



Delft University of Technology

## RustBelt: securing the foundations of the rust programming language

Jung, Ralf; Jourdan, Jacques-Henri; Krebbers, Robbert; Dreyer, Derek

**DOI**

[10.1145/3158154](https://doi.org/10.1145/3158154)

**Publication date**

2018

**Document Version**

Final published version

**Published in**

Proceedings of the ACM on Programming Languages

**Citation (APA)**

Jung, R., Jourdan, J.-H., Krebbers, R., & Dreyer, D. (2018). RustBelt: securing the foundations of the rust programming language. *Proceedings of the ACM on Programming Languages*, 2(POPL), 66:1-66:34. Article 66. <https://doi.org/10.1145/3158154>

**Important note**

To cite this publication, please use the final published version (if applicable). Please check the document version above.

**Copyright**

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

**Takedown policy**

Please contact us and provide details if you believe this document breaches copyrights. We will remove access to the work immediately and investigate your claim.



# RustBelt: Securing the Foundations of the Rust Programming Language

RALF JUNG, MPI-SWS, Germany

JACQUES-HENRI JOURDAN, MPI-SWS, Germany

ROBBERT KREBBERS, Delft University of Technology, The Netherlands

DEREK DREYER, MPI-SWS, Germany

Rust is a new systems programming language that promises to overcome the seemingly fundamental tradeoff between high-level safety guarantees and low-level control over resource management. Unfortunately, none of Rust's safety claims have been formally proven, and there is good reason to question whether they actually hold. Specifically, Rust employs a strong, ownership-based type system, but then extends the expressive power of this core type system through libraries that internally use unsafe features. In this paper, we give the first formal (and machine-checked) safety proof for a language representing a realistic subset of Rust. Our proof is extensible in the sense that, for each new Rust library that uses unsafe features, we can say what verification condition it must satisfy in order for it to be deemed a safe extension to the language. We have carried out this verification for some of the most important libraries that are used throughout the Rust ecosystem.

CCS Concepts: • **Theory of computation** → Programming logic; Separation logic; Operational semantics;

Additional Key Words and Phrases: Rust, separation logic, type systems, logical relations, concurrency

## ACM Reference Format:

Ralf Jung, Jacques-Henri Jourdan, Robbert Krebbers, and Derek Dreyer. 2018. RustBelt: Securing the Foundations of the Rust Programming Language. *Proc. ACM Program. Lang.* 2, POPL, Article 66 (January 2018), 34 pages. <https://doi.org/10.1145/3158154>

## 1 INTRODUCTION

Systems programming languages like C and C++ give programmers low-level control over resource management at the expense of safety, whereas most other modern languages give programmers safe, high-level abstractions at the expense of control. It has long been a “holy grail” of programming languages research to overcome this seemingly fundamental tradeoff and design a language that offers programmers both high-level safety *and* low-level control.

Rust [Matsakis and Klock II 2014; Rust team 2017], developed at Mozilla Research, comes closer to achieving this holy grail than any other industrially supported programming language to date. On the one hand, like C++, Rust supports zero-cost abstractions for many common systems programming idioms and avoids dependence on a garbage collector [Stroustrup 2012; Turon 2015a]. On the other hand, like most modern high-level languages, Rust is type-safe and memory-safe. Furthermore, Rust's type system goes beyond that of the vast majority of safe languages in that it statically rules out data races (which are a form of undefined behavior for concurrent programs in many languages like C++ or Rust), as well as common programming pitfalls like

---

Authors' addresses: Ralf Jung, MPI-SWS\*, [jung@mpi-sws.org](mailto:jung@mpi-sws.org); Jacques-Henri Jourdan, MPI-SWS\*, [jjourdan@mpi-sws.org](mailto:jjourdan@mpi-sws.org); Robbert Krebbers, Delft University of Technology, [mail@robbertkrebbers.nl](mailto:mail@robbertkrebbers.nl); Derek Dreyer, MPI-SWS\*, [dreyer@mpi-sws.org](mailto:dreyer@mpi-sws.org).  
\* Saarland Informatics Campus.



This work is licensed under a Creative Commons Attribution 4.0 International License.

© 2018 Copyright held by the owner/author(s).

2475-1421/2018/1-ART66

<https://doi.org/10.1145/3158154>

iterator invalidation [Gregor and Schupp 2003]. In other words, compared to mainstream “safe” languages, Rust offers both lower-level control *and* stronger safety guarantees.

*At least, that is the hope.* Unfortunately, none of Rust’s safety claims have been formally proven, and there is good reason to question whether they actually hold. In this paper, we make a major step toward rectifying this situation by giving the first formal (and machine-checked) safety proof for a language representing a realistic subset of Rust. Before explaining our contributions in more detail, and in particular what we mean here by “realistic”, let us begin by exploring what makes Rust’s type system so unusual, and its safety so challenging to verify.

### 1.1 Rust’s “Extensible” Approach to Safe Systems Programming

At the heart of the Rust type system is an idea that has emerged in recent years as a unifying concept connecting both academic and mainstream language design: *ownership*. In its simplest form, the idea of ownership is that, although multiple aliases to a resource may exist simultaneously, performing certain actions on the resource (such as reading and writing a memory location) should require a “right” or “capability” that is uniquely “owned” by one alias at any point during the execution of the program. Although the right is uniquely owned, it can be “transferred” from one alias to another—e.g., upon calling a function or spawning a thread, or via synchronization mechanisms like locks. In more complex variations, ownership can be *shared* between aliases, but only in a controlled manner (e.g., shared ownership only permits read access [Boyland 2003]). In this way, ownership allows one to carefully administer the safe usage of potentially aliased resources.

Ownership pervades both academic and mainstream language design for safe(r) systems programming. On the academic side, many proposals have been put forth for using *types* to enforce various ownership disciplines, including “ownership type” systems [Clarke et al. 1998]; region- or typestate-based systems for “safe C” programming in languages like Cyclone [Jim et al. 2002] and Vault [DeLine and Fähndrich 2001]; and substructural type systems like Ynot [Nanevski et al. 2008], Alms [Tov and Pucella 2011], and Mezzo [Balabonski et al. 2016]. Unfortunately, although these languages provide strong safety guarantees, none of them have made it out of academic research into mainstream use. On the mainstream side, “modern C++” (i.e., C++ since the 2011 standard [ISO Working Group 21 2011]) provides several features—e.g., smart pointers, move semantics, and RAII (Resource Acquisition Is Initialization)—that are essentially mechanisms for controlling ownership. However, while these features encourage safer programming idioms, the type system of C++ is too weak to enforce its ownership disciplines statically, so it is still easy to write programs with unsafe or undefined behavior using these features.

In some sense, the key challenge in developing sound static enforcement of ownership disciplines—and the reason perhaps that academic efforts have not taken off in practice—is that no language can account for the safety of every advanced form of low-level programming that one finds in the wild, because there is no practical way to do so while retaining automatic type checking. As a result, previous designs employ type systems that are either too restrictive (i.e., preventing programmers from writing certain kinds of low-level code they want to write) or too expressive (i.e., encoding such a rich logic in the type structure that programmers must do proofs to appease the type checker).

Rust addresses this challenge by taking a hybrid, *extensible* approach to ownership.

The basic ownership discipline enforced by Rust’s type system is a simple one: If ownership of an object (of type  $T$ ) is shared between multiple aliases (“shared references” of type  $\&T$ ), then none of them can be used to directly mutate it. This discipline, which is similar in spirit to (if different in detail from) that of several prior academic approaches, is enforceable automatically and eliminates a wide range of common low-level programming errors, such as “use after free”, data races, and iterator invalidation. However, it is also too restrictive to account for many low-level

data structures and synchronization mechanisms, which fundamentally depend on the ability to mutate aliased state (e.g., to implement mutual exclusion or communication between threads).

Consequently, to overcome this restriction, the implementations of Rust’s standard libraries make widespread use of `unsafe` operations, such as “raw pointer” manipulations for which aliasing is not tracked. The developers of these libraries claim that their uses of `unsafe` code have been properly “encapsulated”, meaning that if programmers make use of the APIs exported by these libraries but otherwise avoid the use of `unsafe` operations themselves, then their programs should never exhibit any unsafe/undefined behaviors. In effect, these libraries extend the expressive power of Rust’s type system by loosening its ownership discipline on aliased mutable state in a modular, controlled fashion: Even though a shared reference of type `&T` may not be used to *directly* mutate the contents of the reference, it may nonetheless be used to *indirectly* mutate them by passing it to one of the observably “safe” (but internally `unsafe`) methods exported by the object’s API.

However, there is cause for concern about whether Rust’s extensible approach is actually sound. Over the past few years, several soundness bugs have been found in Rust, both in the type system itself [Ben-Yehuda 2015a,b; Turon 2015b] and in libraries that use `unsafe` code [Ben-Yehuda 2015c; Biocca 2017; Jung 2017]. Some of these—such as the Leakpocalypse bug [Ben-Yehuda 2015c]—are quite subtle in that they involve an *interaction* of multiple libraries, each of which is (or seems to be) perfectly safe on its own. To make matters worse, the problem cannot easily be contained by blessing a fixed set of standard libraries as primitive and just verifying the soundness of those; for although it is considered a badge of honor for Rust programmers to avoid the use of `unsafe` code entirely, many nevertheless find it necessary to employ a sprinkling of `unsafe` code in their developments. Of course, it is not unusual for safe languages to provide unsafe escape hatches (e.g., Haskell’s `unsafePerformIO`, OCaml’s `Obj.magic`) to work around limitations of their type systems. But `unsafe` code plays such a fundamental role in Rust’s extensible ownership discipline that it cannot simply be swept aside if one wishes to give a realistic formal account of the language.

The question remains: How can we verify that Rust’s extensible approach makes any sense? The standard technique for proving safety properties for high-level programming languages—namely, “progress and preservation” introduced by Wright and Felleisen [1994]—does not apply to languages in which one can mix safe and unsafe code. (Progress and preservation is a *closed-world* method, which assumes the use of a closed set of typing rules. This assumption is fundamentally violated by Rust’s extensible, open-world approach.) So, to account for safe-unsafe interaction, we need a way to specify formally what we are obliged to *prove* if we want to establish that a library employing `unsafe` code constitutes a sound extension of the Rust type system. Luckily, decades of research in semantics and verification have provided us with just the right tools for the job.

## 1.2 RustBelt: An Extensible, Semantic Approach to Proving Soundness of Rust

In this paper, we give the first formal (and machine-checked) account of Rust’s extensible approach to safe systems programming and how to prove it sound.

For obvious reasons of scale, we do not consider the full Rust language, for which no formal description exists anyway. Instead, after beginning (in §2) with an example-driven tour of the most central and distinctive features of the Rust type system, we proceed (in §3) to describe  $\lambda_{\text{Rust}}$ , a continuation-passing style language (of our own design) that formalizes the static and dynamic semantics of these central features. Crucially,  $\lambda_{\text{Rust}}$  incorporates Rust’s notions of *borrowing*, *lifetimes*, and *lifetime inclusion*—which are fundamental to Rust’s ownership discipline—in a manner inspired by Rust’s Mid-level Intermediate Representation (MIR). For simplicity,  $\lambda_{\text{Rust}}$  omits some orthogonal features of Rust such as traits (which are akin to Haskell type classes); it also avoids the morass of exciting complications concerning relaxed memory, instead adopting a simplified memory model featuring only non-atomic and sequentially consistent atomic operations. Nevertheless,  $\lambda_{\text{Rust}}$

is realistic enough that studying it led us to uncover a previously unknown soundness bug in Rust itself [Jung 2017].

Our core contribution is then to develop an *extensible soundness proof* for  $\lambda_{\text{Rust}}$ . The basic idea is to build a *semantic model* of the language—in particular, a *logical relation* [Plotkin 1973; Tait 1967]. The idea of proving soundness semantically is hardly new [Milner 1978], but it fell out of favor after Wright and Felleisen [1994] developed their simpler “syntactic” proof method. The semantic approach to type soundness is more powerful than the syntactic approach, however, because it offers an interpretation of what types *mean* (i.e., what terms inhabit them) that is more general than just “what the syntactic typing rules allow”—it describes when it is *observably* safe to treat a term as having a certain type, even if syntactically that term employs *unsafe* features. Moreover, thanks to the Foundational Proof-Carrying Code project [Ahmed et al. 2010; Appel 2001] and the development of “step-indexed” logical relations [Ahmed 2004; Appel and McAllester 2001] which arose from that project, we now know how to scale the semantic approach to languages with semantically complex features like recursive types and higher-order state.

Here, we follow the style of recent “logical” accounts of step-indexed logical relations [Dreyer et al. 2011, 2010; Krogh-Jespersen et al. 2017; Turon et al. 2013], interpreting  $\lambda_{\text{Rust}}$  types as predicates on values expressed in a rich program logic (see §4 and Challenge #1 below), and interpreting  $\lambda_{\text{Rust}}$  typing judgments as logical entailments between these predicates (see §7). With our semantic model—which we call **RustBelt**—in hand, the proof of safety of  $\lambda_{\text{Rust}}$  divides into three parts:

- (1) Verify that the typing rules of  $\lambda_{\text{Rust}}$  are sound when interpreted semantically, i.e., as lemmas establishing that the semantic interpretations of the premises imply the semantic interpretation of the conclusion. This is called *the fundamental theorem of logical relations*.
- (2) Verify that, if a closed program is *semantically* well-typed according to the model, its execution will not exhibit any *unsafe/undefined* behaviors. This is called *adequacy*.
- (3) For any library that employs *unsafe* code internally, verify that its implementation satisfies the predicate associated with the semantic interpretation of its interface, thus establishing that the *unsafe* code has indeed been safely “encapsulated” by the library’s API. In essence, the semantic interpretation of the interface yields a *library-specific verification condition*.

Together, these ensure that, so long as the only *unsafe* code in a well-typed  $\lambda_{\text{Rust}}$  program is confined to libraries that satisfy their verification conditions, the program is safe to execute.

This proof is “extensible” in the sense, that whenever you have a new library that uses *unsafe* code and that you want to verify as being safe to use in Rust programs, RustBelt tells you the verification condition you need to prove about it. Using the Coq proof assistant [Coq team 2017], we have formally proven the fundamental theorem and adequacy once and for all, and we have also proven the verification conditions for ( $\lambda_{\text{Rust}}$  ports of) several standard Rust libraries that use *unsafe* code, including `Arc`, `Rc`, `Cell`, `RefCell`, `Mutex`, `RwLock`, `mem::swap`, `thread::spawn`, `rayon::join`, and `take_mut`.

Although the high-level structure of our soundness proof is standard [Ahmed 2004; Milner 1978], developing such a proof for a language as subtle and sophisticated as Rust has required us to tackle a variety of technical challenges, more than we can describe in the space of this paper. To focus the presentation, we will therefore not present all these challenges and their solutions in full technical detail (although further details can be found in our technical appendix and Coq development [Jung et al. 2017a]). Rather, we aim to highlight the following key challenges and how we dealt with them.

**Challenge #1: Choosing the right logic for modeling Rust.** The most fundamental design choice in RustBelt was deciding which logic to use as its target, i.e., for defining semantic interpretations of Rust types. There are several desiderata for such a logic, but the most important is that it

should support high-level reasoning about concepts that are central to Rust’s type system, such as ownership and borrowing. The logic we chose, **Iris**, is ideally suited to this purpose.

