100% safe Rust
Verify the client
Goal
P*rust→*i
http://prusti.org
client.rs 1
client.rs > ...
1 pub fn main() {
2 // ...
3 assert!(wrong);
4 }
5
⊗ 1 ⚠ 0 ⊗ Verification of file 'client.rs' failed
Deductive verifier for Rust

Background: fully safe code

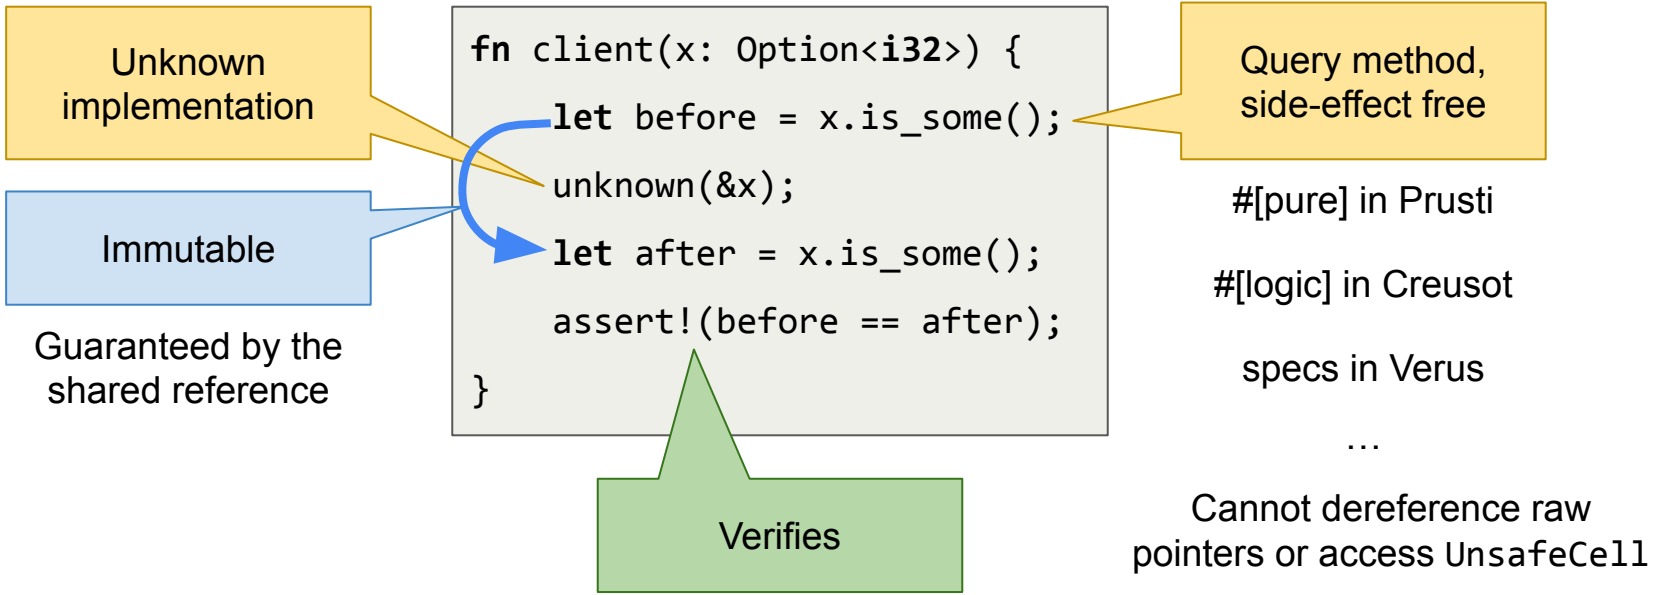

Now: libraries implemented with unsafe code

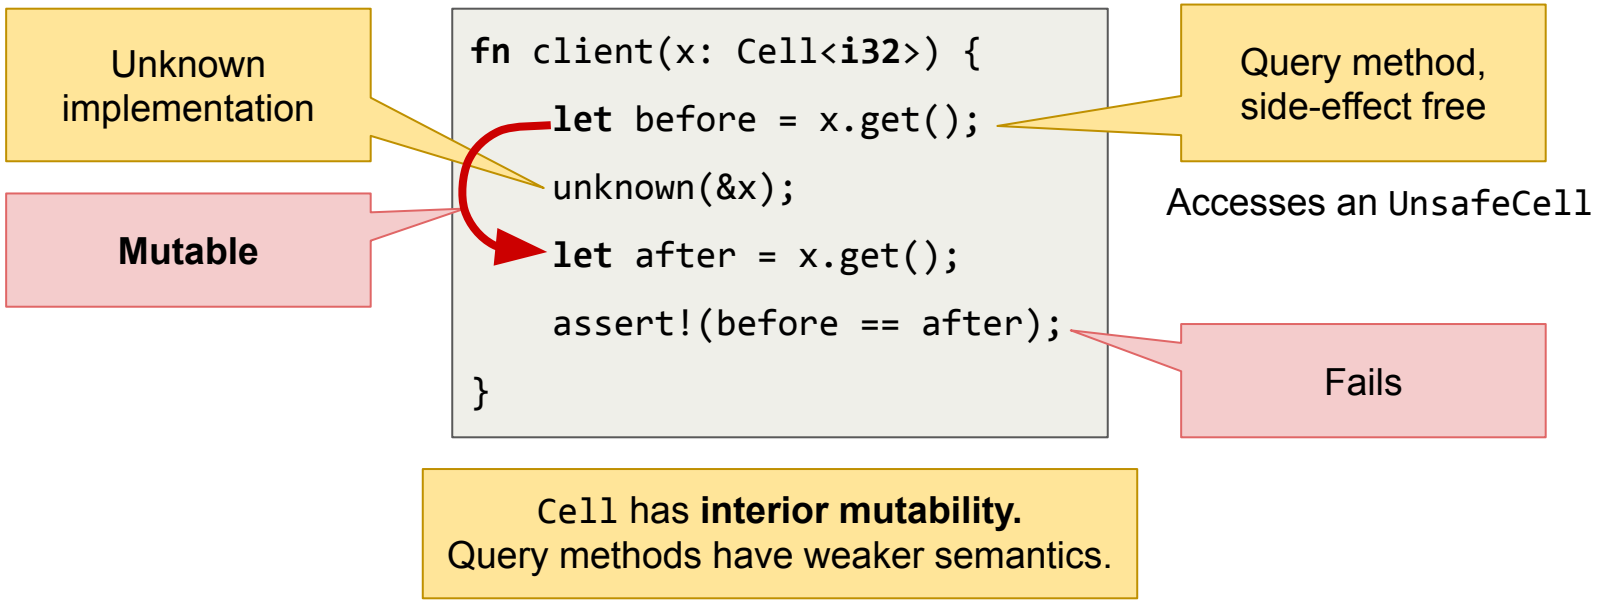

Interior mutability

In this talk:
