


Verifying the Rust Standard Library Using Verus

Elanor Tang, Travis Hance (MPI), Chris Hawblitzel (Microsoft),
Natalie Neamtu, Jake Ginesin, and Bryan Parno


Carnegie Mellon University

2025 New England Systems Verification Day

Why Verify the Rust Standard Library?



...except in **unsafe** code



Rust standard library

20
CVEs

Goal: Provide safety and correctness guarantees



Challenges and Tool Motivation


The Rust standard library is...

Evolving


Large
(500,000 SLOC)

Implemented with **unsafe**
code


Verification tools must...



Provide “**Rust-like**”
proof environment



Provide good
automation



Reason about **raw pointer**
operations, among others


Verus does all of this

Rust: Ownership, References, and Borrowing

- Each value in Rust has a variable that's its **owner**.
- Create an alias by making a **reference**.
 - Called **borrowing**. Done with `&` operator.
 - Think of a reference `&T` as **(T, ptr)**.



Enforced by the Rust ***borrow-checker***.



A reference can be **shared** (immutable) or **mutable**, which determines the read/write access allowed.

The **lifetime** of a reference cannot be longer than the owner's lifetime.



Verus: Ownership Ghost Permissions


- **Ownership ghost permissions** track the **evolving state of a resource** in unsafe code.
- Type-check permissions with Rust's **borrow-checker** to ensure safety.

```
unsafe {  
    let (p, Tracked(mut points_to)) = allocate::<u32>(4);  
}
```

Signifies ownership ghost permission

ptr: *mut T
value: T

Track information about what the pointer points to



(example simplified for demonstration purposes)

Ownership Ghost Permissions: Mutability

Mutability of permission reference must **match** mutability of the pointer operation.

- Signature of `ptr_mut_write` requires a **mutable** reference, to ensure exclusive access to memory.

```
unsafe {  
    let (p, Tracked(mut points_to)) = allocate::<u32>(4);  
    ptr_mut_write(p, Tracked(&mut points_to), 5);  
}
```

(example simplified for demonstration purposes)

Ownership Ghost Permissions: Mutability

Mutability of permission reference must *match* mutability of the pointer operation.

- Signature of `ptr_mut_write` requires a *mutable* reference, to ensure exclusive access to memory.
- Enforced by Rust's ***borrow-checker***.

```
unsafe {  
    let (p, Tracked(mut points_to)) = allocate::<u32>(4);  
    ptr_mut_write(p, Tracked(&points_to), 5);      // FAILS  
}
```

(example simplified for demonstration purposes)

Ownership Ghost Permissions: Lifetime

Permission is valid for ***exactly as long*** as the allocation's lifetime.

- Signature of `deallocate` requires ***ownership transfer*** of `points_to`.

```
unsafe {  
    let (p, Tracked(mut points_to)) = allocate::<u32>(4);  
    ptr_mut_write(p, Tracked(&mut points_to), 5);  
    deallocate(p, 4, Tracked(points_to));  
}
```

(example simplified for demonstration purposes)

Ownership Ghost Permissions: Lifetime

Permission is valid for ***exactly as long*** as the allocation's lifetime.

- Signature of `deallocate` requires ***ownership transfer*** of `points_to`.
- Rust borrow-checker ***forbids references*** after that: it is no longer in scope.

```
unsafe {  
    let (p, Tracked(mut points_to)) = allocate::<u32>(4);  
    ptr_mut_write(p, Tracked(&mut points_to), 5);  
    deallocate(p, 4, Tracked(points_to));  
    ptr_mut_write(p, Tracked(&mut points_to), 5); // FAILS  
}
```

(example simplified for demonstration purposes)


Challenge #1: Handling Pointer Provenance

Provenance captures what you are allowed to *do* with a pointer, based on the source *it was derived from*.

Spatial

Temporal

Mutability



Challenge #1: Handling Pointer Provenance


Provenance captures what you are allowed to do with a pointer, based on the source *it was derived from*.

✓ Spatial → Extend ownership ghost permissions with provenance information

✓ Temporal

✓ Mutability

Addressed by Rust's borrow-checker on the lifetime and mutability of ghost permissions




Example: `ptr::add`

```
pub unsafe fn add(ptr: *const T, count: usize) -> *const T
```

Advance `ptr: *const T` by `count` elements of type `T`

Example: `ptr::add`




```
pub unsafe fn add(ptr: *const T, count: usize) -> *const T
```

If the computed offset is non-zero, then

- ➊ `ptr` must have a valid allocation (not freed).
- ➋ Memory range between `ptr` and the result must be within bounds.