Iris is a language-generic framework for higher-order concurrent separation logic [Jung et al. 2016, 2017b, 2015; Krebbers et al. 2017a], which in the past year has been equipped with tactical support for conducting machine-checked proofs of programs in Coq [Krebbers et al. 2017b] and deployed in several ongoing verification projects [Kaiser et al. 2017; Swasey et al. 2017; Tassarotti et al. 2017; Timany et al. 2018]. By virtue of being a separation logic [O’Hearn 2007; Reynolds 2002], Iris comes with built-in support for reasoning modularly about ownership. Moreover, the main selling point of Iris is its support for *deriving* custom program logics for different domains using only a small set of primitive mechanisms (namely, *higher-order ghost state* and *impredicative invariants*). In the case of RustBelt, we used Iris to derive a novel **lifetime logic**, whose primary feature is a notion of *borrow propositions* that mirrors the “borrowing” mechanism for tracking aliasing in Rust. This lifetime logic, which we describe in some detail in §5, has made it possible for us to give fairly direct interpretations of a number of Rust’s most semantically complex types, and to verify their soundness at a high level of abstraction.

**Challenge #2: Modeling Rust’s extensible ownership discipline.** As explained above, a distinctive feature of Rust is its extensible ownership discipline: Owning a value of shared reference type  $\&T$  confers different privileges depending on the type  $T$ . For many simple types,  $\&T$  confers read-only access to the contents of the reference; but for types defined by libraries that use **unsafe** operations,  $\&T$  may in fact confer mutable access to the contents, indirectly via the API of  $T$ . In Rust lingo, this phenomenon is termed *interior mutability*.

To model interior mutability, RustBelt interprets types  $T$  in two ways: (1) with an *ownership* predicate that says what it means to own a value of type  $T$ , and (2) with a *sharing* predicate that says what it means to own a value of type  $\&T$ . Unlike the ownership predicate, the sharing predicate must be a freely duplicable assertion, since Rust allows values of shared reference type to be freely copied. But otherwise there is a great deal of freedom in how it is defined, thus allowing us to assign very different semantics to  $\&T$  for different types  $T$ . We exploit this freedom in proving semantic soundness of several Rust libraries whose types exhibit interior mutability (see §6).

**Challenge #3: Accounting for Rust’s “thread-safe” type bounds.** Some of Rust’s types that exhibit interior mutability use non-atomic rather than atomic memory accesses to improve performance. As a result, however, they are not “thread-safe”, meaning that if one could transfer ownership of values of these types between threads, it could cause a data race. Rust handles this potential safety problem by restricting cross-thread ownership transfer to types that satisfy certain *type bounds*: the **Send** bound classifies types  $T$  that are thread-safe, and the **Sync** bound classifies types  $T$  such that  $\&T$  is thread-safe.

We account for these type bounds in RustBelt in a simple and novel way. First, we parameterize both the ownership and sharing predicates in the semantics of types by a *thread identifier*, representing the thread that is claiming ownership. We then define  $T$  to be **Send** if  $T$ ’s ownership predicate does not depend on the thread id parameter (and **Sync** if  $T$ ’s sharing predicate does not depend on the thread id parameter). Intuitively, this makes sense because, if ownership of a value  $v$  of type  $T$  is thread-independent, transferring ownership of  $v$  between threads is perfectly safe.

All results in this paper have been fully formalized in the Coq proof assistant [Jung et al. 2017a].

## 2 A TOUR OF RUST

In this section, we give a brief overview of some of the central features of the Rust type system. We do not assume the reader has any prior familiarity with Rust.

## 2.1 Ownership Transfer

A core feature of the Rust type system is to provide thread-safety, *i.e.*, to guarantee the absence of unsynchronized race conditions. Race conditions can only arise from an unrestricted combination of *aliasing* and *mutation* on the same location. In fact, it turns out that ruling out mutation of aliased data also prevents other errors commonplace in low-level pointer-manipulating programs, like use-after-free or double-free. The essential idea of the Rust type system is thus to ensure that aliasing and mutation cannot occur at the same time on any given location, which it achieves by letting *types represent ownership*.

Let us begin with the most basic form of ownership, *exclusive ownership*, in which, at any time, at most one thread is allowed to mutate a given location. Exclusive ownership rules out aliasing entirely, and thus prevents data races. However, just exclusive ownership would not be very expressive, and therefore Rust allows one to *transfer* ownership between threads. To see this principle in practice, consider the following sample program:

```

1  let (snd, rcv) = channel();
2  join(move || {
3      let mut v = Vec::new(); v.push(0); // v: Vec<i32>
4      snd.send(v);
5      // Cannot access v: v.push(1) rejected
6  },
7  move || {
8      let v = rcv.recv().unwrap(); // v: Vec<i32>
9      println!("Received: {:?}", v);
10 });

```

Before we take a detailed look at the way the Rust type system handles ownership here, we briefly discuss syntax: **let** is used to introduce local, stack-allocated variables. These can be made mutable by using **let mut**. The first line uses a pattern to immediately destruct the pair returned by `channel()` into its components. The vertical bars `||` mark the beginning of an anonymous closure; if the closure would take arguments, they would be declared between the bars.

In this example, one thread sends a shallow copy (*i.e.*, not duplicating data behind pointer indirections) of a vector `v` of type `Vec<i32>` (a resizable heap-allocated array of 32-bit signed integers) over a channel to another thread. In Rust, having a value of some type indicates that we are the *exclusive owner* of the data described by said type, and thus that *nobody else* has any kind of access to this array, *i.e.*, no other part of the program can write to or even read from the array. When ownership is passed to a function (*e.g.*, `send`), the function receives a shallow copy of the data.<sup>1</sup> At the same time, ownership of the data is considered to have *moved*, and thus no longer available in the callee—thus, Rust’s variable context is substructural. This is important because the receiver only receives a shallow copy, so if both threads were to use the vector, they could end up racing on the same data.

The function `channel` creates a typed multi-producer single-consumer channel and returns the two endpoints as a pair. The function `join` is essentially parallel composition; it takes two closures and executes them in parallel, returning when both are done.<sup>2</sup> The keyword `move` instructs the type checker to move exclusive ownership of the sending end `snd` and receiving end `rcv` of the channel into the first, and, respectively, second closure.

In this example, the first thread creates a new empty `Vec`, `v`, and pushes an element onto it. Next, it sends `v` over the channel. The `send` function takes type `Vec<i32>` as argument, so the Rust type

<sup>1</sup>Of course, Rust provides a way to do a *deep copy* that actually duplicates the vector, but it will never do this implicitly.

<sup>2</sup>`join` is not in the Rust standard library, but part of Rayon [Stone and Matsakis 2017], a library for parallel list processing.

checker considers `v` to be *moved* after the call to `send`. Any further attempts to access `v` would thus result in a compile-time error. The second thread works on the receiving end of the channel. It uses `recv` in order to receive (ownership of) the vector. However, `recv` is a fallible operation, so we call `unwrap` to trigger a *panic* (which aborts execution of the current thread) in case of failure. Finally, we print a debug representation of the vector (as indicated by the format string "{:?}",).

One aspect of low-level programming that is distinctively absent in the code above is memory management. Rust does not have garbage collection, so it may seem like our example program leaks memory, but that is not actually the case: Due to ownership tracking, Rust can tell when a variable (say, the vector `v`) goes out of scope without having been moved elsewhere. When that is the case, the compiler automatically inserts calls to a *destructor*, called `drop` in Rust. For example, when the second thread finishes in line 10, `v` is dropped. Similarly, the sending and receiving ends of the channel are dropped at the end of their closures. This way, Rust provides automatic memory management without garbage collection, and with predictable runtime behavior.

## 2.2 Mutable References

Ownership transfer is a fairly straightforward mechanism for ensuring data-race freedom and related memory safety properties. However, it is also very restrictive. In fact, close inspection shows that even our first sample program does not strictly follow this discipline. Observe that in line 3, we are calling the method `push` on the vector `v`—and we keep using `v` afterwards. Indeed, it would be very inconvenient if pushing onto a vector required explicitly passing ownership to `push` and back. Rust’s solution to this issue is *borrowing*, which is the mechanism used to handle reference types. The idea is that `v` is not moved to `push`, but instead borrowed, *i.e.*, passed by reference—granting `push` access to `v` for the duration of the function call.

This is expressed in the type of `push: fn(&mut Vec<i32>, i32) -> ()`. (Henceforth, we follow the usual Rust style and omit the return type if it is the unit type `()`.) The syntax `v.push(0)`, as used in the example, is just syntactic sugar for `Vec::push(&mut v, 0)`, where `&mut v` creates a mutable reference to `v`, which is then passed to `push`. A mutable reference grants *temporary exclusive access* to the vector, which in the example means that access is restricted to the duration of the call to `push`. Because the access is temporary, our program can keep using `v` when `push` returns. Moreover, the exclusive nature of this access guarantees that no other party will access the vector in any way during the function call, and that `push` cannot keep copies of the pointer to the vector. Mutable references are always unique pointers.

The type of `send, fn(&mut Sender<Vec<i32>>, Vec<i32>)`, shows another use of mutable references. The first argument is just borrowed, so the caller can use the channel again later. In contrast, the second argument is moved, using ownership transfer as already described above.

## 2.3 Shared References

Rust’s approach to guaranteeing the absence of races and other memory safety is to rule out the combination of aliasing and mutation. So far, we have seen unique ownership (§2.1) and (borrowed) mutable references (§2.2), both of which allow for mutation but prohibit aliasing. In this section we discuss another form of references, namely *shared references*, which form the dual to mutable references: They allow aliasing but prohibit mutation.

Like mutable references, shared references grant *temporary* access to a data structure, and operationally correspond to just pointers. The difference is in the guarantees and permissions provided to the receiver of the reference. While mutable references are exclusive (non-duplicable), shared references can be *duplicated*. In other words, shared references permit aliasing. As a consequence, to ensure data-race freedom and memory safety, shared references are *read-only*.

Practically speaking, shared references behave like unrestricted variables in linear type systems, *i.e.*, just like integers, they can be “copied” (as opposed to just being “moved”, which is possible



with variables of all types). Rust expresses such properties of types using *bounds*, and the bound that describes unrestricted types is called *Copy*. Specifically, if a type is *Copy*, it means that doing a shallow copy (which, remember, is what Rust does to pass arguments) suffices to duplicate elements of the type. Both `&T` and `i32` are *Copy* (for any `T`)—however, `Vec<i32>` is not! The reason for this is that `Vec<i32>` stores data on the heap, and a shallow copy does not duplicate this heap data.

We can see shared references in action in the following example:

```
1 let mut v = Vec::new(); v.push(1);
2 join(|| println!("Thread 1: {:?}", &v), || println!("Thread 2: {:?}", &v));
3 v.push(2);
```

This program starts by creating and initializing a vector `v`. It uses a shared reference `&v` to the vector in two threads, which concurrently print the contents of the vector. This time, the closures are not marked as *move*, which leads to `v` being captured by-reference, *i.e.*, at type `&Vec<i32>`. As discussed above, this type is *Copy*, so the type checker accepts using `&v` in both threads.

The concurrent accesses to `v` use *non-atomic* reads, which have no synchronization. This is safe because when a function holds a shared reference, it can rely on the data-structure not being mutated—so there cannot be any data races. (Notice that this is a much stronger guarantee than what C provides with `const` pointers: In C, `const` pointers prevent mutation by the current function, however, they do *not* rule out mutation by *other* functions.)

Finally, when `join` returns, the example program re-gains full access to the vector `v` and can mutate `v` again in line 3. This is safe because `join` will only return when both threads have finished their work, so there cannot be a race between the `push` and the `println`. This demonstrates that shared references are powerful enough to *temporarily* share a data structure and permit unrestricted copying of the pointer, but regain exclusive access later.

## 2.4 Lifetimes

As previously explained, (mutable and shared) references *borrow* ownership and thus grant *temporary* access to a data structure. This immediately raises the question: “How long is *temporary*?” In Rust, this question is answered by equipping every reference with a *lifetime*. The full form of a reference type is actually `&'a mut T` or `&'a T`, where `'a` is the lifetime of the reference. Rust uses a few conventions so that lifetimes can be elided in general, which is why they did not show up in the programs and types we considered so far. However, lifetimes play a crucial role in explaining what happens when the following function is type-checked:

```
1 fn example(v: &'a mut Vec<i32>) {
2     v.push(21);                               Lifetime 'c
3     { let mut head : &'b mut i32 = v.index_mut(0);
4       // Cannot access v: v.push(2) rejected
5       *head = 23; }                           Lifetime 'b
6     v.push(42);
7     println!("{:?}", v); // Prints [23, ..., 42]
8 }
```

Here we define a function `example` that takes an argument of type `&mut Vec<i32>`. The function uses `index_mut` to obtain a pointer to the first element inside the vector. Writing to `head` in line 5 changes the first element of the vector, as witnessed by the output in line 7. Such pointers directly into a data structure are sometimes called *deep* or *interior* pointers. One has to be careful when using deep pointers because they are a form of aliasing: When `v` is deallocated, `head` becomes a dangling pointer. In fact, depending on the data structure, *any* modification of the data structure could lead to deep pointers being invalidated. (One infamous instance of this issue is iterator invalidation,

troubling not only low-level languages like C++, but also safe languages like Java.) This is why the call to `push` in line 4 is rejected.

How does Rust manage to detect this problem and reject line 4 above? To understand this, we have to look at the type of `index_mut`: `for<'b> fn(&'b mut Vec<i32>, usize) -> &'b mut i32`.<sup>3</sup> The `for` is a universal quantifier, making `index_mut` generic in the lifetime `'b`. The caller can use any `'b` to instantiate this generic function, limited only by an implicit requirement that `'b` must last at least as long as the call to `index_mut`. Crucially, `'b` is used for *both* the reference passed to the function *and* the reference returned.

In our example, Rust has to infer the lifetime `'b` left implicit when calling `index_mut`. Because the result of `index_mut` is stored in `head`, the type checker infers `'b` to be the scope of `head`, *i.e.*, lines 3-5. As a consequence, based on the type of `index_mut`, the vector must be borrowed *for the same lifetime*. So Rust knows that `v` is mutably borrowed for lines 3-5, which makes the access in line 4 invalid: The lifetime of the reference needed by `push` would overlap with the lifetime of the reference passed to `index_mut`, which violates the rule that mutable references must be unique.

Lifetimes were not visible in the examples discussed so far, but they are always present implicitly. For example, the full type of `push` is given by `for<'c> push(&'c mut Vec<i32>, i32)`. The type checker thus has the freedom to pick *any* lifetime for the reference to the vector, constrained only by the implicit requirement that `'c` has to cover at least the duration of the function call. This is why the vector can be used again immediately after `push` returned.

Notice that, unlike in the previous examples, `v` in this example is just a mutable reference to begin with. Just like `push`, the type of `example` actually involves a generic lifetime `'a`, and `v` has type `&'a mut Vec<i32>`. Despite not being the original owner of `v`, we can still borrow `v` to someone else—a phenomenon dubbed *reborrowing*. All we have to check is that the reborrow ends *before* the lifetime of our reference ends. In other words, the lifetime of the reborrow (the `'b` used for `index_mut`, *i.e.*, the scope of `head`) has to be *included* in the lifetime of the reference (`'a`). In this case, we know this to be true by making use of the implicit assumption that `'a` includes this function call, so in particular, it includes `'b`, which is entirely contained within the function call.