implemented with `UnsafeCell` or raw pointers

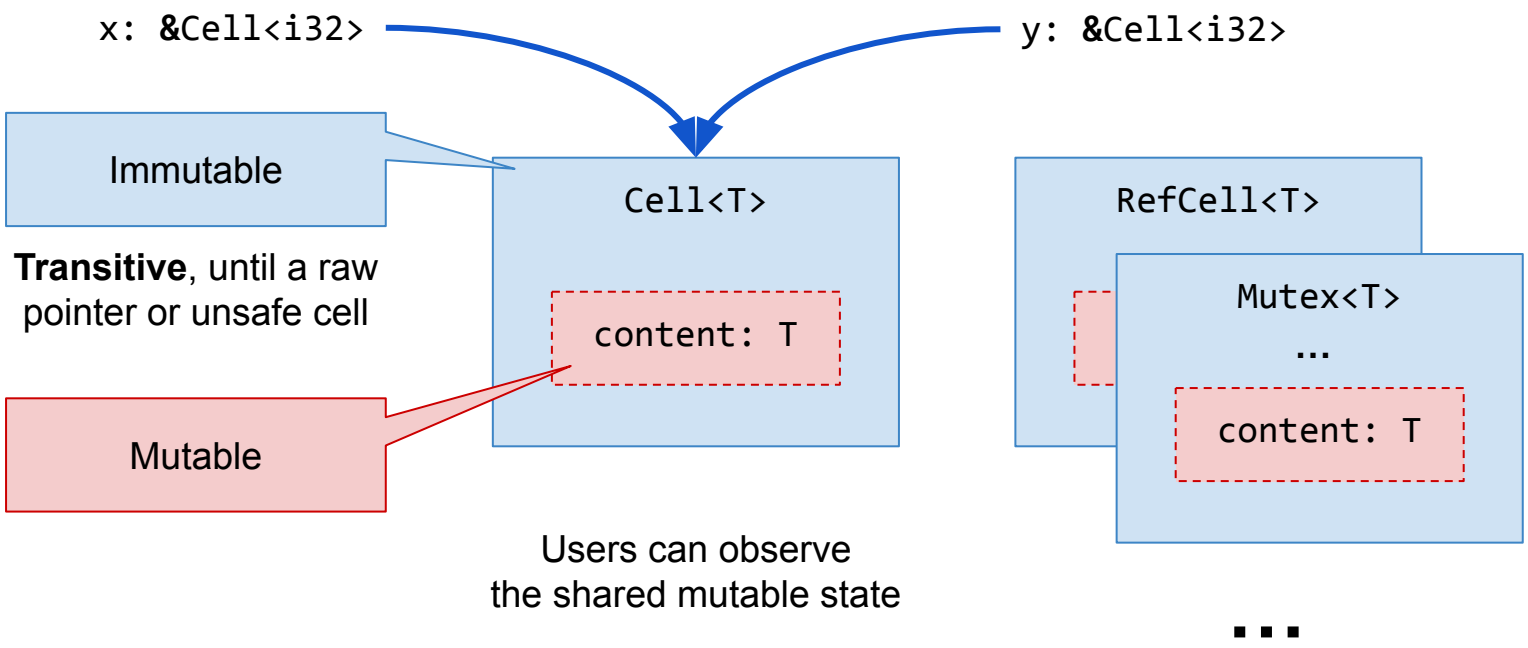

Capabilities (non zero-size types)

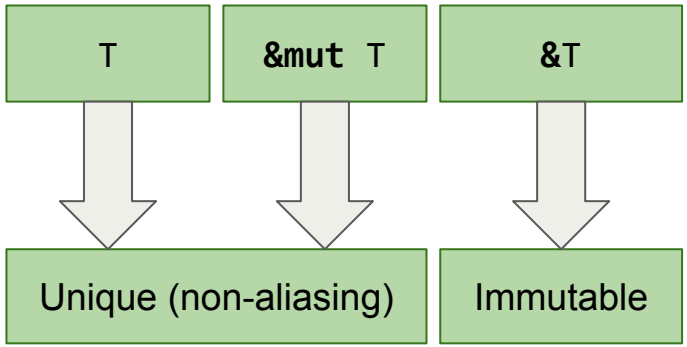

Content of:

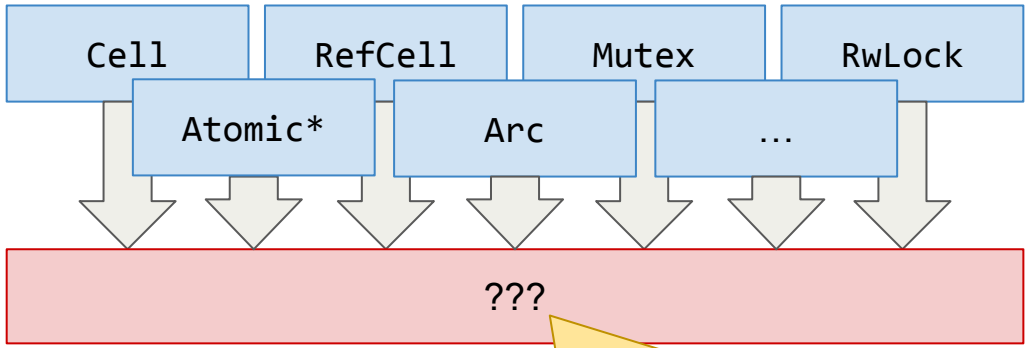

This talk:

They have capabilities

- Provided by auxiliary types (e.g. `MutexGuard`, `Ref`…)
- Depending on the state (e.g. reference count)
- Not always expressible using Rust types

Capabilities (non zero-size types)