Example: `ptr::add`



```
pub unsafe fn add_verus<T>(ptr: *const T, count: usize, Tracked(perm): Tracked<&PointsToRaw>) -> *const T
```

If the computed offset is non-zero, then


- ➊ `ptr` must have a valid allocation (not freed).

provenance: `alloc_start: usize`
`alloc_length: usize`

Tells us that this memory has not been deallocated

- ➋ Memory range between `ptr` and the result must be within bounds.

Example: `ptr::add`



```
pub unsafe fn add_verus<T>(ptr: *const T, count: usize, Tracked(perm): Tracked<&PointsToRaw>) -> *const T
```

- If the computed offset is non-zero, then
- 1 `ptr` must have a valid allocation (not freed).


```
perm.provenance() == ptr.provenance
```

- 2 Memory range between `ptr` and the result must be within bounds.

provenance: `alloc_start: usize`
`alloc_length: usize`

Tells us that this memory has not been deallocated

Example: `ptr::add`



```
pub unsafe fn add_verus<T>(ptr: *const T, count: usize, Tracked(perm): Tracked<&PointsToRaw>) -> *const T
```

- If the computed offset is non-zero, then
- 1 `ptr` must have a valid allocation (not freed).

```
perm.provenance() == ptr.provenance
```


provenance: `alloc_start: usize`
`alloc_length: usize`

Tells us that this memory has not been deallocated

- 2 Memory range between `ptr` and the result must be within bounds.

```
ptr.in_bounds(perm.provenance.alloc_start(),  
              perm.provenance.alloc_start() + perm.provenance.alloc_length(),  
              count)
```

Example: `ptr::add`



```
pub unsafe fn add_verus<T>(ptr: *const T, count: usize, Tracked(perm): Tracked<&PointsToRaw>) -> *const T
```

If the computed offset is non-zero, then

- 1 `ptr` must have a valid allocation (not freed)

The only provenance information we needed to add was `alloc_start` and `alloc_length`

provenance: `alloc_start: usize`
`alloc_length: usize`
memory has not moved

- 2 Memory range between `ptr` and the result must be within bounds.


```
ptr.in_bounds(perm.provenance.alloc_start(),  
              perm.provenance.alloc_start() + perm.provenance.alloc_length(),  
              count)
```

Challenge #2: Shared Reference SMT Encoding

Think of $\&T$ as (T, ptr)

- In most cases, we only care about T

Encoding A




Add `ptr_info(v: &T)` function to get pointer information as needed

➤ ***Simpler for users***

➤ ***Harder to track and update pointer information internally***


Encoding B



Use $\&T$ only when you actually need the pointer


➤ ***Straightforward to implement***

➤ ***Often unavoidable to have $\&T$ when we do not need the pointer***



Challenge #3: Ergonomically Incorporating Spec/Proof Code into Existing Rust Code

Need to use Verus versions of functions



Solution: Support in progress for attribute-based syntax

```
# [with_ghost_arg(Tracked(perm): Tracked<&PointsToRaw>)]
pub unsafe fn add<T>(ptr: *const T, count: usize) -> *const T
```

Challenge #3: Ergonomically Incorporating Spec/Proof Code into Existing Rust Code

Need to use Verus versions of functions

`add(ptr, count)`  `add_verus(ptr, count, Tracked(&perm))`


**Need to do this in a way that still enables type-checking,
so we can keep borrow-checking our permissions**

Solution: Support in progress for attribute-based syntax

```
# [with_ghost_arg(Tracked(perm): Tracked<&PointsToRaw>)]
pub unsafe fn add<T>(ptr: *const T, count: usize) -> *const T
```

Results

- Verified
- Specified
- Unspecified



Results

Verified

Specified

Vec::new

RawVec::new

str::run_utf8_validation

str::split_at_unchecked

str::next_code_point_verus

```
assert(ch == (((x & 0x07) as u32) << 18) |
        (((y & 0x3f) as u32) << 12) |
        (((z & 0x3f) as u32) << 6) |
        ((w & 0x3f) as u32)) by (bit_vector)

requires
  x >= 0xF0,
  init == (x & 0x7Fu8 >> (2 as u8)) as u32,
  y_z == (((y & 0b0011_1111) as u32) << 6) | (z & 0b0011_1111) as u32,
  ch == (init & 7) << 18 | ((y_z << 6) | (w & 0b0011_1111) as u32),
;
```

usize::count_ones

intrinsics::ctpop

unchecked::
<NonZero<T>, T>

provenance
mut

transmute::<Alignment,
NonZero<usize>>

transmute::<usize,
*mut T>

Recap

- Rely on Rust's ***borrow-checker*** and ***ownership ghost types***.
- Straightforward ***pointer provenance*** model.
- Capable of ***verifying complex, real-world code***.