## 2.5 Interior Mutability

So far, we have seen how Rust ensures memory safety and data-race freedom by ruling out the combination of aliasing and mutation. However, there are cases where shared mutable state is actually needed to (efficiently) implement an algorithm or a data structure. To support these use-cases, Rust provides some primitives providing shared mutable state. All of these have in common that they permit *mutation* through a *shared reference*—a concept called *interior mutability*.

At this point, you may be wondering—how does this fit together with the story of mutation and aliasing being the root of all memory and thread safety problems? The key point is that these primitives have a carefully controlled API surface. Even though mutation through a shared reference is unsafe *in general*, it can still be safe when appropriate restrictions are enforced by either static or run-time checks. This is where we can see Rust’s “extensible” approach to safety in action. Interior mutability is not wired into the type system; instead, the types we are discussing here are implemented in the standard library using `unsafe` code (which we will verify in §6).

**2.5.1 Cell.** The simplest type with interior mutability is `Cell`. Consider the following example:

```
1 let c1 : &Cell <i32> = &Cell::new(0);
2 let c2 : &Cell <i32> = c1;
3 c1.set(2);
4 println!("{:?}", c2.get()); // Prints 2
```

<sup>3</sup>`usize` is an unsigned integer type of platform-dependent size large enough to cover the address space.

The `Cell<i32>` type provides operations for storing and obtaining its content: `set` has type `fn(&Cell<i32>, i32)`, and `get` has type `fn(&Cell<i32>) -> i32`. Both of these only take a shared reference, so they can be called in the presence of arbitrary aliasing. So after we just spent several pages explaining that safety in Rust arises from ruling out aliasing and mutation, now we have `set`, which seems to completely violate this principle. How can this be safe?

The answer to this question has two parts. First of all, `Cell` only allows getting a *copy* of the content via `get`; it is not possible to obtain a pointer into the content. This rules out deep pointers into the `Cell`, making mutation safe. Unsurprisingly, `get` requires the content of the `Cell` to be `Copy`. In particular, `get` cannot be used with cells that contain non-`Copy` types like `Vec<i32>`.

However, there is still a potential source of problems, which arises from Rust’s support for multithreading. In particular, the following program must *not* be accepted:

```
1 let c = &Cell::new(0);
2 join(|| c.set(1), || println!("{:?}", c.get()));
```

The threads perform conflicting unsynchronized accesses to `c`, *i.e.*, this program has a data race.

To rule out programs like the one above, Rust has a notion of types being “sendable to another thread”. Such types satisfy the `Send` bound. The type of `join` demands that the environment captured by the closure satisfies `Send`. For example, `Vec<i32>` is `Send` because when the vector is moved to another thread, the previous owner is no longer allowed to access the vector—so it is fine for the new owner, in a different thread, to perform any operation whatsoever on the vector.

In the case above, the closure captures a shared reference to `c` of type `&Cell<i32>`. To check whether shared references are `Send`, there is another bound called `Sync`, with the property that type `&T` is `Send` if and only if `T` is `Sync`. Intuitively, a type is `Sync` if it is safe to have shared references to the same instance of the type in different threads. In other words, all the operations available on `&T` have to be thread-safe. For example, `Vec<i32>` is `Sync` because shared references only permit reading the vector, and it is fine if multiple threads do that at the same time. However, `Cell<i32>` is *not* `Sync` because `set` is not thread-safe. As a consequence, `&Cell<i32>` is not `Send`, which leads to the program above being rejected.

**2.5.2 Mutex.** The `Cell` type is a great example of interior mutability and a zero-cost abstraction as it comes with no overhead: `get` and `set` compile to plain unsynchronized accesses, so the compiled program is just as efficient as a C program using shared mutable state. However, as we have seen, `Cell` pays for this advantage by not being thread-safe. The Rust standard library also provides primitives for thread-safe shared mutable state, one being `Mutex`, which implements mutual exclusion (via a standard lock) for protecting access to one shared memory location. Consider the following example:

```
1 let mutex = Mutex::new(Vec::new());
2 join( || { let mut guard = mutex.lock().unwrap();
3           guard.deref_mut().push(0) },
4       || { let mut guard = mutex.lock().unwrap();
5           println!("{:?}", guard.deref_mut()) } );
```

This program starts by creating a mutex of type `Mutex<Vec<i32>>` initialized with an empty vector. The mutex is then shared between two threads (implicitly relying on `Mutex<Vec<i32>>` being `Sync`). The first thread acquires the lock, and pushes an element to the vector. The second thread acquires the lock just to print the contents of the vector.

The `guard` variables are of type `MutexGuard<'a, Vec<i32>>` where `'a` is the lifetime of the shared mutex reference passed to `lock` (this ensures that the mutex itself will stay around for at least as long as the guard). Mutex guards serve two purposes. Most importantly, if a thread owns a

guard, that means it holds the lock. To this end, guards provide a method `deref_mut` which turns a mutable reference of `MutexGuard` into a mutable reference of `Vec<i32>` with the same lifetime. Very much unlike `Cell`, the `Mutex` type permits obtaining deep pointers into the data guarded by the lock. In fact, the compiler will insert calls to `deref_mut` automatically where appropriate, making `MutexGuard<'a, Vec<i32>>` behave essentially like `'a mut Vec<i32>`.

Moreover, the guards are set up to release the lock when their destructors are called, which will happen automatically when the guards go out of scope. This is safe because, just like with `index_mut` (§2.4), the compiler ensures that deep pointers obtained through `deref_mut` have all expired by the time the guard is dropped.

### 3 THE $\lambda_{\text{Rust}}$ LANGUAGE AND TYPE SYSTEM

In this section, we introduce  $\lambda_{\text{Rust}}$ : our formal version of Rust. The Rust surface language comes with significant syntactic sugar (some of which we have already seen). To simplify the formalization,  $\lambda_{\text{Rust}}$  features only a small set of primitive constructs, and requires the advanced sugar of Rust’s surface language to be desugared into primitive constructs. Indeed, something very similar happens in the compiler itself, where surface Rust is lowered into the *Mid-level Intermediate Representation (MIR)* [Matsakis 2016a].  $\lambda_{\text{Rust}}$  is much closer to MIR than to surface Rust.

Before we present the syntax (§3.1), operational semantics (§3.2) and type system (§3.3) of  $\lambda_{\text{Rust}}$ , we highlight some of its key features:

- Programs are represented in continuation-passing style. This choice enables us to represent complex control-flow constructs, like labeled `break` and early `return`, as present in the Rust surface language. Furthermore, following the correspondence of CPS and control-flow graphs [Appel 2007], this makes  $\lambda_{\text{Rust}}$  easier to relate to MIR.
- The individual instructions of our language perform a single operation. By keeping the individual instructions simple and avoiding large composed expressions, it becomes possible to describe the type system in a concise way.
- The memory model of  $\lambda_{\text{Rust}}$  supports pointer arithmetic and ensures that programs with data races or illegal memory accesses can reach a stuck state in the operational semantics. In particular, programs that cannot get stuck in any execution—a guarantee established by the adequacy theorem of our type system (Theorem 7.2)—are data-race free.

#### 3.1 The Syntax

The syntax of  $\lambda_{\text{Rust}}$  is as follows:

$$\begin{aligned}
 \text{Path } \ni p &::= x \mid p.n \\
 \text{Val } \ni v &::= \mathbf{false} \mid \mathbf{true} \mid z \mid \ell \mid \mathbf{funrec} f(\bar{x}) \mathbf{ret} k := F \\
 \text{Instr } \ni I &::= v \mid p \mid p_1 + p_2 \mid p_1 - p_2 \mid p_1 \leq p_2 \mid p_1 = p_2 \mid \mathbf{new}(n) \mid \mathbf{delete}(n, p) \\
 &\quad \mid *p \mid p_1 := p_2 \mid p_1 :=_n *p_2 \mid p :=_{\text{inj } i} () \mid p_1 :=_{\text{inj } i} p_2 \mid p_1 :=_n^{\text{inj } i} *p_2 \mid \dots \\
 \text{FuncBody } \ni F &::= \mathbf{let} x = I \mathbf{in} F \mid \mathbf{letcont} k(\bar{x}) := F_1 \mathbf{in} F_2 \mid \mathbf{newlft}; F \mid \mathbf{endlft}; F \\
 &\quad \mid \mathbf{if} p \mathbf{then} F_1 \mathbf{else} F_2 \mid \mathbf{case} *p \mathbf{of} \bar{F} \mid \mathbf{jump} k(\bar{x}) \mid \mathbf{call} f(\bar{x}) \mathbf{ret} k
 \end{aligned}$$

We let path offsets  $n$  and integer literals  $z$  range over the integers, and sum indices  $i$  range over the natural numbers. The language has two kinds of variables: program variables, which are written as  $x$  or  $f$ , and continuation variables, which are written as  $k$ .

We distinguish four classes of expressions: *function bodies*  $F$  consist of *instructions*  $I$  that operate on *paths*  $p$  and *values*  $v$ . Only the most basic values can be written as literals: the Booleans `false` and `true`, integers  $z$ , locations  $\ell$  (see §3.2 for further details), and functions `funrec  $f(\bar{x})$  ret  $k := F$` .

There are no literals for products or sums, as these only exist in memory, represented by sequences of values and tagged unions, respectively. Paths are used to express the values that instructions operate on. The common case is to directly refer to a local variable  $x$ . Beyond this, paths can refer to parts of a compound data structure laid out in memory: Offsets  $p.n$  perform pointer arithmetic, incrementing the pointer expressed by  $p$  by  $n$  memory cells.

Function bodies mostly serve to chain instructions together and manage control flow, which is handled through continuations. Continuations are declared using **letcont**  $k(\bar{x}) := F_1$  **in**  $F_2$ , and called using **jump**  $k(\bar{x})$ . The parameters  $\bar{x}$  are instantiated when calling the continuation. We allow continuations to be recursive, in order to model looping constructs like **while** and **for**.

The “ghost instructions” **newlft** and **endlft** start and end lifetimes. These instructions have interesting typing rules, but do not do anything operationally.

Functions can be declared using **funrec**  $f(\bar{x})$  **ret**  $k := F$ , where  $f$  is a binder for the recursive call,  $\bar{x}$  is a list of binders for the arguments, and  $k$  is a binder for the return continuation. The return continuation takes one argument for the return value. Functions can be called using **call**  $f(\bar{x})$  **ret**  $k$ , where  $\bar{x}$  is the list of parameters and  $k$  is the continuation that should be called when the function returns.

Local variables of  $\lambda_{\text{Rust}}$ —as represented by **let** bindings—are pure values. This is different from local variables in Rust (and MIR), which are mutable and addressable. Hence, to correctly model Rust’s local variables, we allocate them on the heap. Similar to prior work on low-level languages [Krebbers 2015; Leroy et al. 2012], we do not make a distinction between the stack and the heap. In practice, this looks as follows:

```

fn option_as_mut<'a>
  (x: &'a mut Option<i32>) ->
  Option<&'a mut i32> {
  match *x {
    None => None,
    Some(ref mut t) => Some(t)
  }
}

funrec option_as_mut(x) ret ret :=
  let r = new(2) in
  letcont k() := delete(1, x); jump ret(r) in
  let y = *x in case *y of
  - r :=inj0 (); jump k()
  - r :=inj1 y.1; jump k()

```

We see that the function argument  $x$  is a pointer, which is dereferenced when used and deallocated before the function returns. In this case, since the Rust program takes a pointer,  $x$  actually is a pointer to a pointer. Similarly, a pointer  $r$  is allocated for the return value.

The  $\lambda_{\text{Rust}}$  language has instructions for the usual arithmetic operations, memory allocation, and deallocation, as well as loading from memory ( $*p$ ) and storing a value into memory ( $p_1 := p_2$ ). The memcopy-like instruction  $p_1 :=_n *p_2$  copies the contents of  $n$  memory locations from  $p_2$  to  $p_1$ . All of these accesses are *non-atomic*, i.e., they are not thread-safe. We will come back to this point in §3.2.

The example above also demonstrates the handling of sums. Values of the `Option<i32>` type are represented by a sequence of two base values: an integer value that represents the tag (0 for `None` and 1 for `Some`) and, if the tag is 1, a value of type `i32` for the argument  $t$  of `Some(t)`. If the tag is 0, the second value can be anything. The instructions  $p_1 :=^{\text{inj } i} p_2$  and  $p_1 :=^{\text{inj } i}_n *p_2$  can be used to assign to a pointer  $p_1$  of sum type, setting both the tag  $i$  and the value associated with this variant of the union, while  $p_1 :=^{\text{inj } i} ()$  is used for variants that have no data associated with them (like `None`). The **case** command is used to perform case-distinction on the tag, jumping to the  $n$ -th branch for tag  $n$ .

There are more instructions available in the underlying core language, e.g., instructions to spawn threads or perform atomic accesses, including CAS (compare-and-swap). However, the type system does not provide any typing rules for these instructions, so they can only be used by `unsafe` code.

### 3.2 The Operational Semantics

The operational semantics of  $\lambda_{\text{Rust}}$  is given by translation into a core language. The core language is a lambda calculus equipped with primitive values, pointer arithmetic, and concurrency. We define the semantics this way for three reasons. First of all, we can model some of the  $\lambda_{\text{Rust}}$  constructs (e.g.,  $p_1 :=_n * p_2$ ) as sequences of simpler instructions in the core language. Secondly, we can reduce both continuations and functions to plain lambda terms. Finally, the core language supports a substitution-based semantics, which makes reasoning more convenient, whereas the CPS grammar given above is not actually closed under substitution. The details of the core language are fully spelled out in our technical appendix [Jung et al. 2017a].

The memory model is inspired by CompCert [Leroy et al. 2012] in order to properly support pointer arithmetic. On top of this, we want the memory model to detect and rule out data races. Following C++11 [ISO Working Group 21 2011], we provide both *non-atomic* memory accesses, on which races are considered undefined behavior, and *atomic* accesses, which may be racy. However, for simplicity, we only provide sequentially consistent (SC) atomic operations, avoiding consideration of C++11's relaxed atomics in this paper. Notice that, like in C++, atomicity is a property of the individual memory *access*, not of the memory location. The same location can be subject to both atomic and non-atomic accesses. We consider a program to have a data race if there are ever two concurrent accesses to the same location, at least one of which is a write, and at least one of which is non-atomic. To detect such data races, every location is equipped with some additional state (resembling a reader-writer lock), which is checked dynamically to see if a particular memory access is permitted. We have shown in Coq that if a program has a data race, then it has an execution where these checks fail. As a consequence, if we prove that a program cannot get stuck (which implies that the checks always succeed, in all executions), then the program is data-race free.

In our handling of uninitialized memory, we follow Lee et al. [2017]. Upon allocation, memory holds a *poison* value  $\text{\textcircled{X}}$  that will cause the program to get stuck if it is ever used for a computation or a conditional branch. The only safe operations on  $\text{\textcircled{X}}$  are loading from and storing to memory.