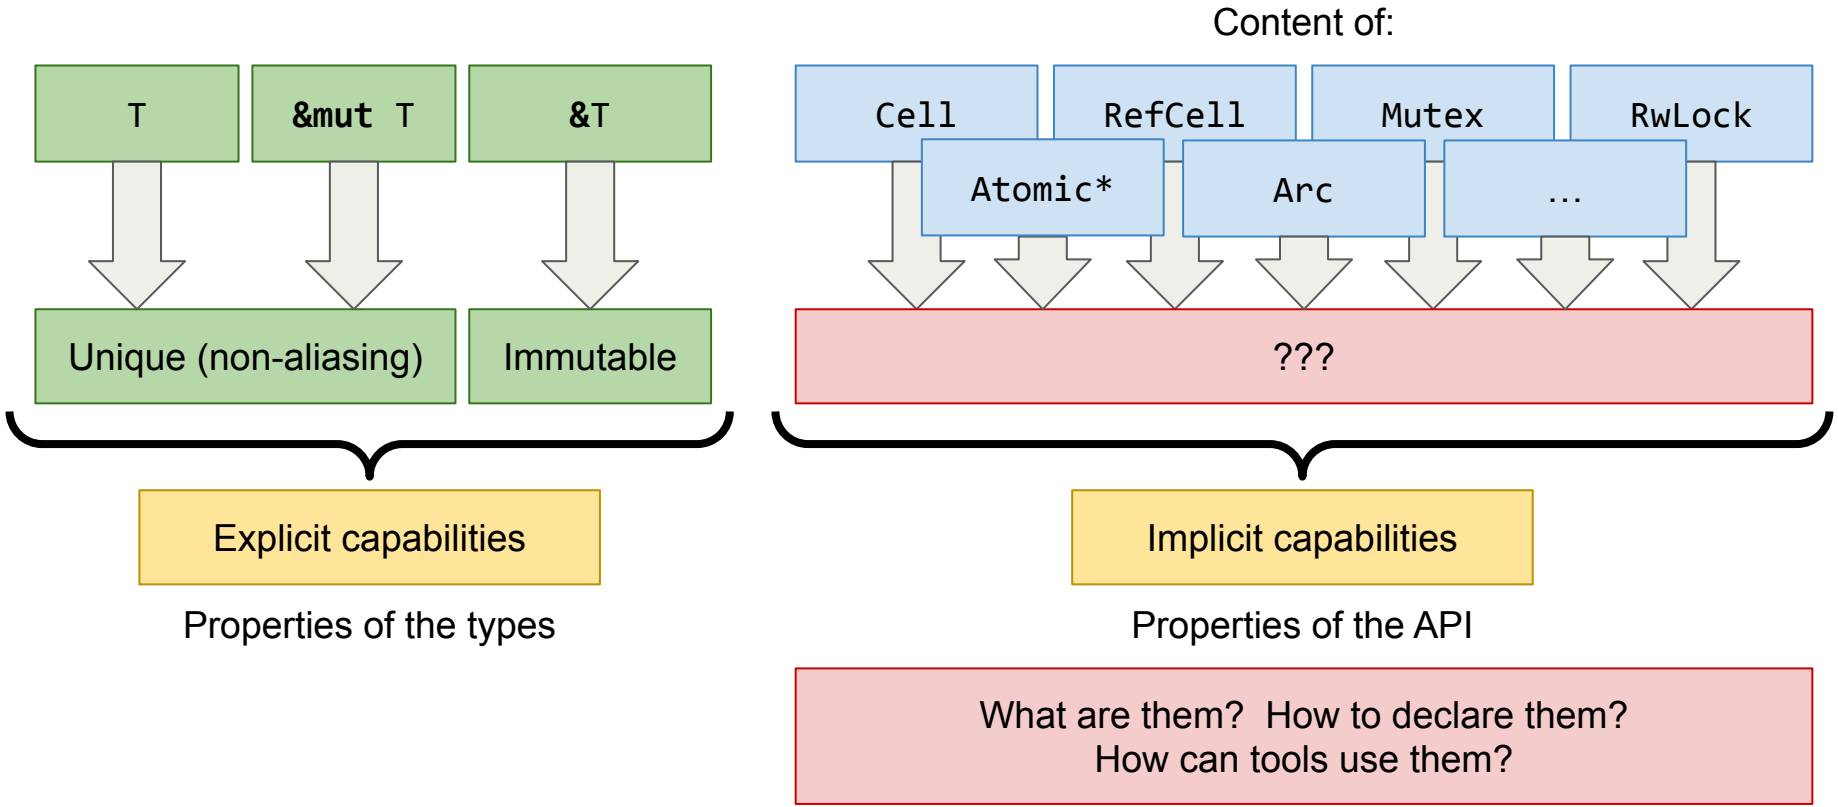

Example: RefCell

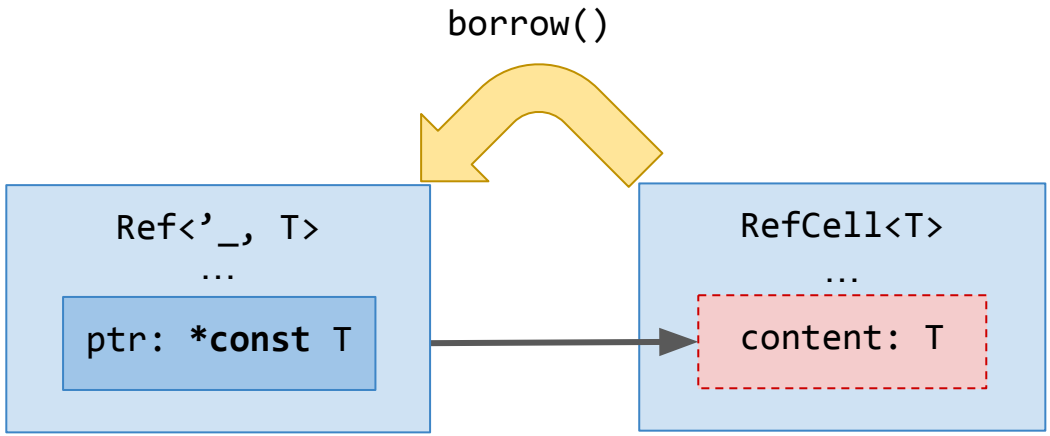

Read-only API

Immutability

Example: RefCell

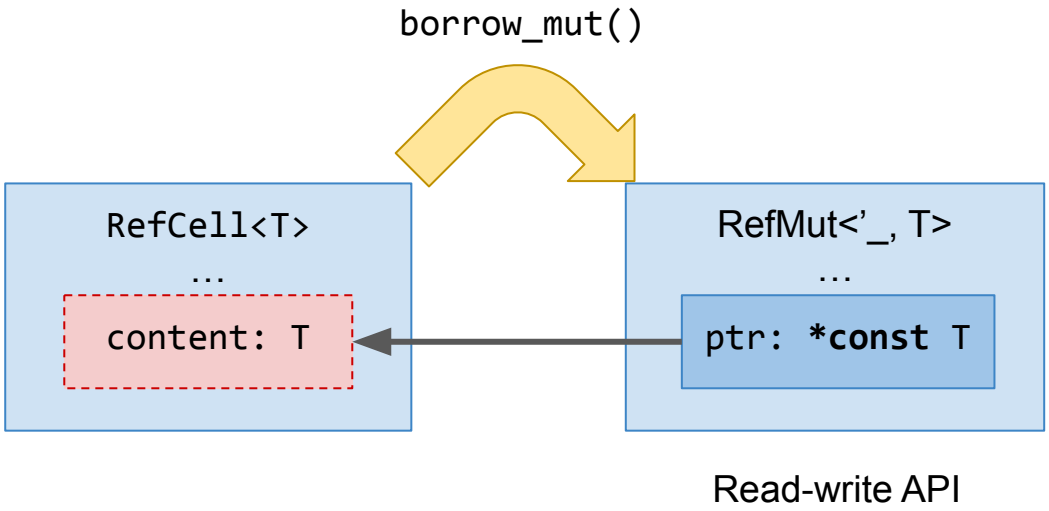

Read-write API

Example: RefCell

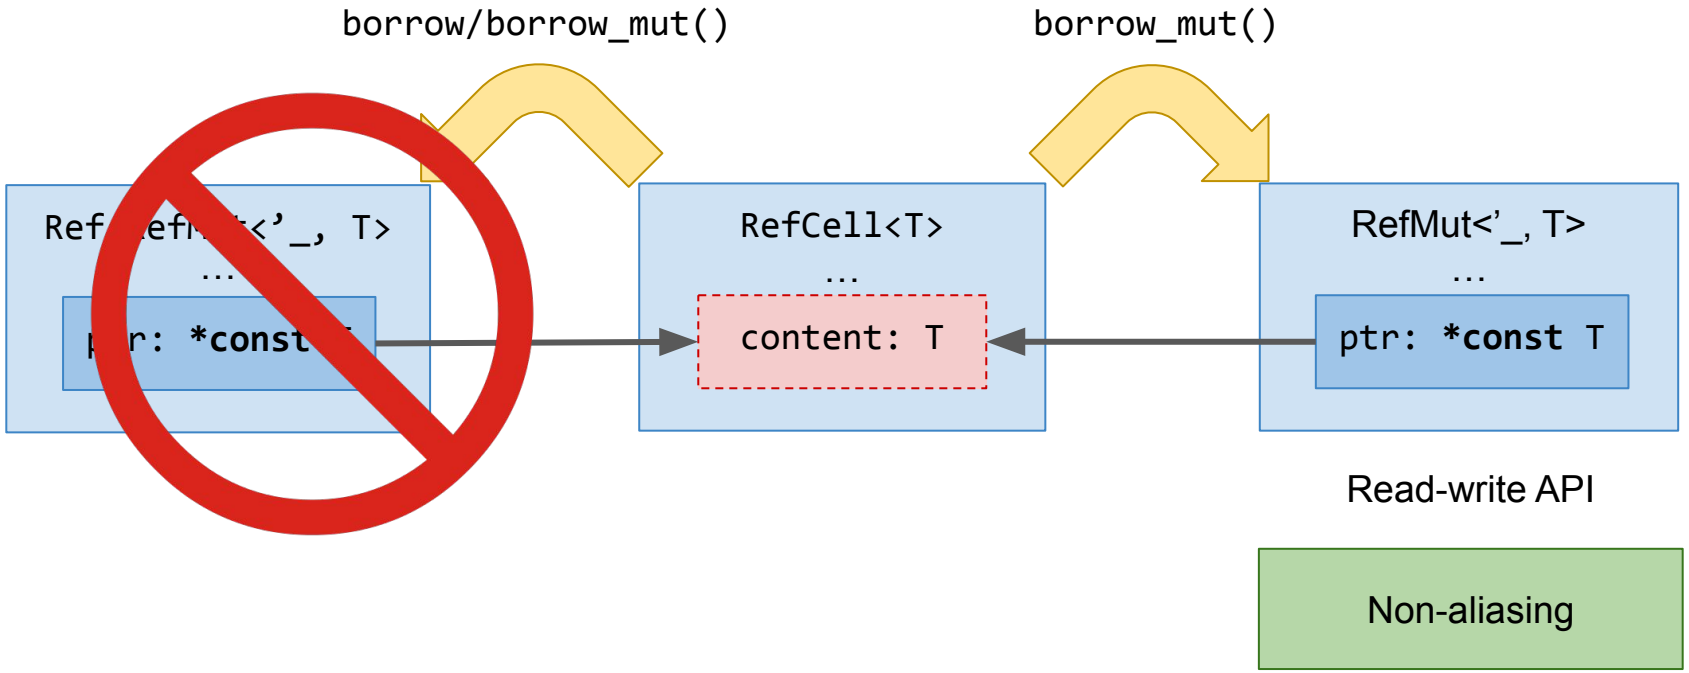

Implicit capability annotation

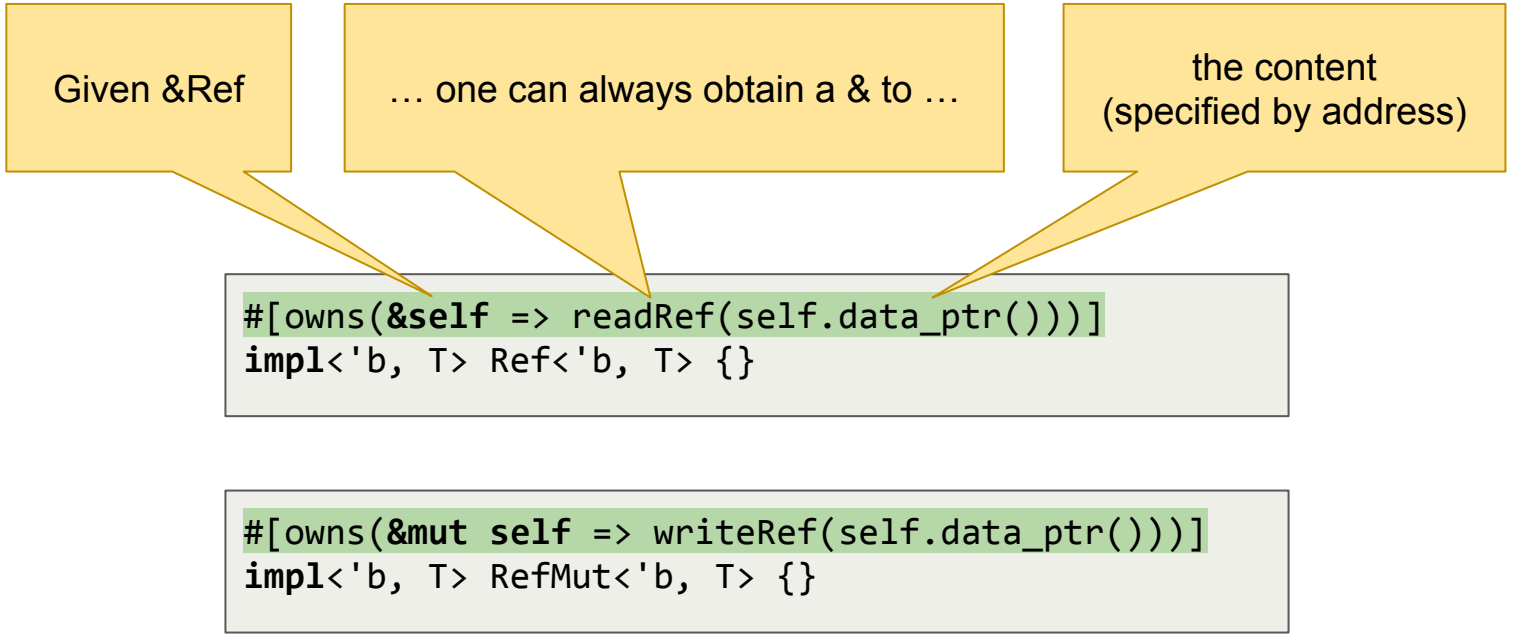