### 3.3 The Type System

The types and contexts of  $\lambda_{\text{Rust}}$  are as follows:

$$\begin{array}{lll}
 \text{Lft} \ni \kappa ::= \alpha \mid \mathbf{static} & \mathbf{E} ::= \emptyset \mid \mathbf{E}, \kappa \sqsubseteq_e \kappa' & \mathbf{T} ::= \emptyset \mid \mathbf{T}, p \triangleleft \tau \mid \mathbf{T}, p \triangleleft^{\dagger \kappa} \tau \\
 \text{Mod} \ni \mu ::= \mathbf{mut} \mid \mathbf{shr} & \mathbf{L} ::= \emptyset \mid \mathbf{L}, \kappa \sqsubseteq_l \bar{\kappa} & \mathbf{K} ::= \emptyset \mid \mathbf{K}, k \triangleleft \mathbf{cont}(\mathbf{L}; \bar{x}. \mathbf{T}) \\
 \text{Type} \ni \tau ::= T \mid \mathbf{bool} \mid \mathbf{int} \mid \mathbf{own}_n \tau \mid \&_{\mu}^{\kappa} \tau \mid \downarrow_n \mid \Pi \bar{\tau} \mid \Sigma \bar{\tau} \mid \forall \bar{\alpha}. \mathbf{fn}(\bar{c} : \mathbf{E}; \bar{\tau}) \rightarrow \tau \mid \mu T. \tau
 \end{array}$$

Selected typing rules are shown in Figure 1 and Figure 2. We first discuss the types provided by the system, before looking at some examples.

There are two kinds of pointer types: *owned pointers*  $\mathbf{own}_n \tau$  and (*borrowed*) *references*  $\&_{\mu}^{\kappa} \tau$ .

Owned pointers  $\mathbf{own}_n \tau$  are used to represent full ownership of (some part of) a heap allocation. Because we model the stack using heap allocations, owned pointers also represent Rust's local, stack-allocated variables. As usual,  $\tau$  is the type of the pointee. Furthermore,  $n$  tracks the size of the entire allocation. This can be different from the size of  $\tau$  for *inner pointers* that point into a larger data structure.<sup>4</sup> Still, most of the time,  $n$  is the size of  $\tau$ , in which case we omit the subscript.

References  $\&_{\mu}^{\kappa} \tau$  are qualified by a *modifier*  $\mu$ , which is either **mut** (for mutable references, which are unique) or **shr** (for shared references), and a *lifetime*  $\kappa$ . References  $\&_{\mu}^{\kappa} \tau$  are borrowed for lifetime  $\kappa$  and, as such, can only be used as long as the lifetime  $\kappa$  is alive, *i.e.*, still ongoing. Lifetimes begin and end at the **newlft** and **endlft** ghost instructions, following **F-NEWLFT** and **F-ENDLFT**.

<sup>4</sup>Such pointers can be obtained using **C-SPLIT-OWN**.

Furthermore, the special lifetime **static** lasts for the execution of the entire program (corresponding to **'static** in Rust, which plays the same role). The type system is able to abstract over lifetimes, so most of the time, we will work with lifetime variables  $\alpha$ .

The type  $\downarrow_n$  describes arbitrary sequences of  $n$  base values. This type represents uninitialized memory. For example, when allocating an owned pointer (rule **S-NEW**), its type is **own**  $\downarrow_n$ . Owned pointers permit strong updates, which means their type  $\tau$  can change when the memory gets (re-)initialized. Similarly, the type changes back to **own**  $\downarrow_n$  when data is moved out of the owned pointer (rule **TREAD-OWN-MOVE**). Note that this is sound because ownership of owned pointers is unique.

The types  $\Pi\bar{\tau}$  and  $\Sigma\bar{\tau}$  represent  $n$ -ary products and sums, respectively. In particular, this gives rise to a unit type **()** (the empty product  $\Pi[]$ ) and the empty type **!** (the empty sum  $\Sigma[]$ ). We use  $\tau_1 \times \tau_2$  and  $\tau_1 + \tau_2$  as notation for binary products ( $\Pi[\tau_1, \tau_2]$ ) and sums ( $\Sigma[\tau_1, \tau_2]$ ), respectively.

Function types  $\forall\bar{\alpha}. \mathbf{fn}(\bar{f} : \mathbf{E}; \bar{\tau}) \rightarrow \tau$  can be polymorphic over lifetimes  $\bar{\alpha}$ . The external lifetime context  $\mathbf{E}$  can be used to demand that one lifetime parameter be included in another one. The lifetime  $\bar{f}$  here is a binder that can be used in  $\mathbf{E}$  to refer to the lifetime of this function. For example,  $\forall\alpha. \mathbf{fn}(\bar{f} : \bar{f} \sqsubseteq \alpha; \&_{\mathbf{mut}}^\alpha \mathbf{int}) \rightarrow \mathbf{()}$  is the type of a function that takes a mutable reference to an integer with any lifetime that covers this function call (matching the implicit assumption Rust makes), and returns unit. Note that, to allow passing and returning objects of arbitrary size, both the parameters and the return value are transmitted via owned pointers; this calling convention is universally applied and hence does not show up in the function type.

Finally,  $\lambda_{\text{Rust}}$  supports *recursive types*  $\mu T. \tau$ , with the restriction (enforced by the well-formedness judgment shown in the appendix [Jung et al. 2017a]) that  $T$  only appears in  $\tau$  below a pointer type or within a function type.

To keep the type system of  $\lambda_{\text{Rust}}$  focused on our core objective (modeling borrowing and lifetimes), there is no support for type-polymorphic functions. Instead, we handle polymorphism on the meta-level: In our shallow embedding of the type system in Coq, we can quantify any definition and theorem over arbitrary semantic types (§4). We exploit this flexibility when verifying the safety of Rust libraries that use **unsafe** features (§6). These libraries are typically polymorphic, and by keeping the verification similarly polymorphic, we can prove that functions and libraries are safe to use at any instantiation of their type parameters.

*Type-checking the example.* The typing judgments for function bodies  $F$  and instructions  $I$  have the shape  $\Gamma \mid \mathbf{E}; \mathbf{L} \mid \mathbf{K}; \mathbf{T} \vdash F$  and  $\Gamma \mid \mathbf{E}; \mathbf{L} \mid \mathbf{T}_1 \vdash I \dashv x. \mathbf{T}_2$ . To see these judgments in action, we will go through part of the typing derivation of the example from §3.1. The code, together with some annotated type and continuation contexts, is repeated in Figure 3. Overall, we will want to derive a judgment for the body of `option_as_mut`, in the following initial contexts:

$$\begin{aligned} \Gamma_1 &:= x : \mathbf{val}, \mathbf{ret} : \mathbf{val}, \alpha : \mathbf{lft}, \bar{f} : \mathbf{lft} \\ \mathbf{E}_1 &:= \bar{f} \sqsubseteq_e \alpha \\ \mathbf{L}_1 &:= \bar{f} \sqsubseteq_l [] \\ \mathbf{K}_1 &:= \mathbf{ret} \triangleleft \mathbf{cont}(\bar{f} \sqsubseteq_l []); r. r \triangleleft \mathbf{own} \mathbf{() + \&_{mut}^\alpha \mathbf{int}} \\ \mathbf{T}_1 &:= x \triangleleft \mathbf{own} \&_{\mathbf{mut}}^\alpha \mathbf{() + \mathbf{int}} \end{aligned}$$

The first context, the *variable context*  $\Gamma$ , is the only binding context. It introduces all variables that are free in the judgment and keeps track of whether they are program variables ( $x : \mathbf{val}$ ; this also covers continuations), lifetime variables ( $\alpha : \mathbf{lft}$ ), or type variables<sup>5</sup> ( $T : \mathbf{type}$ ). All the remaining contexts state facts and assert ownership related to variables introduced here, but they do not introduce additional binders.

<sup>5</sup>Type variables can only occur in the definition of recursive types.

**Rules for lifetimes:**

$$\begin{array}{c}
\text{LINCL-STATIC} \\
\frac{}{\Gamma \mid \mathbf{E}; \mathbf{L} \vdash \kappa \sqsubseteq \mathbf{static}}
\end{array}
\quad
\begin{array}{c}
\text{LINCL-LOCAL} \\
\frac{\kappa \sqsubseteq_1 \bar{\kappa} \in \mathbf{L} \quad \kappa' \in \bar{\kappa}}{\Gamma \mid \mathbf{E}; \mathbf{L} \vdash \kappa \sqsubseteq \kappa'}
\end{array}
\quad
\begin{array}{c}
\text{LINCL-EXTERN} \\
\frac{\kappa \sqsubseteq_e \kappa' \in \mathbf{E}}{\Gamma \mid \mathbf{E}; \mathbf{L} \vdash \kappa \sqsubseteq \kappa'}
\end{array}
\quad
\begin{array}{c}
\text{LINCL-REFL} \\
\frac{(\Gamma \mid \mathbf{E}; \mathbf{L} \vdash \kappa_1 \sqsubseteq \kappa_2 \text{ and } \Gamma \mid \mathbf{E}; \mathbf{L} \vdash \kappa \text{ alive})}{\Gamma \mid \mathbf{E}; \mathbf{L} \vdash \kappa \sqsubseteq \kappa}
\end{array}$$

$$\begin{array}{c}
\text{LALIVE-LOCAL} \\
\frac{\kappa \sqsubseteq_1 \bar{\kappa} \in \mathbf{L} \quad \forall i. \mathbf{E}; \mathbf{L} \vdash \bar{\kappa}_i \text{ alive}}{\Gamma \mid \mathbf{E}; \mathbf{L} \vdash \kappa \text{ alive}}
\end{array}$$

$$\begin{array}{c}
\text{LALIVE-INCL} \\
\frac{\Gamma \mid \mathbf{E}; \mathbf{L} \vdash \kappa \text{ alive} \quad \Gamma \mid \mathbf{E}; \mathbf{L} \vdash \kappa' \sqsubseteq \kappa'}{\mathbf{E}; \mathbf{L} \vdash \kappa' \text{ alive}}
\end{array}$$

$$\begin{array}{c}
\text{LINCL-TRANS} \\
\frac{\Gamma \mid \mathbf{E}; \mathbf{L} \vdash \kappa \sqsubseteq \kappa' \quad \Gamma \mid \mathbf{E}; \mathbf{L} \vdash \kappa' \sqsubseteq \kappa''}{\Gamma \mid \mathbf{E}; \mathbf{L} \vdash \kappa \sqsubseteq \kappa''}
\end{array}$$

**Rules for subtyping and type coercions:**

$$\begin{array}{c}
\text{T-BOR-LFT} \\
\frac{\Gamma \mid \mathbf{E}; \mathbf{L} \vdash \kappa \sqsubseteq \kappa'}{\Gamma \mid \mathbf{E}; \mathbf{L} \vdash \&_{\mu}^{\kappa'} \tau \Rightarrow \&_{\mu}^{\kappa} \tau}
\end{array}
\quad
\begin{array}{c}
\text{C-SUBTYPE} \\
\frac{\Gamma \mid \mathbf{E}; \mathbf{L} \vdash \tau \Rightarrow \tau'}{\Gamma \mid \mathbf{E}; \mathbf{L} \vdash p \triangleleft \tau \stackrel{\text{ctx}}{\Rightarrow} p \triangleleft \tau'}
\end{array}
\quad
\begin{array}{c}
\text{C-COPY} \\
\frac{\tau \text{ copy}}{\Gamma \mid \mathbf{E}; \mathbf{L} \vdash p \triangleleft \tau \stackrel{\text{ctx}}{\Rightarrow} p \triangleleft \tau, p \triangleleft \tau}
\end{array}$$

$$\begin{array}{c}
\text{C-SPLIT-OWN} \\
\mathbf{E}; \mathbf{L} \vdash p \triangleleft \mathbf{own}_n \tau_1 \times \tau_2 \stackrel{\text{ctx}}{\Leftrightarrow} p.0 \triangleleft \mathbf{own}_n \tau_1, \triangleleft \mathbf{own}_n \tau_2
\end{array}
\quad
\begin{array}{c}
\text{C-SHARE} \\
\frac{\Gamma \mid \mathbf{E}; \mathbf{L} \vdash \kappa \text{ alive}}{\Gamma \mid \mathbf{E}; \mathbf{L} \vdash p \triangleleft \&_{\mathbf{mut}}^{\kappa} \tau \stackrel{\text{ctx}}{\Rightarrow} p \triangleleft \&_{\mathbf{shr}}^{\kappa} \tau}
\end{array}$$

$$\begin{array}{c}
\text{C-BORROW} \\
\Gamma \mid \mathbf{E}; \mathbf{L} \vdash p \triangleleft \mathbf{own}_n \tau \stackrel{\text{ctx}}{\Rightarrow} p \triangleleft \&_{\mathbf{mut}}^{\kappa} \tau, p \triangleleft \dagger^{\kappa} \mathbf{own}_n \tau
\end{array}
\quad
\begin{array}{c}
\text{C-REBORROW} \\
\frac{\Gamma \mid \mathbf{E}; \mathbf{L} \vdash \kappa' \sqsubseteq \kappa}{\Gamma \mid \mathbf{E}; \mathbf{L} \vdash p \triangleleft \&_{\mathbf{mut}}^{\kappa} \tau \stackrel{\text{ctx}}{\Rightarrow} p \triangleleft \&_{\mathbf{mut}}^{\kappa'} \tau, p \triangleleft \dagger^{\kappa'} \&_{\mathbf{mut}}^{\kappa} \tau}
\end{array}$$

**Rules for reading and writing:**

$$\begin{array}{c}
\text{TREAD-OWN-COPY} \\
\frac{\tau \text{ copy}}{\Gamma \mid \mathbf{E}; \mathbf{L} \vdash \mathbf{own}_n \tau \circ^{-\tau} \mathbf{own}_n \tau}
\end{array}
\quad
\begin{array}{c}
\text{TREAD-OWN-MOVE} \\
\frac{n = \text{size}(\tau)}{\Gamma \mid \mathbf{E}; \mathbf{L} \vdash \mathbf{own}_m \tau \circ^{-\tau} \mathbf{own}_m \not\downarrow n}
\end{array}
\quad
\begin{array}{c}
\text{TREAD-BOR} \\
\frac{\tau \text{ copy} \quad \Gamma \mid \mathbf{E}; \mathbf{L} \vdash \kappa \text{ alive}}{\Gamma \mid \mathbf{E}; \mathbf{L} \vdash \&_{\mu}^{\kappa} \tau \circ^{-\tau} \&_{\mu}^{\kappa} \tau}
\end{array}$$

$$\begin{array}{c}
\text{TWRITE-OWN} \\
\frac{\text{size}(\tau) = \text{size}(\tau')}{\Gamma \mid \mathbf{E}; \mathbf{L} \vdash \mathbf{own}_n \tau' \circ^{-\tau} \mathbf{own}_n \tau}
\end{array}
\quad
\begin{array}{c}
\text{TWRITE-BOR} \\
\frac{\Gamma \mid \mathbf{E}; \mathbf{L} \vdash \kappa \text{ alive}}{\Gamma \mid \mathbf{E}; \mathbf{L} \vdash \&_{\mathbf{mut}}^{\kappa} \tau \circ^{-\tau} \&_{\mathbf{mut}}^{\kappa} \tau}
\end{array}$$

**Rules for typing of instructions:**

$$\begin{array}{c}
\text{S-NUM} \\
\Gamma \mid \mathbf{E}; \mathbf{L} \mid \emptyset \vdash z \div x. x \triangleleft \mathbf{int}
\end{array}
\quad
\begin{array}{c}
\text{S-NAT-LEQ} \\
\Gamma \mid \mathbf{E}; \mathbf{L} \mid p_1 \triangleleft \mathbf{int}, p_2 \triangleleft \mathbf{int} \vdash p_1 \leq p_2 \div x. x \triangleleft \mathbf{bool}
\end{array}$$

$$\begin{array}{c}
\text{S-NEW} \\
\Gamma \mid \mathbf{E}; \mathbf{L} \mid \emptyset \vdash \mathbf{new}(n) \div x. x \triangleleft \mathbf{own}_n \not\downarrow n
\end{array}
\quad
\begin{array}{c}
\text{S-DELETE} \\
\frac{n = \text{size}(\tau)}{\Gamma \mid \mathbf{E}; \mathbf{L} \mid p \triangleleft \mathbf{own}_n \tau \vdash \mathbf{delete}(n, p) \div \emptyset}
\end{array}$$

$$\begin{array}{c}
\text{S-DEREF} \\
\frac{\Gamma \mid \mathbf{E}; \mathbf{L} \vdash \tau_1 \circ^{-\tau} \tau'_1 \quad \text{size}(\tau) = 1}{\Gamma \mid \mathbf{E}; \mathbf{L} \mid p \triangleleft \tau_1 \vdash *p \div x. p \triangleleft \tau'_1, x \triangleleft \tau}
\end{array}
\quad
\begin{array}{c}
\text{S-SUM-ASSGN} \\
\frac{\bar{\tau}_i = \tau \quad \tau_1 \circ^{-\sum \bar{\tau}} \tau'_1}{\mathbf{E}; \mathbf{L} \mid p_1 \triangleleft \tau_1, p_2 \triangleleft \tau \vdash p_1 := \text{inj } i \quad p_2 \div p_1 \triangleleft \tau'_1}
\end{array}$$

Fig. 1. A selection of the typing rules of  $\lambda_{\text{Rust}}$  (helper judgments and instructions).



## Rules for typing of function bodies:

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

$$\frac{\text{F-CONSEQUENCE} \quad \Gamma \mid \mathbf{E}; \mathbf{L} \vdash \mathbf{T} \stackrel{\text{ctx}}{\Rightarrow} \mathbf{T}' \quad \mathbf{K}' \subseteq \mathbf{K} \quad \Gamma \mid \mathbf{E}; \mathbf{L}' \mid \mathbf{K}'; \mathbf{T}' \vdash F}{\Gamma \mid \mathbf{E}; \mathbf{L} \mid \mathbf{K}; \mathbf{T} \vdash F}$$

$$\frac{\text{F-LET} \quad \Gamma \mid \mathbf{E}; \mathbf{L} \mid \mathbf{T}_1 \vdash I \rightarrow x. \mathbf{T}_2 \quad \Gamma, x : \mathbf{val} \mid \mathbf{E}; \mathbf{L} \mid \mathbf{K}; \mathbf{T}_2, \mathbf{T} \vdash F}{\Gamma \mid \mathbf{E}; \mathbf{L} \mid \mathbf{K}; \mathbf{T}_1, \mathbf{T} \vdash \mathbf{let} \ x = I \ \mathbf{in} \ F}$$

$$\frac{\text{F-LETCONT} \quad \Gamma, k, \bar{x} : \mathbf{val} \mid \mathbf{E}; \mathbf{L}_1 \mid \mathbf{K}, k \triangleleft \mathbf{cont}(\mathbf{L}_1; \bar{x}. \mathbf{T}'); \mathbf{T}' \vdash F_1 \quad \Gamma, k : \mathbf{val} \mid \mathbf{E}; \mathbf{L}_2 \mid \mathbf{K}, k \triangleleft \mathbf{cont}(\mathbf{L}_1; \bar{x}. \mathbf{T}'); \mathbf{T} \vdash F_2}{\Gamma \mid \mathbf{E}; \mathbf{L}_2 \mid \mathbf{K}; \mathbf{T} \vdash \mathbf{letcont} \ k(\bar{x}) := F_1 \ \mathbf{in} \ F_2}$$

$$\frac{\text{F-JUMP} \quad \Gamma \mid \mathbf{E}; \mathbf{L} \vdash \mathbf{T} \stackrel{\text{ctx}}{\Rightarrow} \mathbf{T}'[\bar{y}/\bar{x}]}{\Gamma \mid \mathbf{E}; \mathbf{L} \mid k \triangleleft \mathbf{cont}(\mathbf{L}; \bar{x}. \mathbf{T}'); \mathbf{T} \vdash \mathbf{jump} \ k(\bar{y})}$$

$$\frac{\text{F-NEWLFT} \quad \Gamma, \alpha : \mathbf{lft} \mid \mathbf{E}; \mathbf{L}, \alpha \sqsubseteq_1 \bar{\kappa} \mid \mathbf{K}; \mathbf{T} \vdash F}{\Gamma \mid \mathbf{E}; \mathbf{L} \mid \mathbf{K}; \mathbf{T} \vdash \mathbf{newlft}; F}$$

$$\frac{\text{F-ENDLFT} \quad \Gamma \mid \mathbf{E}; \mathbf{L} \mid \mathbf{K}; \mathbf{T}' \vdash F \quad \mathbf{T} \Rightarrow^{\dagger \kappa} \mathbf{T}'}{\Gamma \mid \mathbf{E}; \mathbf{L}, \kappa \sqsubseteq_1 \bar{\kappa} \mid \mathbf{K}; \mathbf{T} \vdash \mathbf{endlft}; F}$$

$$\frac{\text{F-CASE-BOR} \quad \Gamma \mid \mathbf{E}; \mathbf{L} \vdash \kappa \ \mathbf{alive} \quad \forall i. (\Gamma \mid \mathbf{E}; \mathbf{L} \mid \mathbf{K}; \mathbf{T}, p.1 \triangleleft \&_{\mu}^{\kappa} \tau_i \vdash F_i)}{\Gamma \mid \mathbf{E}; \mathbf{L} \mid \mathbf{K}; \mathbf{T}, p \triangleleft \&_{\mu}^{\kappa} \Sigma \bar{\tau} \vdash \mathbf{case}^* p \ \mathbf{of} \ \bar{F}}$$

$$\frac{\text{F-CALL} \quad \Gamma \mid \mathbf{E}; \mathbf{L} \vdash \mathbf{T} \stackrel{\text{ctx}}{\Rightarrow} \bar{x} \triangleleft \mathbf{own} \ \bar{\tau}, \mathbf{T}' \quad \Gamma \mid \mathbf{E}; \mathbf{L} \vdash \bar{\kappa} \ \mathbf{alive} \quad \Gamma, \varphi : \mathbf{lft} \mid \mathbf{E}, \varphi \sqsubseteq_e \bar{\kappa}; \mathbf{L} \vdash \mathbf{E}'}{\Gamma \mid \mathbf{E}; \mathbf{L} \mid k \triangleleft \mathbf{cont}(\mathbf{L}; y. y \triangleleft \mathbf{own} \ \tau, \mathbf{T}'); \mathbf{T}, f \triangleleft \mathbf{fn}(\varphi : \mathbf{E}'; \bar{\tau}) \rightarrow \tau \vdash \mathbf{call} \ f(\bar{x}) \ \mathbf{ret} \ k}$$

Fig. 2. A selection of the typing rules of  $\lambda_{\text{Rust}}$  (function bodies).

Our initial variable context consists of the parameter  $x$ , our return continuation  $\text{ret}$ , the lifetime  $\alpha$  (corresponding to `'a`), and the lifetime  $\varphi$ , which (by convention) is always the name of the lifetime of the current function. This lifetime is used in the *external lifetime context*  $\mathbf{E}$  to state that  $\alpha$  outlives the current function call. Rust does not have a direct equivalent of  $\varphi$  in its surface syntax; instead it always implicitly assumes that lifetime parameters like `'a` outlive the current function.

The *typing context*  $\mathbf{T}$  is in charge of describing ownership of local variables. It mostly contains type assignments  $p \triangleleft \tau$ . It is important to stress that the typing context is substructural: Type assignments can only be duplicated if the type satisfies  $\tau \ \text{copy}$  (**C-COPY**), corresponding to Rust's **Copy** bound. In this case, we have a single variable  $x$  (our argument), which is an owned pointer to  $\&_{\text{mut}}^{\alpha} (\ () + \mathbf{int})$ , the  $\lambda_{\text{Rust}}$  equivalent of the Rust type `&'a mut Option<i32>`. As already mentioned, the additional owned pointer indirection here models the fact that  $x$  on the Rust side has an address in memory.

We mostly use **F-LET** to type-check the function one instruction at a time. The first instruction is **new**, so we use **S-NEW**. That extends our typing context with  $r$  being an uninitialized owned pointer:

$$x \triangleleft \mathbf{own} \ \&_{\text{mut}}^{\alpha} (\ () + \mathbf{int}), r \triangleleft \mathbf{own} \ \zeta_2$$

Next, we declare a continuation (**letcont**  $k() := \dots$ ). Continuations are tracked in the *continuation context*  $\mathbf{K}$ . Initially, we already have our return continuation  $\text{ret}$  of type  $\mathbf{cont}(\varphi \sqsubseteq_1 []; r. r \triangleleft \mathbf{own} (\ () + \&_{\text{mut}}^{\alpha} \mathbf{int}))$  in that context. This says that  $\text{ret}$  expects one argument  $r$  of our return type, `Option<'a mut i32>`.

The continuation also makes assumptions about the local lifetime context  $\mathbf{L}$  at the call site, which we will discuss soon. As usual with CPS, since the return type is given by the return continuation, the function judgment does not have a notion of a return type itself.

The function `option_as_mut` declares a continuation  $k$  to represent the merging control flow after the **case**. Following **F-LETCONT**, we have to pick  $\mathbf{T}'$ , the typing context at the call site of the

```

funrec option_as_mut(x) ret ret :=
  {K : ret ◁ cont( $\mathcal{F} \sqsubseteq_1 []$ ; r. r ◁ own () +  $\&_{\text{mut}}^\alpha \text{int}$ ); T : x ◁ own  $\&_{\text{mut}}^\alpha$  () + int}
  let r = new(2) in
  {T : x ◁ own  $\&_{\text{mut}}^\alpha$  () + int, r ◁ own  $\not\downarrow_2$ }
  letcont k() := delete(1, x); jump ret(r) in
  {K : ret ◁ ..., k ◁ cont( $\mathcal{F} \sqsubseteq_1 []$ ; r ◁ own () +  $\&_{\text{mut}}^\alpha \text{int}$ ), x ◁ own  $\not\downarrow_1$ }
  let y = *x in
  {T : x ◁ own  $\not\downarrow_1$ , r ◁ own  $\not\downarrow_2$ , y ◁  $\&_{\text{mut}}^\alpha$  () + int}
  case *y of
  - r :=inj0 (); jump k()
  - {T : x ◁ own  $\not\downarrow_1$ , r ◁ own  $\not\downarrow_2$ , y.1 ◁  $\&_{\text{mut}}^\alpha \text{int}$ }
    r :=inj1 y.1;
    {T : x ◁ own  $\not\downarrow_1$ , r ◁ own () +  $\&_{\text{mut}}^\alpha \text{int}$ }
    jump k()

```

Fig. 3. Example code with annotated type and continuation contexts.

continuation. It turns out that the right choice is  $r \triangleleft \mathbf{own} () + \&_{\text{mut}}^\alpha \text{int}$ ,  $x \triangleleft \mathbf{own} \not\downarrow_1$ . Let us omit checking that the continuation actually has this type, and continue on with the following new item in our continuation context:

$$k \triangleleft \mathbf{cont}(\mathcal{F} \sqsubseteq_1 []; r \triangleleft \mathbf{own} () + \&_{\text{mut}}^\alpha \text{int}), x \triangleleft \mathbf{own} \not\downarrow_1$$

Next, the code dereferences the argument (**let**  $y = *x$ ), which unwraps the additional owned pointer indirection that got inserted in the translation. Dereferencing is type-checked using **S-DEREF**. This rule uses a helper judgment:  $\Gamma \mid \mathbf{E}; \mathbf{L} \vdash \tau_1 \overset{r}{\dashv} \tau_2$  means that we can read a value of type  $\tau$  (the pointee) from a pointer of type  $\tau_1$ , and doing so will change the type of the pointer to  $\tau_2$ . In this case, we derive  $\mathbf{own} \&_{\text{mut}}^\alpha () + \mathbf{int} \overset{r}{\dashv} \mathbf{own} \not\downarrow_1$  from **TREAD-OWN-MOVE**. The type of the pointer changes because we *moved* the content out of the owned pointer. Effectively,  $x$  is now no longer initialized. After this instruction, our typing context becomes:

$$x \triangleleft \mathbf{own} \not\downarrow_1, r \triangleleft \mathbf{own} \not\downarrow_2, y \triangleleft \&_{\text{mut}}^\alpha () + \mathbf{int}$$

Next, we have to type-check the **case** using **F-CASE-BOR**, which involves loading the tag from  $y$ . Because we are dereferencing a reference (as opposed to an owned pointer) here, the type system requires us to show that the lifetime ( $\alpha$ ) is still alive. This is where the lifetime contexts **E** and **L** come in: We have to show  $\mathbf{E}; \mathbf{L} \vdash \alpha$  alive.

To this end, we first make use of the *external lifetime context* **E**, which tracks inclusions between lifetime parameters and the lifetime  $\mathcal{F}$  of the current function. Concretely, we make use of  $\mathcal{F} \sqsubseteq_e \alpha$  and apply **LALIVE-INCL**, which reduces the goal to  $\mathbf{E}; \mathbf{L} \vdash \mathcal{F}$  alive: Because  $\mathcal{F}$  is shorter than  $\alpha$ , it suffices to show that  $\mathcal{F}$  is still alive. In the second step, we employ our *local lifetime context* **L**, which tracks lifetimes that we control. Elements of this context are of the form  $\kappa \sqsubseteq_1 \bar{\kappa}$ , indicating that  $\kappa$  is a local lifetime with its superlifetimes listed in  $\bar{\kappa}$ . The rule **LALIVE-LOCAL** expresses that  $\kappa$  is alive as long as all its superlifetimes are alive. Because  $\mathcal{F}$  has no superlifetimes ( $\mathcal{F} \sqsubseteq_1 []$ ), this finishes the proof that  $\mathcal{F}$  is alive, and so is  $\alpha$ .

The local lifetime context also appears in the types of continuations: Both of our continuation expect the local lifetime context at their call site to be  $\varphi \sqsubseteq []$ . In other words,  $\varphi$  has to be still alive when the continuation is invoked. In particular, this means that `option_as_mut` cannot end  $\varphi$ , or else it would not be able to call its return continuation.

Having discharged the first premise of **F-CASE-BOR**, let us now come to the second premise: showing that all the branches of the case distinction are well-typed. The case distinction operates on a pointer to **() + int**, so in the branches, we can assume that  $y.1$  (the data stored in the sum) is a pointer to **()** or **int**, respectively. The second case is the more interesting one, where we go on with the following typing context:

$$x \triangleleft \mathbf{own} \downarrow_1, r \triangleleft \mathbf{own} \downarrow_2, y.1 \triangleleft \&_{\mathbf{mut}}^{\alpha} \mathbf{int}$$

The next instruction is  $r := \text{inj}_1 y.1$ , which is type-checked using **S-SUM-ASSGN**. Again the main work of adjusting the types is offloaded to a helper judgment:  $\Gamma \mid \mathbf{E}; \mathbf{L} \vdash \tau_1 \xrightarrow{\circ^r} \tau_2$  means that we can write a value of type  $\tau$  to a pointer of type  $\tau_1$ , changing the type of the pointer to  $\tau_2$ . In this case, we derive  $\Gamma \mid \mathbf{E}; \mathbf{L} \vdash \mathbf{own} \downarrow_2 \xrightarrow{0+\&_{\mathbf{mut}}^{\alpha} \mathbf{int}} \mathbf{own} \left( () + \&_{\mathbf{mut}}^{\alpha} \mathbf{int} \right)$  using **TWRITE-OWN**. This is a strong update, which changes the type of  $r$  from uninitialized to the return type of our example function. Our context thus becomes:

$$x \triangleleft \mathbf{own} \downarrow_1, r \triangleleft \mathbf{own} \left( () + \&_{\mathbf{mut}}^{\alpha} \mathbf{int} \right)$$

Notice that  $y.1$  disappeared from the context; it was used up when we moved it into  $r$ .