Properties of implicit capabilities

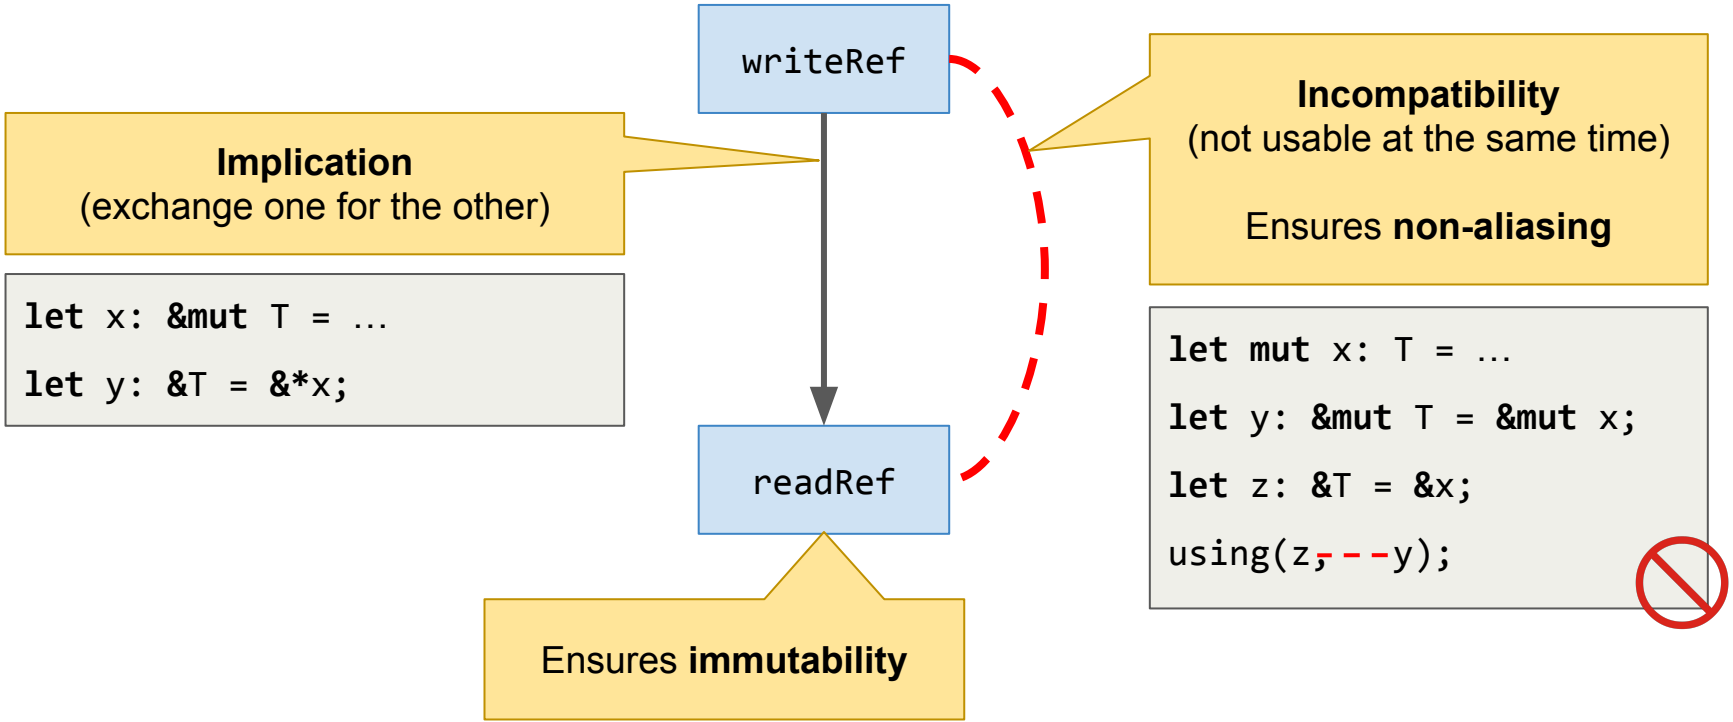

Verification example 1

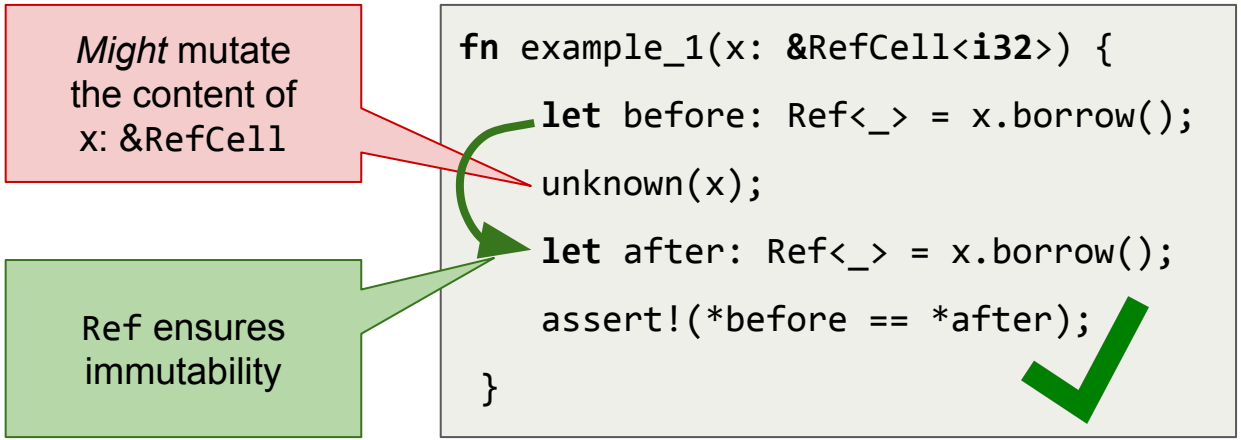

Verification example 2

The guards cannot refer to the content of the same `RwLock`

```
fn example_2(a: &RwLock<i32>) {
    let Ok(guard_1) = a.write() else { return; };
    let Ok(guard_2) = a.read() else { return; };
    unreachable!();
}
```

Verification example 3

```
fn example_3(x: Arc<i32>, y: Arc<i32>) {
    if Arc::strong_count(&x) != 1 {
        assert!(Arc::strong_count(&x) != 1);
    } else {
        assert!(Arc::strong_count(&x) == 1);
        assert!(Arc::as_ptr(&x) != Arc::as_ptr(&y));
    }
}
```

What verifies and what does not?