Finally, we jump to the continuation  $k$  that we declared earlier. This is type-checked using **F-JUMP**, which verifies that our current typing context  $\mathbf{T}$  and local lifetime context  $\mathbf{L}$  match what is expected by the continuation.

*Further noteworthy type system features.* Besides the type assignments we have already seen, the type context can also contain *lifetime-blocked* type assignments  $p \triangleleft^{\dagger\kappa} \tau$ . Such assignments are introduced when creating a reference (**C-BORROW**, **C-REBORROW**), which blocks the referent until the lifetime of the reference ends (**F-ENDLFT**), as expressed by the unblocking judgment  $\mathbf{T} \Rightarrow^{\dagger\kappa} \mathbf{T}'$ .

External lifetime context satisfaction  $\Gamma \mid \mathbf{E}; \mathbf{L} \vdash \mathbf{E}'$  is used on function calls to check the assumptions made by the callee (**F-CALL**). The  $\overline{\triangleleft}$  in **F-CALL** indicates that we are requiring a list of type assignments in the context, matching a list of variables ( $\overline{x}$ ) with an equal-length list of types ( $\mathbf{own} \overline{\tau}$ ).

Subtyping is described by  $\Gamma \mid \mathbf{E}; \mathbf{L} \vdash \tau_1 \Rightarrow \tau_2$ . The main forms of subtyping supported in Rust are lifetime inclusion (**T-BOR-LFT**) and (un)folding recursive types. Apart from that, there are the usual structural rules witnessing covariance and contravariance of type constructors. On the type context level,  $\Gamma \mid \mathbf{E}; \mathbf{L} \vdash \mathbf{T}_1 \xrightarrow{\text{ctx}} \mathbf{T}_2$  lifts subtyping (**C-SUBTYPE**) while also adding a few coercions that can only be applied at the top-level type. Most notably, a mutable reference can be coerced into a shared reference (**C-SHARE**), an owned pointer can be borrowed (**C-BORROW**) to create a mutable reference, and a mutable reference can be reborrowed (**C-REBORROW**).<sup>6</sup>

#### 4 RUSTBELT: A SEMANTIC MODEL OF $\lambda_{\text{Rust}}$ TYPES IN IRIS

Our proof of soundness of  $\lambda_{\text{Rust}}$  proceeds by defining a *logical relation*, which interprets the types and typing judgments of  $\lambda_{\text{Rust}}$  as logical predicates in an appropriate semantic domain. We focus here on the interpretation of types, leaving the interpretation of typing judgments and the statements of our main results to §7. First, in §4.1, we give a simplified version of the semantic domain of types. In §4.2, we give the semantic interpretation of some representative  $\lambda_{\text{Rust}}$  types. Finally, in §4.3, we focus on the interpretation of shared reference types. It will turn out that we have to generalize our semantic domain of types to account for them.

<sup>6</sup>There is no need to reborrow shared references because they are duplicable, and hence using subtyping to a shorter lifetime does not lose any information.

#### 4.1 A Simplified Semantic Domain of Types

The semantic domain of types answers the question “What is a type?”. Usually, the answer is that a *type denotes a set of values*—or, equivalently, a predicate over values. Fundamentally, this is also the case for  $\lambda_{\text{Rust}}$ , but the details get somewhat more complicated. First of all, our model of the type system of  $\lambda_{\text{Rust}}$  expresses types not as predicates in “plain mathematics” (e.g., the usual higher-order logic), but as predicates in *Iris*. As discussed in the introduction, *Iris* is a higher-order separation logic designed to prove correctness of complex concurrent programs. Using *Iris* to express types has the advantage that concepts like ownership are already built into the underlying framework, so the model itself does not have to take care of them.

Rather than try to explain all the features of *Iris* here, we will introduce them *en passant*, as needed. However, one that is worth mentioning up front is the ability to define predicates by *guarded recursion*. This means that a predicate can refer to itself recursively, but only below a  $\triangleright$  (“later”) modality [Appel et al. 2007] or some other appropriate “guard”. The use of a guard ensures that the circular definition can be solved—regardless of whether the recursive reference occurs positively, negatively, or both—using the technique of “step-indexing” [Appel and McAllester 2001]. For this reason,  $\triangleright$  appears in various places in our model; the placement of these  $\triangleright$ ’s is important for soundness, but is not otherwise relevant to our high-level exposition, so we will mostly ignore it in the rest of the paper.

Our interpretation of types associates to every type  $\tau$  an *Iris* predicate  $\llbracket \tau \rrbracket.\text{own} \in \text{Tid} \times \text{list}(\text{Val}) \rightarrow i\text{Prop}$ . This predicate takes two parameters and returns an *Iris* proposition (of type *iProp*). The second parameter is the *list* of values we are considering. It turns out that types in Rust do not just cover a single value: In general, data is laid out in memory and spans multiple locations. However, we have to impose some restrictions on the lists of values accepted by a type: we require that every type has a fixed *size*  $\llbracket \tau \rrbracket.\text{size}$ . This size is used to compute the layout of compound data structures, e.g., for product types. We require that a type only accepts lists whose length matches the size:

$$\llbracket \tau \rrbracket.\text{own}(t, \bar{v}) \Rightarrow |\bar{v}| = \llbracket \tau \rrbracket.\text{size} \quad (\text{TY-SIZE})$$

Furthermore, for *Copy* types we require that  $\llbracket \tau \rrbracket.\text{own}(t, \bar{v})$  be *persistent*. In *Iris*, a proposition is considered persistent if it does not describe ownership of any exclusive right or resource, and can therefore be freely copied and shared among several parties.

The first parameter of the predicate (of type *Tid*) permits types to moreover depend on the *thread identifier* of the thread that claims ownership. This is used for types like `&Cell` that cannot be sent to another thread. In other words, *ownership is (in general) thread-relative*. As we explained in §1.2, this provides a very natural way of modeling *Send*: Semantically speaking, a type  $\tau$  is *Send* if  $\llbracket \tau \rrbracket.\text{own}$  does *not* depend on the thread id. We will see more details about this in §6.1, when we give the interpretation of `Cell`, a type that cannot be shared across threads.

#### 4.2 Interpreting Types

Now that we have a semantic domain of types, we can define their semantic interpretation as a function from syntactic types  $\tau$  into the semantic domain. In this paper, we focus on the most representative types. The full interpretation can be found in the technical appendix [Jung et al. 2017a].

*Booleans*. To get started, let us consider a very simple type: `bool`. It should not come as a surprise that  $\llbracket \text{bool} \rrbracket.\text{size} := 1$ . The semantic predicate of a Boolean is defined as follows:

$$\llbracket \text{bool} \rrbracket.\text{own}(t, \bar{v}) := \bar{v} = [\text{true}] \vee \bar{v} = [\text{false}]$$

In other words, a Boolean can only be a singleton list (which is already expressed by its size), and that list has to contain either **true** or **false**.

Unsurprisingly, the semantic interpretation of integers is similar and equally straightforward.

*Products.* Given two types  $\tau_1$  and  $\tau_2$ , we define the semantics of their binary product  $\tau_1 \times \tau_2$  as that of the two types laid out one after the other in memory. This definition can be iterated to yield the interpretation of  $n$ -ary products.

For the size, we have  $\llbracket \tau_1 \times \tau_2 \rrbracket.\text{size} := \llbracket \tau_1 \rrbracket.\text{size} + \llbracket \tau_2 \rrbracket.\text{size}$ . The semantic predicate associated with  $\tau_1 \times \tau_2$  uses *separating conjunction* ( $P * Q$ ), the defining feature of separation logic, to join the semantic predicates of both types. The separating conjunction ensures that they describe ownership of disjoint pieces of memory. (Here,  $\#$  is list concatenation.)

$$\llbracket \tau_1 \times \tau_2 \rrbracket.\text{own}(t, \bar{v}) := \exists \bar{v}_1, \bar{v}_2. \bar{v} = \bar{v}_1 \# \bar{v}_2 * \llbracket \tau_1 \rrbracket.\text{own}(t, \bar{v}_1) * \llbracket \tau_2 \rrbracket.\text{own}(t, \bar{v}_2)$$

*Owned pointers.* In order to give a semantic interpretation to the type  $\mathbf{own}_n \tau$  of owned pointers, we use the standard *points-to* proposition of separation logic,  $\ell \mapsto \bar{v}$ . It states that, starting at location  $\ell$ , the memory contains the values  $\bar{v}$ , and asserts *ownership* of this memory region. With this ingredient, the interpretation is given by  $\llbracket \mathbf{own}_n \tau \rrbracket.\text{size} = 1$  and the following semantic predicate:

$$\llbracket \mathbf{own}_n \tau \rrbracket.\text{own}(t, \bar{v}) := \exists \ell. \bar{v} = [\ell] * \exists \bar{w}. \ell \mapsto \bar{w} * \triangleright \llbracket \tau \rrbracket.\text{own}(t, \bar{w}) * \triangleright \text{DeallocSize}(\ell, n, \llbracket \tau \rrbracket.\text{size})$$

Rust supports recursive types whenever the recursive occurrence is below a pointer indirection. To properly model this using Iris’s guarded recursive definitions, we have to make sure that all uses of  $\tau$  are guarded—in this case, by adding a  $\triangleright$ .

The proposition  $\text{DeallocSize}(\ell, n, \llbracket \tau \rrbracket.\text{size})$  in the semantic predicate above manages the right to deallocate the location  $\ell$ . These details can be found spelled-out in our technical appendix [Jung et al. 2017a].

*Mutable references.* Mutable references, like owned pointers, are unique pointers to something of type  $\tau$ . The key difference is that mutable references are *borrowed*, not owned, and hence they come with a lifetime indicating when they expire. In standard separation logic, an assertion always represents ownership of some part of the heap, for an unlimited duration (or until the owner actively decides to give it to another party). Instead, a mutable reference in Rust represents ownership *for a limited period of time*. When this lifetime of the reference is over, a mutable reference becomes useless, because the original owner gets back the full ownership.

To handle this new notion of “ownership with an expiry date”, we developed a custom logic for reasoning about lifetimes and borrowing. It is called the *lifetime logic*. This logic is embedded and proven correct in Iris, and we describe it in §5. Most importantly, for an Iris assertion  $P$  and a lifetime  $\kappa$ , the lifetime logic defines an assertion  $\&_{\text{full}}^{\kappa} P$ , called a *full borrow*, representing ownership of  $P$  for the duration of lifetime  $\kappa$ . Using full borrows, the interpretation of the type of mutable references is as follows:

$$\llbracket \&_{\text{mut}}^{\kappa} \tau \rrbracket.\text{size} := 1 \quad \llbracket \&_{\text{mut}}^{\kappa} \tau \rrbracket.\text{own}(t, \bar{v}) := \exists \ell. \bar{v} = [\ell] * \&_{\text{full}}^{\llbracket \kappa \rrbracket} (\exists \bar{w}. \ell \mapsto \bar{w} * \llbracket \tau \rrbracket.\text{own}(t, \bar{w}))$$

This is very similar to the interpretation of  $\mathbf{own}_n \tau$ , except that the assertion describing ownership of the contents of the reference ( $\exists \bar{w}. \ell \mapsto \bar{w} * \llbracket \tau \rrbracket.\text{own}(t, \bar{w})$ ) is wrapped in a full borrow (at lifetime  $\kappa$ ) instead of being owned directly. Finally, it turns out that  $\&_{\text{full}}^{\llbracket \kappa \rrbracket} P$  already functions as a guard of  $P$ , so there is no need for us to add any extra later modality  $\triangleright$ .

### 4.3 Interpreting Shared References

The interpretation of shared references  $\&_{\text{shr}}^{\kappa} \tau$  requires more work than the types we considered so far. Usually, we would proceed as we did above: Define  $\llbracket \&_{\text{shr}}^{\kappa} \tau \rrbracket.\text{own}$  based on  $\llbracket \tau \rrbracket.\text{own}$  such that

all the typing rules for  $\&_{\text{shr}}^{\kappa} \tau$  work out. Most of the time, this does not leave much room for choice; the primitive operations available for the type almost define it uniquely. This is decidedly not the case for shared references, for it turns out that, in Rust, there are hardly any primitive operations on  $\&T$ . The only properties that hold for  $\&T$  in general is that it can be created from a  $\&\text{mut } T$ , it is `Copy`, it has size 1, and its values have to be memory locations. However, as we have seen in §2, types like `Cell` or `Mutex` *do* provide some very interesting operations on shared references, e.g., providing indirect mutable access through a shared reference.

To account for this freedom, we permit every type to pick its own *sharing predicate*. We then use the sharing predicate of  $\tau$  to define  $\llbracket \&_{\text{shr}}^{\kappa} \tau \rrbracket.\text{own}$ . This permits, for every type, a different set of operations on its shared references. For example, the sharing predicate for basic types like `bool` allows read-only access, while the sharing predicate for `Mutex<T>` allows read and write accesses to the underlying object of type `T` once the lock has been acquired.

More formally, we extend the semantic domain of types and associate to each of them another predicate  $\llbracket \tau \rrbracket.\text{shr} \in Lft \times TId \times Loc \rightarrow iProp$ , and use it directly to model shared references:

$$\llbracket \&_{\text{shr}}^{\kappa} \tau \rrbracket.\text{size} := 1 \quad \llbracket \&_{\text{shr}}^{\kappa} \tau \rrbracket.\text{own}(t, \bar{v}) := \exists \ell. \bar{v} = [\ell] * \llbracket \tau \rrbracket.\text{shr}(\llbracket \kappa \rrbracket, t, \ell)$$

The  $\llbracket \tau \rrbracket.\text{shr}$  predicate takes three parameters: the lifetime  $\kappa$  of the shared reference, the thread identifier  $t$ , and the location  $\ell$  constituting the shared reference itself. Just like `Send` expresses that  $\llbracket \tau \rrbracket.\text{own}$  does not actually depend on the thread identifier (see §4.1), we define `Sync` to mean that  $\llbracket \tau \rrbracket.\text{shr}$  does not depend on the thread identifier. To support the aforementioned primitive operations on  $\&T$ , the sharing predicate has to satisfy the following properties:

$$\begin{aligned} & \text{persistent}(\llbracket \tau \rrbracket.\text{shr}(\kappa, t, \ell)) && \text{(TY-SHR-PERSIST)} \\ & (\&_{\text{full}}^{\kappa} (\exists \bar{w}. \ell \mapsto \bar{w} * \llbracket \tau \rrbracket.\text{own}(t, \bar{w})) * [\kappa]_q) \Rightarrow \star (\llbracket \tau \rrbracket.\text{shr}(\kappa, t, \ell) * [\kappa]_q) && \text{(TY-SHARE)} \\ & \kappa' \sqsubseteq \kappa \wedge \llbracket \tau \rrbracket.\text{shr}(\kappa, t, \ell) \Rightarrow \llbracket \tau \rrbracket.\text{shr}(\kappa', t, \ell) && \text{(TY-SHR-MONO)} \end{aligned}$$

First, `TY-SHR-PERSIST` requires that  $\llbracket \tau \rrbracket.\text{shr}$  be persistent, which implies that  $\llbracket \&_{\text{shr}}^{\kappa} \tau \rrbracket.\text{own}(t, \bar{v})$  is persistent. This corresponds to the fact that, in Rust, shared references are always `Copy`.

Second, `TY-SHARE` asserts that shared references can be created from mutable references: This is the main ingredient for proving the rule `C-SHARE` of the type system. Looking at this rule more closely, its first premise is a full borrow of an owned pointer to  $\tau$ . This is exactly  $\llbracket \&_{\text{mut}}^{\kappa} \tau \rrbracket.\text{own}(t, [\ell])$ . Its second premise is a *lifetime token*  $[\kappa]_q$ , which, as we will explain in §5, witnesses that the lifetime is alive and permits accessing borrows. Given these premises, `TY-SHARE` states that we can perform an *update*, denoted by the Iris connective  $\Rightarrow \star$ .<sup>7</sup> This update will safely transform the resources described by the premises into those described by the conclusion, namely  $\tau$ 's sharing predicate along with the same lifetime token that was passed in.

Third, `TY-SHR-MONO` requires that  $\llbracket \tau \rrbracket.\text{shr}$  be monotone with respect to the lifetime parameter. This is important for proving the subtyping rule `T-BOR-LFT`.

The addition of the sharing predicate completes our description of the semantic domain of types: Each type  $\tau$  is interpreted by a tuple  $\llbracket \tau \rrbracket = (\text{size}, \text{own}, \text{shr})$  of a natural number and two Iris predicates that satisfy `TY-SIZE`, `TY-SHR-PERSIST`, `TY-SHARE` and `TY-SHR-MONO`. Let us now go back to the types we already considered above and define their sharing predicates.

*Sharing predicate for products.* The sharing predicate for products is simply the separating conjunction of the sharing predicates of the two components:

$$\llbracket \tau_1 \times \tau_2 \rrbracket.\text{shr}(\kappa, t, \ell) := \llbracket \tau_1 \rrbracket.\text{shr}(\kappa, t, \ell) * \llbracket \tau_2 \rrbracket.\text{shr}(\kappa, t, \ell + \llbracket \tau_1 \rrbracket.\text{size})$$

The location used for the second component is shifted by  $\llbracket \tau_1 \rrbracket.\text{size}$ , reflecting the memory layout.

<sup>7</sup>The connective  $P \Rightarrow \star Q$  is in fact a shorthand for  $P * \Rightarrow Q$  in Iris [Jung et al. 2017b].

*Sharing predicate for simple types.* It turns out that there is a common pattern for defining the sharing predicates of many basic types: Indeed, when no interior mutability is at play, a shared reference provides read-only access. This sharing predicate can be used for any `Copy` type  $\tau$  of size 1. In this case,  $[[\tau]].\text{own}(t, \bar{v})$  can be written in the following form:

$$[[\tau]].\text{own}(t, \bar{v}) = \exists v. \bar{v} = [v] * \Phi_\tau(t, v)$$

where  $\Phi_\tau$  is a persistent predicate. This is the case, for example, for `bool`, `int`, function types, and shared references  $\&_{\text{shr}}^\kappa \tau$  themselves. For these types, we use the following sharing predicate:

$$[[\tau]].\text{shr}(\kappa, t, \ell) := \exists v. \&_{\text{frac}}^\kappa(\lambda q. \ell \xrightarrow{q} v) * \triangleright \Phi_\tau(t, v)$$

This definition says that there exists a fixed value  $v$  (the current value the reference points to) such that  $\Phi_\tau$  holds under the later modality  $\triangleright$  (recall that shared references are pointers, and hence occurrences of  $\tau$  need to be guarded to enable construction of recursive types), and that we have a *fractured borrow*  $\&_{\text{frac}}^\kappa(\lambda q. \ell \xrightarrow{q} v)$  of the ownership of the memory cell.

Fractured borrows are another notion provided by the lifetime logic: Similarly to full borrows, they represent temporary ownership of some resource, limited by a given lifetime. The difference is that they are persistent, but only grant *some fraction* of the content. Fortunately, that is all that is needed in order to support a read of the shared reference.

## 5 LIFETIME LOGIC

In §4, we gave a semantic model for  $\lambda_{\text{Rust}}$  types, but we left some important notions undefined. In particular, we used the notion of a *full borrow*  $\&_{\text{full}}^\kappa P$  in the interpretation of mutable references to reflect that this kind of ownership is temporary and will “expire” when lifetime  $\kappa$  ends; we mentioned *lifetime tokens*  $[\kappa]_q$  as a resource used to witness that a lifetime is ongoing; and we employed *fractured borrows*  $\&_{\text{frac}}^\kappa \Phi$  in the sharing predicate of simple types.

In this section, we describe the *lifetime logic*, a library we have developed in Iris to support these notions. In the paper, we focus on discussing the proof rules provided by the library and show how the lifetime logic can be used to model temporary and potentially shared ownership of Iris resources. More details can be found in our technical appendix and in our Coq development [Jung et al. 2017a].

We start by presenting the two core notions of *lifetimes* and *full borrows* in §5.1. We then continue in §5.2, explaining how lifetimes can be *compared* and *intersected*. Finally, in §5.3, we present *fractured borrows*, which we have already seen as being useful for defining sharing predicates.

### 5.1 Full Borrows and Lifetime Tokens

Figure 4 shows the main rules of the lifetime logic. We explain them by referring to the following Rust example, similar to the one in §2.4:

```
1 let mut v = Vec::new(); v.push(0);
2 { let mut head = v.index_mut(0); *head = 23; }
3 println!("{:?}", v);
```

Recall the type of `index_mut`: `for<'a> fn(&'a mut Vec<i32>, usize) -> &'a mut i32`. To call this function, we need a borrow at some lifetime  $\kappa$  (which we will use to instantiate `'a`). To get started, we need to create this lifetime. This is the role of `LFTL-BEGIN`: it lets us perform an Iris update to create a fresh lifetime  $\kappa$  and gives us the full *lifetime token*  $[\kappa]_1$  witnessing that this lifetime is ongoing. (This token can then be split into fractional lifetime tokens  $[\kappa]_q$ —see below.) It

<p><b>LFTL-BEGIN</b>  <math>\text{True} \Rightarrow \star \exists \kappa. [\kappa]_1 * ([\kappa]_1 \Rightarrow \star [\dagger\kappa])</math></p>	<p><b>LFTL-TOK-FRACT</b>  <math>[\kappa]_{q+q'} \Leftrightarrow [\kappa]_q * [\kappa]_{q'}</math></p>	<p><b>LFTL-NOT-OWN-END</b>  <math>[\kappa]_q * [\dagger\kappa] \Rightarrow \text{False}</math></p>
<p><b>LFTL-END-PERSIST</b>  <math>\text{persistent}([\dagger\kappa])</math></p>	<p><b>LFTL-BORROW</b>  <math>\triangleright P \Rightarrow \star \&amp;_{\text{full}}^{\kappa} P * ([\dagger\kappa] \Rightarrow \star \triangleright P)</math></p>	<p><b>LFTL-BOR-SPLIT</b>  <math>\&amp;_{\text{full}}^{\kappa} (P * Q) \Rightarrow \star \&amp;_{\text{full}}^{\kappa} P * \&amp;_{\text{full}}^{\kappa} Q</math></p>
<p><b>LFTL-BOR-ACC</b>  <math>\&amp;_{\text{full}}^{\kappa} P * [\kappa]_q \Rightarrow \star \triangleright P * (\triangleright P \Rightarrow \star \&amp;_{\text{full}}^{\kappa} P * [\kappa]_q)</math></p>	<p><b>LFTL-BOR-SHORTEN</b>  <math>\kappa' \sqsubseteq \kappa * \&amp;_{\text{full}}^{\kappa} P \Rightarrow \&amp;_{\text{full}}^{\kappa'} P</math></p>	<p><b>LFTL-INCL-ISECT</b>  <math>\kappa \sqcap \kappa' \sqsubseteq \kappa</math></p>
<p><b>LFTL-INCL-GLB</b>  <math>\kappa \sqsubseteq \kappa' * \kappa \sqsubseteq \kappa'' \Rightarrow \kappa \sqsubseteq \kappa' \sqcap \kappa''</math></p>	<p><b>LFTL-TOK-INTER</b>  <math>[\kappa \sqcap \kappa']_q \Leftrightarrow [\kappa]_q * [\kappa']_q</math></p>	<p><b>LFTL-END-INTER</b>  <math>[\dagger\kappa \sqcap \kappa'] \Leftrightarrow [\dagger\kappa] \vee [\dagger\kappa']</math></p>
<p><b>LFTL-TOK-UNIT</b>  <math>\text{True} \Rightarrow [\varepsilon]_q</math></p>	<p><b>LFTL-END-UNIT</b>  <math>[\dagger\varepsilon] \Rightarrow \text{False}</math></p>	<p><b>LFTL-REBORROW</b>  <math>\kappa' \sqsubseteq \kappa * \&amp;_{\text{full}}^{\kappa} P \Rightarrow \star \&amp;_{\text{full}}^{\kappa'} P * ([\dagger\kappa'] \Rightarrow \star \&amp;_{\text{full}}^{\kappa} P)</math></p>

Fig. 4. Selected rules of the lifetime logic.

also provides the update  $[\kappa]_1 \Rightarrow \star [\dagger\kappa]$ : we will use this update later to end  $\kappa$  by exchanging the full lifetime token  $[\kappa]_1$  for a *dead token*, written  $[\dagger\kappa]$ , indicating that  $\kappa$  has ended.<sup>8</sup>

Once the lifetime has been created, we can borrow the vector  $\mathbf{v}$  at the lifetime  $\kappa$  in order to pass a borrowed reference to `index_mut`. This is allowed by **LFTL-BORROW**, really the core rule of the lifetime logic. This rule splits ownership of a resource  $P$  (in our example, the vector  $\mathbf{v}$ ) into the separating conjunction of a *full borrow*  $\&_{\text{full}}^{\kappa} P$  and an *inheritance*  $[\dagger\kappa] \Rightarrow \star \triangleright P$ . The borrow grants access to  $P$  *during* the lifetime  $\kappa$ , while the inheritance allows us to retrieve ownership of  $P$  *after*  $\kappa$  has ended. In other words, **LFTL-BORROW** *splits ownership in time*. The separating conjunction indicates that the two operands are “disjoint”, which means we can safely transfer ownership of the borrow to `index_mut` and keep ownership of the inheritance for ourselves to use later. Except here, this is not disjointness in space (e.g., in the memory), since both the borrow and the inheritance grant access to the same shared resource. Rather, it is disjointness *in time*: The lifetime  $\kappa$  is *either* ongoing *or* ended, so the borrow and the inheritance are never useful at the same time.

We do not give the actual implementation of `index_mut` in this paper. However, here is what `index_mut` does with respect to ownership. First, the ownership of the memory used by the vector (“inside” the full borrow) is split into two parts: (1) The ownership of the accessed vector position, and (2) the ownership of the rest of the vector. Then, the rule **LFTL-BOR-SPLIT** is used to split the full borrow into two full borrows dedicated to each of these parts. The full borrow of part (1) is returned to the caller; this matches the return type of `index_mut`. On the other hand, the full borrow of part (2) is dropped.<sup>9</sup> This means that the ownership of the rest of the vector is effectively lost until the lifetime ends, at which point it can be recovered using the inheritance.

The next step of our program is the write to `*head` on line 2. Recall that the type of `head` is `&mut i32`, which represents ownership of a full borrow of a single memory location. In order to perform this write, we need to *access* this full borrow and get the resource it contains (in particular, the maps-to predicate  $\ell \mapsto \bar{v}$ ). This is what **LFTL-BOR-ACC** does: If we give it a full borrow  $\&_{\text{full}}^{\kappa} P$  and a lifetime token  $[\kappa]_q$ , witnessing that  $\kappa$  is alive, then we get the resource  $P$ . Moreover, we also

<sup>8</sup>Note that the ending update uses an “update that takes a step”  $\Rightarrow \star$  rather than a normal update  $\Rightarrow$ . This connective, which is defined in the appendix [Jung et al. 2017a] and is required for technical reasons related to step-indexing, restricts the update to only be used in conjunction with reasoning about a physical step of computation.

<sup>9</sup>Iris is an *affine* logic, in which it is possible to give up ownership of resources at any time, i.e., Iris has the law  $P * Q \vdash P$ .



get the update  $\triangleright P \Rightarrow^* \&_{\text{full}}^{\kappa} P * [\kappa]_q$ : This can later be used when we are done with  $P$ , in order to reconstitute the full borrow and get back the lifetime token. In our proofs, we will always be forced to give back all the lifetime tokens that we obtained; this makes sure that we properly close all borrows again. This can be seen, for example, in **TY-SHARE**: A token is provided as a premise to this update, but the same token must also be returned again in the conclusion.

Finally, at the end of line 2 of our example, **head** goes out of scope and it is time to end  $\kappa$ . To this end, we apply the update  $[\kappa]_1 \Rightarrow^* [\dagger\kappa]$  that we obtained when  $\kappa$  was created. In doing so, we have to give up the lifetime token  $[\kappa]_1$  (ensuring that all borrows are closed again), but we get back the dead token  $[\dagger\kappa]$ , which can be used to prove that  $\kappa$  has indeed ended. Now that  $\kappa$  has ended, we can use our inheritance  $[\dagger\kappa] \Rightarrow^* \triangleright P$  to get back the ownership of  $v$  before printing it. Note that the dead token  $[\dagger\kappa]$  is persistent (**LFTL-END-PERSIST**), so it can be used multiple times—this is important since there may be many borrows (and thus many inheritances we wish to use) at the same lifetime. Each inheritance, however, may only be used once.

One important feature of the lifetime logic that this example does not demonstrate is the parameter  $q$ , a fraction. Lifetime tokens can always be split into smaller parts, in a reversible fashion (**LFTL-TOK-FRACT**). This is needed when we want to access several full borrows with the same lifetime at the same time, or to witness that a lifetime is ongoing in several threads simultaneously. Moreover, unsurprisingly, a lifetime cannot be both dead and alive at the same time (**LFTL-NOT-OWN-END**).

## 5.2 Lifetime Inclusion

In §2 and §3, we have seen that Rust relates lifetimes by *lifetime inclusion*. This is used for subtyping (**T-BOR-LFT**) and reborrowing (**C-REBORROW**).

What does it mean for a lifetime  $\kappa$  to be “included” in another  $\kappa'$ ? The key property of lifetime inclusion is that when the shorter  $\kappa$  is still alive, then so is the longer  $\kappa'$ . From the perspective of lifetime tokens, this means that, given a token for  $\kappa$ , we should be able to obtain a token for  $\kappa'$ . Conversely, given a dead token for  $\kappa'$ , we should be able to obtain a dead token for  $\kappa$ , as well. This is reflected in the definition of lifetime inclusion:

$$\kappa \sqsubseteq \kappa' \quad := \quad \square \left( \left( \forall q. [\kappa]_q \Rightarrow^* \exists q'. [\kappa']_{q'} * ([\kappa']_{q'} \Rightarrow^* [\kappa]_q) \right) * \left( [\dagger\kappa'] \Rightarrow^* [\dagger\kappa] \right) \right)$$

The first part says that we can trade a fraction of the token of  $\kappa$  for a potentially different fraction of the token of  $\kappa'$ . It also provides a way to revert this trading to recover the original token of  $\kappa$ , so that no token is permanently lost. The second part of this definition is the analogue for dead tokens. Note that since dead tokens are persistent, it is not necessary to provide a way to recover the dead token that is passed in. The entire definition is wrapped in Iris’s persistence modality  $\square$  to make lifetime inclusion a persistent assertion that can be reused as often as needed.