Assume no weak references

Demo

Demo: clients

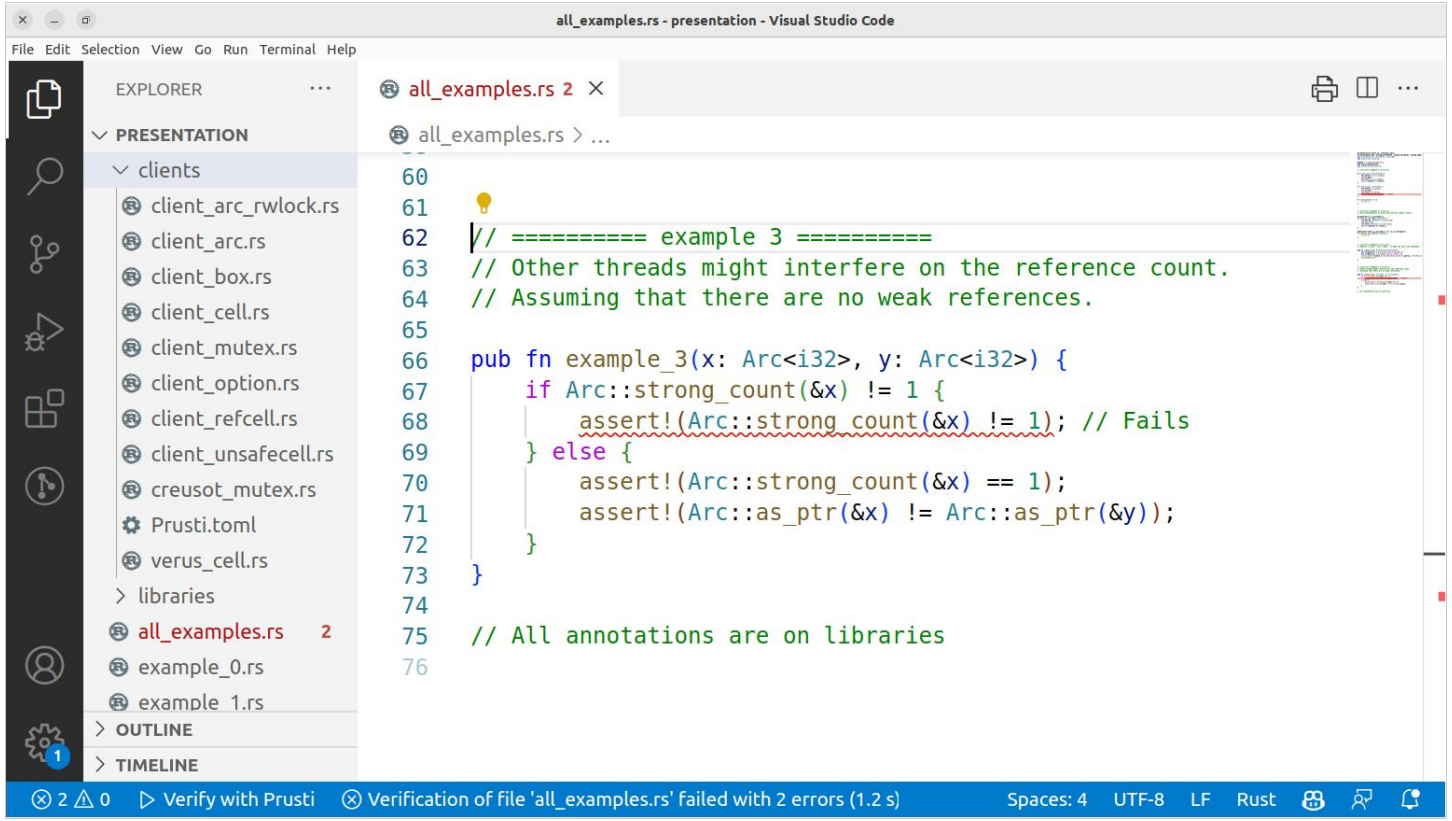

Demo: library annotations

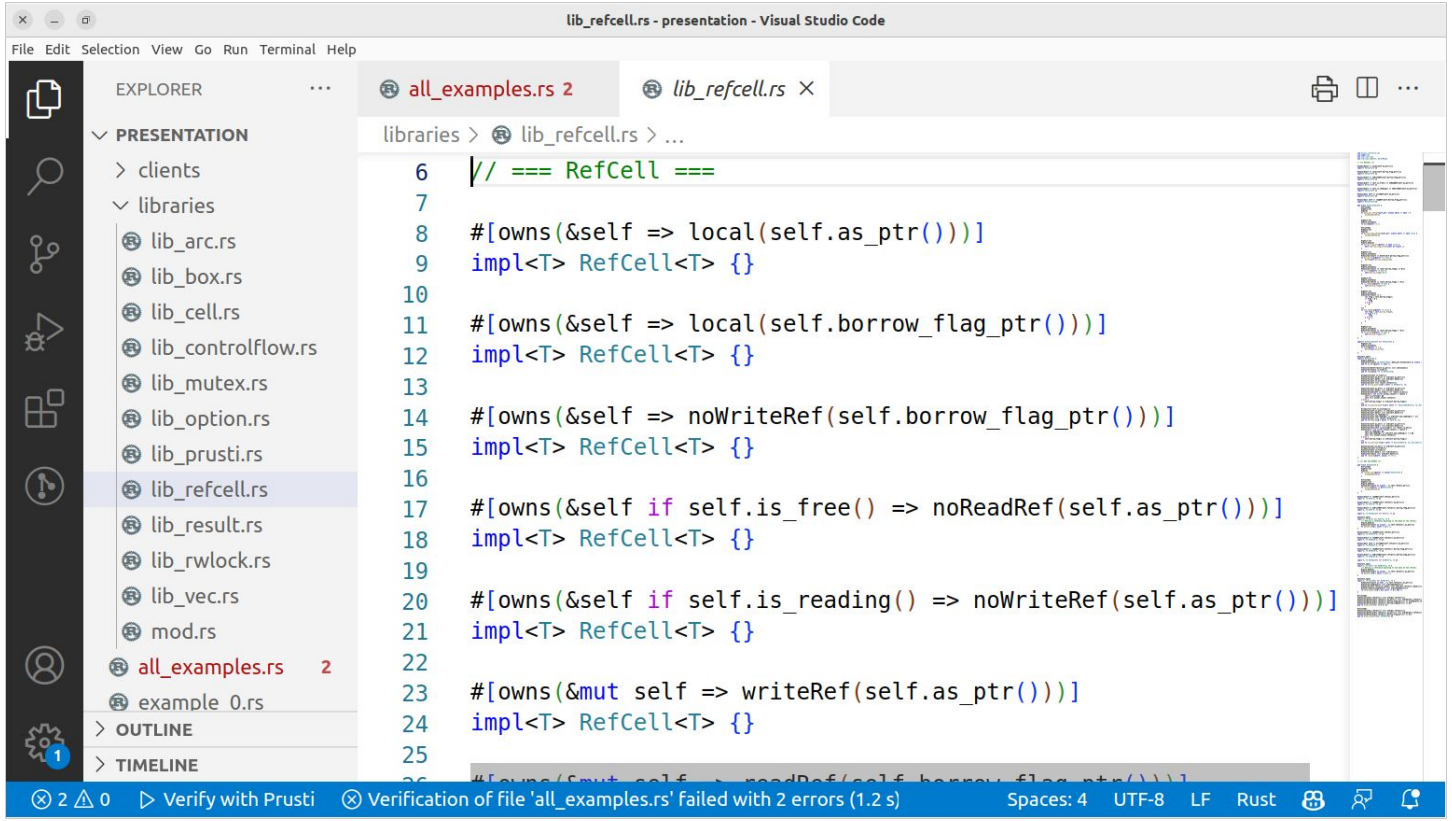

More types…

RefCell

Mutex

RwLock

Cell, Arc, Rc, Atomic, …

```
#[owns(&self => rea Ref(self.data_p ()))]
impl<'b, T> RefMut< b, T> {}

#[owns(&mut self
impl<'b, T> RefMu

…
```

```
#[owns(&mut self => writeRef(self.data_ptr()))]
impl<'a, T> MutexGuard<'a, T> {}

#[owns(&self =>
impl<'a, T> Mut

…
```

```
#[owns(&mut self => writeRef(self.data_ptr()))]
impl<'a, T> RwLock<'a, T> {}

#[owns(&self => readRef(self.data_ptr()))]
impl<'a, T> RwLockReadGuard<'a, T> {}

#[owns(&mut self => writeRef(self.data_ptr()))]
impl<'a, T> RwLockWriteGuard<'a, T> {}

…
```

More capabilities…

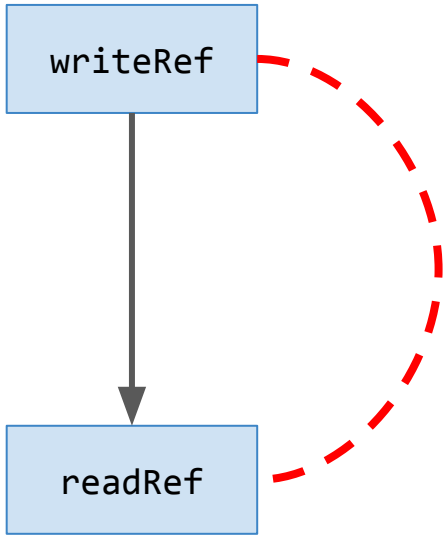

More capabilities…

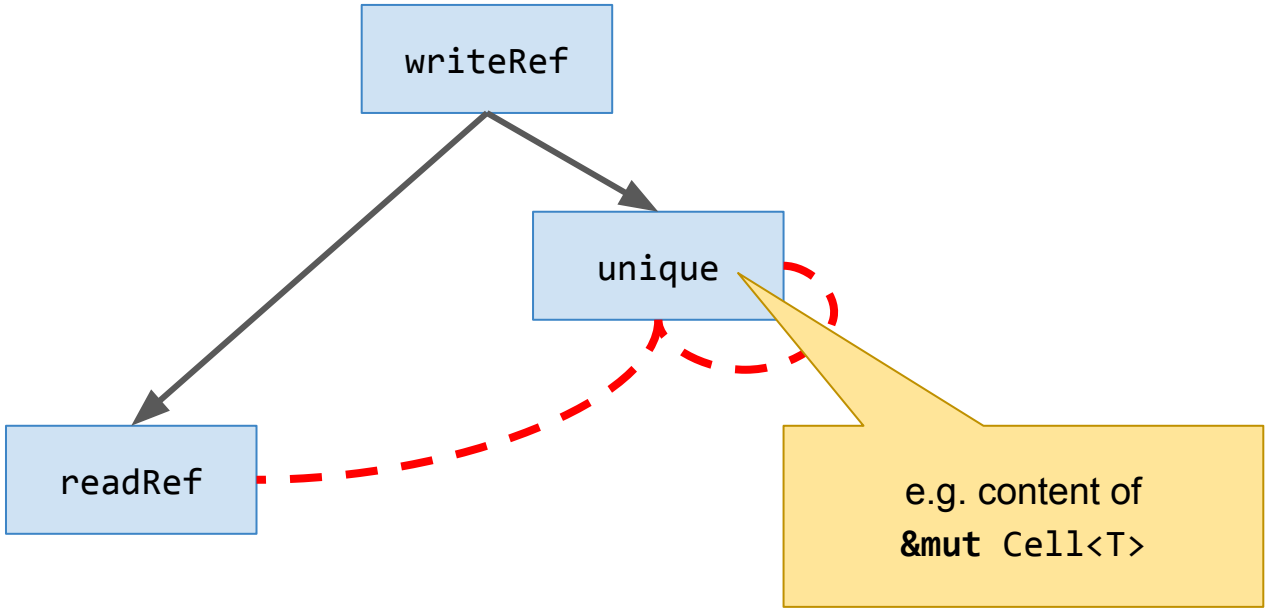

More capabilities…

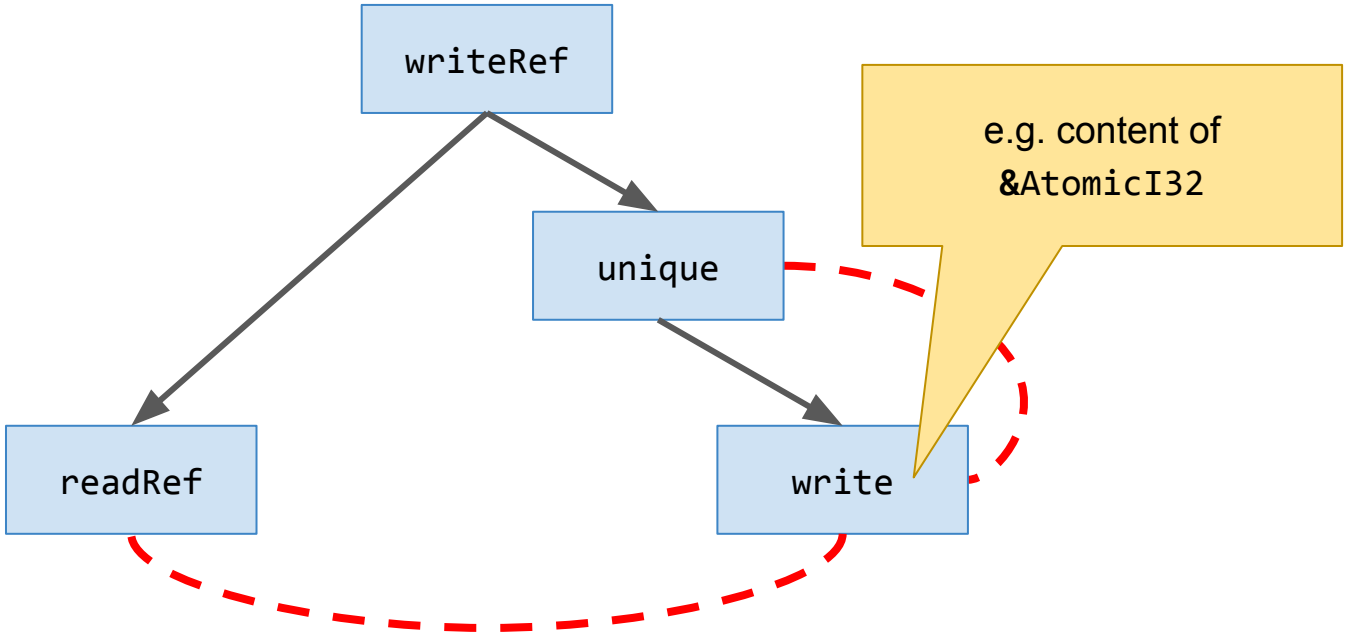

More capabilities…

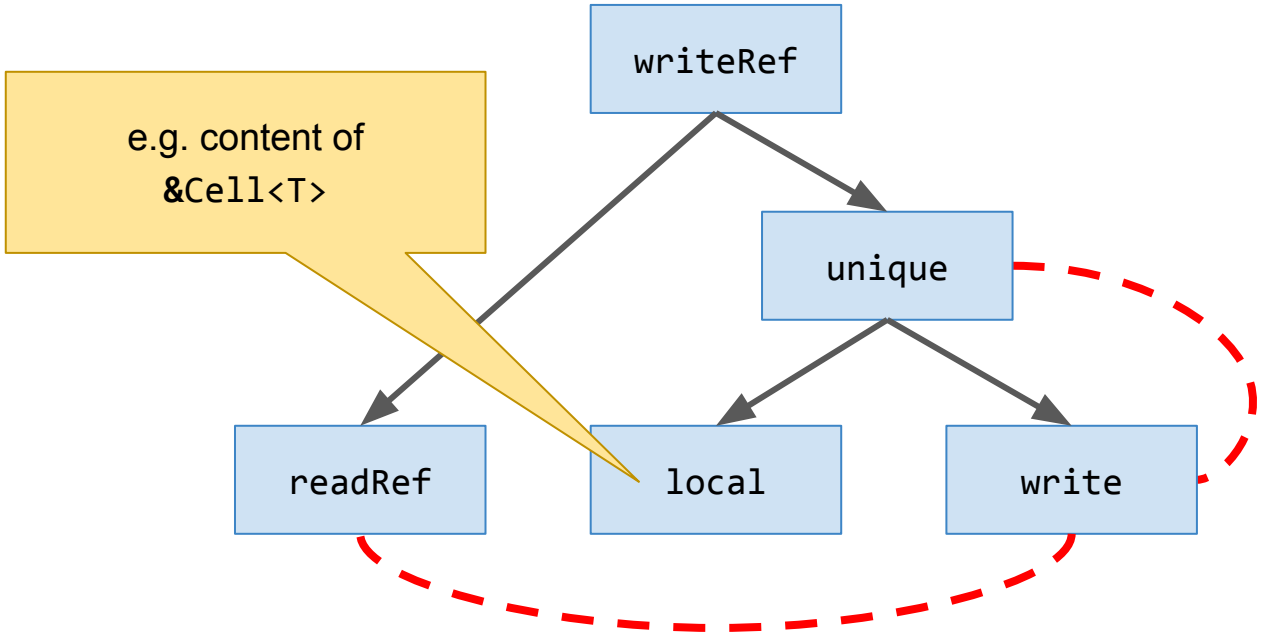

More capabilities…

Shared references whose target is not reachable by other threads.

e.g. y after
`let x: T = …; let y = &x;`

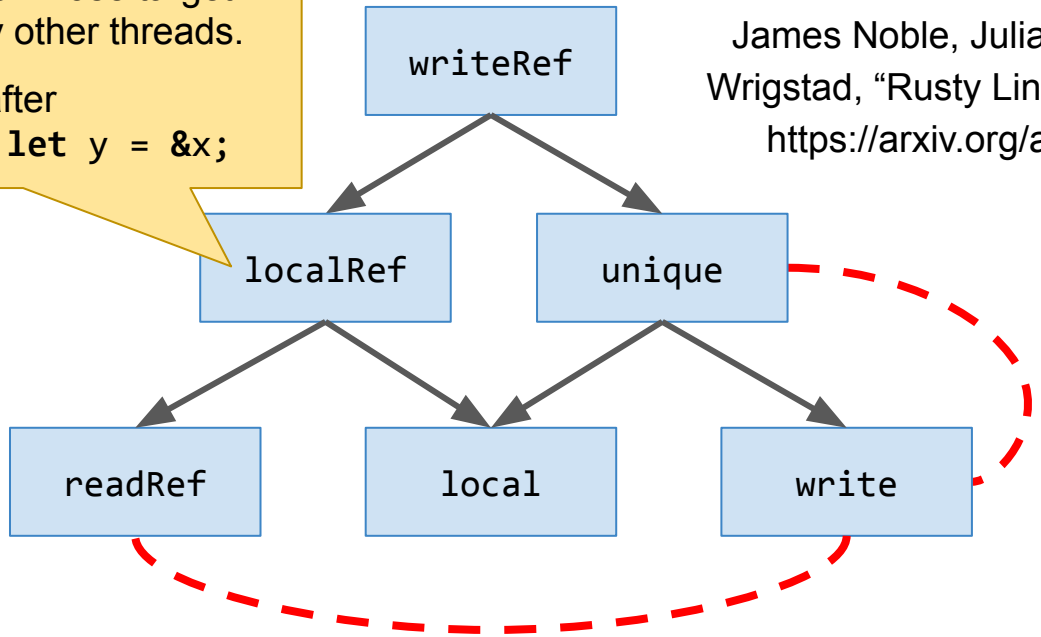

It's the "**`&loc`** `T`" of
James Noble, Julian Mackay, Tobias Wrigstad, "Rusty Links in Local Chains"
https://arxiv.org/abs/2205.00795

More capabilities…

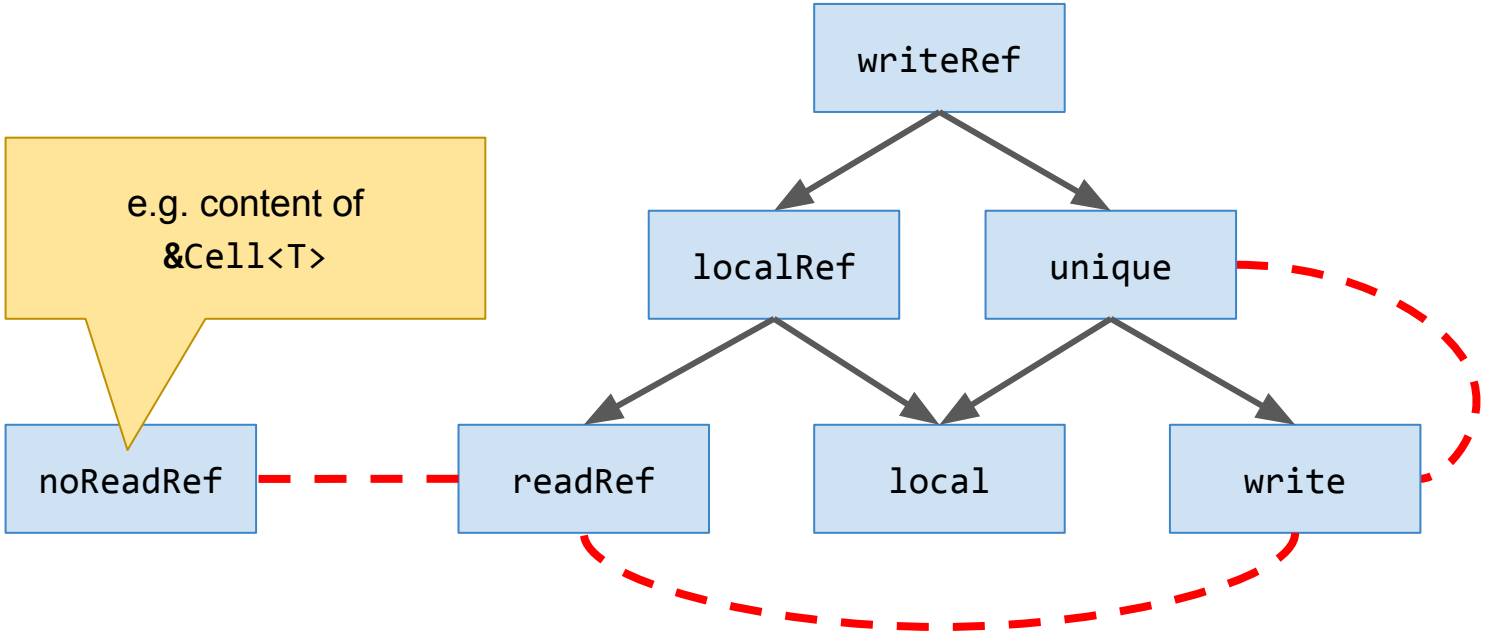

More capabilities…

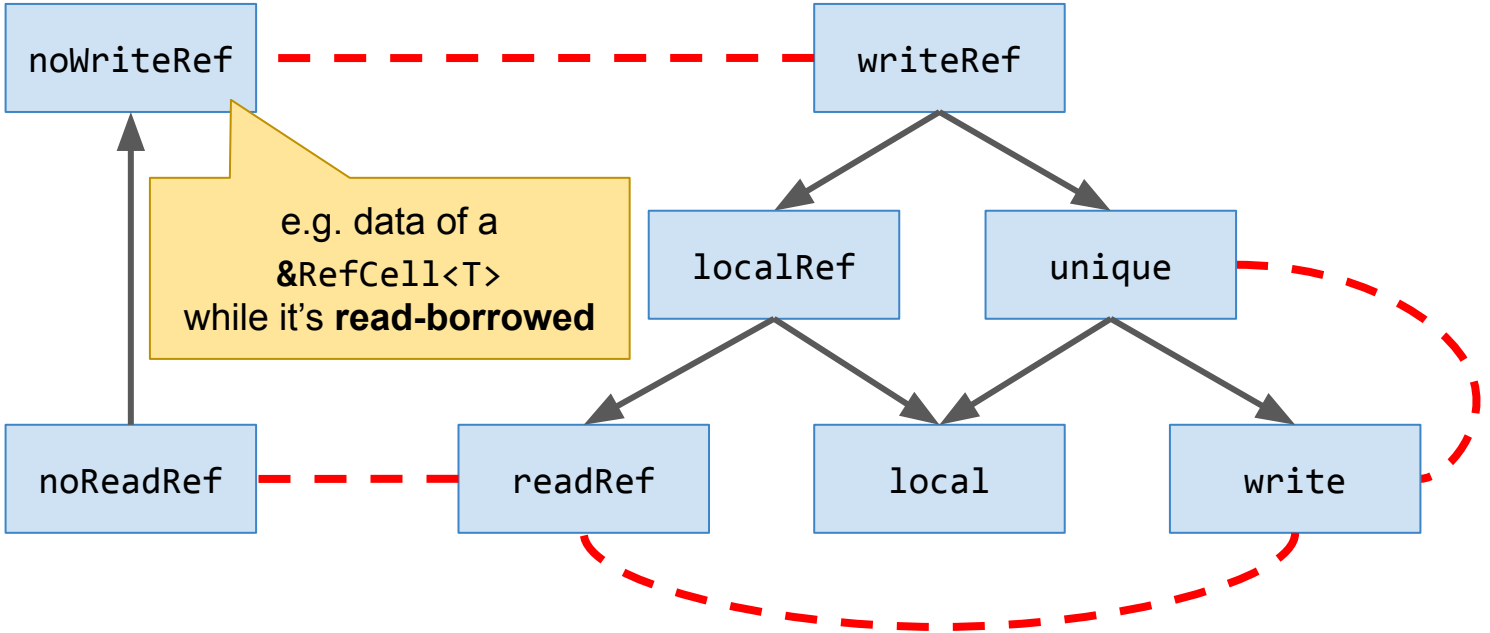

Implicit capabilities with runtime conditions

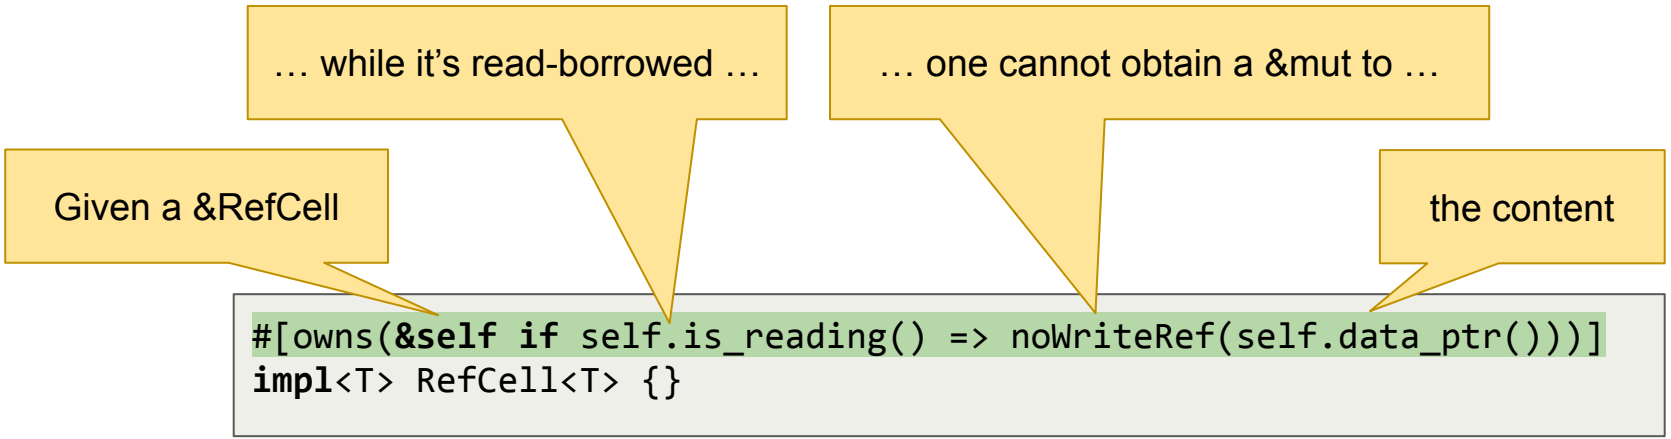

Encoding sketch

Type system && Capability annotations

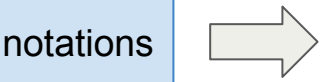

Axioms

```
fn example(x: &RefCell<i32>, y: RefMut<i32>) {
    // t1: {x, y}
    let before: Ref = x.borrow(); // …
    // t2: {x, y, before}
    unknown(x); // unused: {y, before}
    // t3: {x, y, before}
    let after: Ref = x.borrow(); // …
    // t4: {x, y, before, after}
    assert!(...);
}
```

Immutability:

```
content(before, t2)
== content(before, t3)
```

Non-aliasing:

```
content_address(before, t4)
 != content_address(y, t4)
```

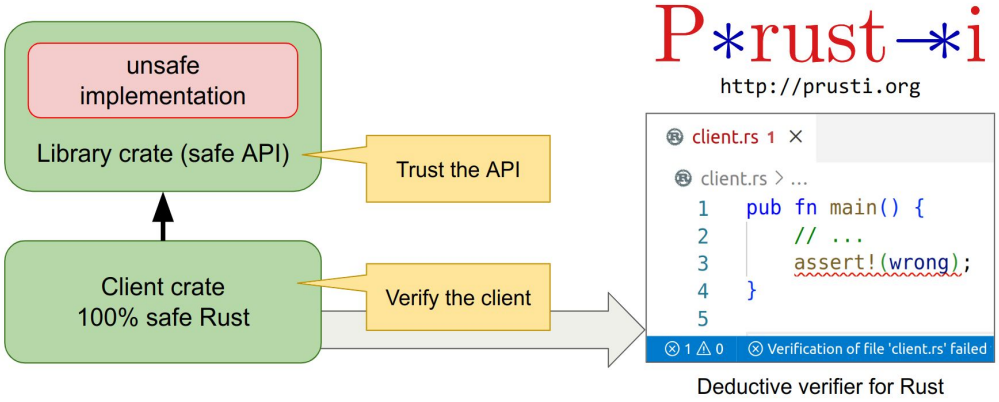

Thanks! Questions?