It is easy to show that lifetime inclusion is a preorder. Inclusion can be used to *shorten* a full borrow (**LFTL-BOR-SHORTEN**): If a full borrow is valid for a long lifetime, then it should also be valid for the shorter one. This rule justifies subtyping based on lifetimes in  $\lambda_{\text{Rust}}$ .

An even stronger use of lifetime inclusion is *reborrowing*, expressed by **LFTL-REBORROW**. This rule is used to prove the reborrowing rule in the type system, **C-REBORROW**. Unlike shortening, reborrowing provides an inheritance to regain the initial full borrow after the shorter lifetime has ended. This may sound intuitively plausible, but turns out to be extremely subtle. In fact, most of the complexity in the model of the lifetime logic arises from reborrowing.

**5.2.1 Lifetime intersection.** Beyond having a preorder, it turns out that lifetimes also have a greatest lower bound: Given two lifetimes  $\kappa$  and  $\kappa'$ , their *intersection*  $\kappa \sqcap \kappa'$  is the lifetime that ends whenever either of the operands ends.

$$\begin{array}{c}
\text{LFTL-BOR-FRACTURE} \\
& \&_{\text{full}}^{\kappa} \Phi(1) \equiv \star \&_{\text{frac}}^{\kappa} \Phi \\
\text{LFTL-FRACT-ACC} \\
& \frac{\forall q_1, q_2. \Phi(q_1 + q_2) \Leftrightarrow \Phi(q_1) * \Phi(q_2)}{\&_{\text{frac}}^{\kappa} \Phi * [\kappa]_q \equiv \star \exists q'. \triangleright \Phi(q') * (\triangleright \Phi(q') \equiv \star [\kappa]_q)}
\end{array}$$

Fig. 5. Selected rules for fractured borrows.

Lifetime intersection is particularly useful to create a fresh lifetime that is a sublifetime of some existing  $\kappa$ . We invoke the rule **LFTL-BEGIN** to create an auxiliary lifetime  $\alpha_0$ , and then we use the intersection  $\alpha := \alpha_0 \sqcap \kappa$  as our new lifetime. It follows that  $\alpha \sqsubseteq \kappa$ . In the type system, we use this in the proof of **F-NEWLFT** to create a new lifetime  $\alpha$  that is shorter than all the lifetimes in  $\bar{\kappa}$ .

Intersection of lifetimes interacts well with lifetime tokens: A token of the intersection is composed of tokens of both operands, at the same fraction (**LFTL-TOK-INTER**). In other words, in order to prove that an intersection is alive, we have to prove that both operands are alive. Similarly, in order to prove that an intersection has ended, it suffices to prove that either operand has ended (**LFTL-END-INTER**). These laws let us do the token trading required by lifetime inclusion, showing that intersection indeed is the greatest lower bound for  $\sqsubseteq$  (**LFTL-INCL-ISECT**, **LFTL-INCL-GLB**).

Furthermore, intersection has a unit  $\varepsilon$ . This lifetime never ends (**LFTL-END-UNIT**) and we can freely get tokens for it (**LFTL-TOK-UNIT**). We use  $\varepsilon$  to model the **static** lifetime.

### 5.3 Fractured Borrows

Full borrows and lifetimes are powerful tools for modeling temporary ownership in Iris. However, they cannot be used as-is for modeling Rust's shared references. In §4.3, we used the notion of *fractured borrows* as our key notion for defining the default read-only sharing predicate. Figure 5 gives the main reasoning rules for fractured borrows.

To make it possible to use them as a sharing predicate, fractured borrows are persistent and, just like full borrows (**LFTL-BOR-SHORTEN**), they can be shortened. Because they are persistent, fractured borrows can potentially be accessed simultaneously by several parties. As such, they cannot provide access to the full underlying resource. Instead, **LFTL-FRACT-ACC** provides access only to *some fraction* of the borrowed content.

To express this, fractured borrows work on a *predicate*  $\Phi$  over fractions that has to be compatible with addition:  $\Phi(q_1 + q_2) \Leftrightarrow \Phi(q_1) * \Phi(q_2)$ . When using **LFTL-FRACT-ACC** to access the content of the fractured borrow, we get  $\Phi(q)$  for some unknown fraction  $q$ . This works because no matter how many threads access the same fractured borrow at the same time, it is always possible to give out *some* tiny fraction of  $\Phi$  and keep some remainder available for the next thread. Similarly to full borrows, **LFTL-FRACT-ACC** requires a lifetime token for witnessing that the lifetime is alive, and gives back the lifetime token only when the resource is returned.

Fractured borrows can be created from a full borrow of  $\Phi(1)$  using the **LFTL-BOR-FRACTURE** rule.

**5.3.1 Lifetime inclusion and fractured borrows.** Fractured borrows have an interesting interaction with lifetime inclusion. Assume we have a fractured borrow of lifetime token for another lifetime  $\kappa'$ . That is, assume  $\Phi(q') = [\kappa']_{q'}$ . The rule **LFTL-FRACT-ACC** for accessing fractured borrows turns out to be exactly the first part of the token trading scheme that we used for defining the lifetime inclusion  $\kappa \sqsubseteq \kappa'$ . In fact, by using some further properties of fractured borrows (see our technical appendix [Jung et al. 2017a]), we can also prove the trading scheme for dead tokens, so that we have:

$$\&_{\text{frac}}^{\kappa} \lambda q'. [\kappa']_{q'} \Rightarrow \kappa \sqsubseteq \kappa'$$

This can be generalized to fracturing just a part of the token, *i.e.*,  $\Phi(q') = [\kappa']_{q'.q}$ . The reason this makes sense is that with the token of a longer lifetime  $\kappa'$  being borrowed at some shorter lifetime  $\kappa$ , it is impossible to end  $\kappa'$  while  $\kappa$  is still ongoing: Ending  $\kappa'$  needs the full token, but part of that token is stuck in  $\kappa$  and can only be recovered through an inheritance.

Deriving lifetime inclusions from fractured borrows significantly expands the power of lifetime inclusion. So far, we have seen that we can use lifetime intersection to make a fresh  $\alpha$  a sublifetime of some existing  $\kappa$ ; however, for this to work out, we have to decide *in advance* which other lifetimes  $\alpha$  is going to be a sublifetime of. Using fractured borrows, we can establish additional lifetime inclusions *dynamically*, when the involved lifetimes are already ongoing and in active use. It turns out that interior mutable types like `RefCell<T>` or `RwLock<T>` allow sharing data structures for a lifetime that cannot be established in advance, and we thus found this new scheme for proving lifetime inclusion crucial in proving the safety of such types.

## 6 MODELING TYPES WITH INTERIOR MUTABILITY

As we have discussed in §2.5, the standard library of Rust provides types with *interior mutability*. These types, written in Rust using `unsafe` features, can nonetheless be used safely because the interface they provide to client code encapsulates these unsafeties behind well-typed abstractions. We have proven the safety of several such libraries, namely: `Cell`, `RefCell`, `Mutex`, `RwLock`, `Rc`, and `Arc`.<sup>10</sup> To fulfill this goal, we had to first pick semantic interpretations for the abstract types exported by these libraries (*e.g.*, `Cell<T>`). We then proved that each publicly exported function from these libraries satisfies the semantic interpretation of its type.

Usually, when modeling types with interior mutability, the most difficult definition is that of the sharing predicate  $[[\tau]].shr$ . Indeed, these types use a sharing predicate which is different from the default, read-only one that we described in §4.3. The sharing predicates vary greatly depending on which operations are allowed. Most of them use a new variant of borrow propositions, called *persistent borrows*, which we present in this section. As it turns out, all the variants of borrow propositions (including full and fractured borrows) are encodable in terms of a single internal mechanism, called *indexed borrows*, but the explanation of this encoding would take us too far afield (details are explained in the technical appendix [Jung et al. 2017a]). We focus our explanations on two representative forms of interior mutability that we have already presented in §2: `Cell` and `Mutex`.

### 6.1 Cell

In §2.5.1, we have seen that `Cell<T>` stores values of type `T` and provides two functions: `get` and `set`, which can be used for reading from and writing to the cell. It turns out that ownership and size of `cell( $\tau$ )`, the equivalent of `Cell<T>` in  $\lambda_{\text{Rust}}$ , are the same as  $\tau$ . In fact, Rust's standard library provides two functions for converting between `T` and `Cell<T>`, `Cell::new` and `Cell::into_inner`, both of which are effectively the identity function.

The sharing predicate is where things get interesting. Remember that `get` and `set` can be called *even if you only have a shared reference to a `Cell<T>`*. This means that `Cell<i32>` must use a very different sharing predicate than `i32`, which just provides read-only access. In contrast, to verify `set`, we need temporary full access for the duration of the function call. However, it is also important that all shared references to a `Cell` are confined to a single thread, since the `get` and `set` operations are not thread-safe. Recall that Rust enforces this by declaring that `Cell` is not `Sync`, which is

<sup>10</sup>Note that some simplifications of our setup make the proof of some of these libraries simpler. More precisely, we are not handling unwinding after panics, and all atomic memory operations are sequentially consistent, while Rust's standard library uses weaker atomic accesses.

$$\begin{array}{ll}
\text{LFTL-BOR-NA} & \text{LFTL-NA-ACC} \\
& \&_{\text{full}}^{\kappa} P \Rightarrow \star \&_{\text{na}}^{\kappa/t} P & \&_{\text{na}}^{\kappa/t} P * [\kappa]_q * [\text{Na} : t] \Rightarrow \star \triangleright P * \left( \triangleright P \Rightarrow \star [\kappa]_q * [\text{Na} : t] \right)
\end{array}$$

Fig. 6. Selected rules for non-atomic persistent borrows.

equivalent to saying that `&Cell` is not `Send`, so that shared references to it cannot be sent to another thread—they must stay in the thread they have initially been created in.

In order to encode this idea, we use *non-atomic persistent borrows*, another kind of borrow derived from the lifetime logic. Some of their rules are presented in Figure 6. Like fractured borrows, non-atomic persistent borrows are persistent, can be created from full borrows, and support shortening. However, the rule to access the borrows is different: `LFTL-NA-ACC` gives *full* access to the borrowed content, so it is important that concurrent threads not be allowed to access the same borrow simultaneously. To this end, the borrows depend on a thread identifier  $t$ . Accessing them requires a *non-atomic token*  $[\text{Na} : t]$  bound to that thread identifier. This token is created at the birth of the thread, and threaded through all of its control flow. That is, every function receives it and has to return it. The token is required to open a borrow, and not returned until the borrow is closed, making it impossible to open it twice at the same time.

We can now use non-atomic persistent borrows to give the sharing predicate of `cell`( $\tau$ ):

$$\llbracket \text{cell}(\tau) \rrbracket.\text{shr}(\kappa, t, \ell) := \&_{\text{na}}^{\kappa/t} (\exists \bar{v}. \ell \mapsto \bar{v} * \llbracket \tau \rrbracket.\text{own}(t, \bar{v}))$$

Note, in particular, that our model of `Cell`< $T$ > reflects the fact that it is never `Sync`, since its sharing predicate depends on the thread identifier  $t$ . However, if  $\llbracket \tau \rrbracket.\text{own}(t, \bar{v})$  does not depend on  $t$ , then neither does  $\llbracket \text{cell}(\tau) \rrbracket.\text{own}(t, \bar{v})$ —just like in Rust, `Cell`< $T$ > is `Send` if and only if  $T$  is.

## 6.2 Mutex

`Mutex` is the other example of interior mutability that we presented in §2.5. `Mutex`< $T$ > uses a lock to safely grant multiple threads read and write access to a shared object of type  $T$ .

We start by giving its size and ownership predicate:

$$\llbracket \text{mutex}(\tau) \rrbracket.\text{size} := 1 + \llbracket \tau \rrbracket.\text{size} \quad \llbracket \text{mutex}(\tau) \rrbracket.\text{own}(t, \bar{v}) := \llbracket \text{bool} \times \tau \rrbracket.\text{own}(t, \bar{v})$$

That is, when it is not shared, `mutex`( $\tau$ ) is exactly the same as a pair of a `bool` (representing the status of the lock<sup>11</sup>), and of an object of type  $\tau$  (the content).

The sharing predicate is more complex: It cannot use fractured borrows, because we cannot afford getting only a fraction of ownership, and it cannot use non-atomic persistent borrows, because mutexes are thread-safe. Instead, it uses yet another kind of borrow, *atomic persistent borrows*, whose rules can be found in the appendix [Jung et al. 2017a]. Again, they are distinguished from the other borrows in the rule granting access to the borrowed content. Here, the mechanism used to prevent two threads accessing the same borrow at the same time is *atomicity*: The proof rules enforce that an atomic persistent borrow cannot be opened for longer than a single, atomic instruction. Thus, during the execution of any given instruction, only one thread can be accessing the borrow. Returning to `Mutex`'s sharing predicate, the content of its borrow will only get accessed when changing the status of the lock, and doing so will require atomic memory accesses. Of course, this corresponds to the fact that, in our spinlock implementation, we are only using atomic sequentially consistent instructions to read or write the status flag. Using non-atomic accesses would lead to data races.

<sup>11</sup>The actual implementation in Rust uses the locking primitives of the operating system. We use our own spinlock-based implementation to model that.

Using atomic persistent borrows, we can give the sharing predicate for mutexes:

$$\llbracket \mathbf{mutex}(\tau) \rrbracket.\text{shr}(\kappa, t, \ell) := \exists \kappa'. \kappa \sqsubseteq \kappa' * \\ \&_{\text{at}}^{\kappa} (\ell \mapsto \mathbf{true} \quad \vee \quad \ell \mapsto \mathbf{false} * \&_{\text{full}}^{\kappa'} (\exists \bar{v}. (\ell + 1) \mapsto \bar{v} * \llbracket \tau \rrbracket.\text{own}(t, \bar{v})))$$

This is quite a mouthful: First, we use an existential quantification at the beginning to close the predicate under shorter lifetimes, satisfying **TY-SHR-MONO**. We use an atomic persistent borrow to share ownership of the status flag at location  $\ell$ . This defines an invariant that is maintained *until*  $\kappa$  ends. The invariant can be in one of two states: In the first state, the flag is **true**, in which case the lock is locked and no other resource is stored in the borrow. Ownership of the content is currently held by whichever thread acquired the lock. In the second state, the flag is **false**. This means the lock is unlocked, and the borrow also stores the ownership of the content at type  $\tau$  at location  $\ell + 1$ . When acquiring or releasing the lock, we can atomically open the persistent borrow and change the branch of the disjunction, thus acquiring or releasing ownership of the content.

Curiously, ownership of the content is wrapped in a full borrow. One might expect instead that it should be directly contained in the outer persistent borrow. In this case, acquiring the lock would result in acquiring full (unborrowed) ownership of the content of the mutex. That, however, does not work: Imagine  $\kappa'$  ends while the lock is held. (That is possible, for example, if the `MutexGuard` is leaked and hence its destructor never gets called.) In this case, ownership of the content would never be returned to the borrow. However, when  $\kappa'$  ends, the `Mutex` is again fully owned by someone, which means *they* expect to be the exclusive owner of the content! This is why the full borrow is necessary: When taking the lock, one gets the inner resource only under a borrow at lifetime  $\kappa'$ , guaranteeing that ownership is returned when  $\kappa'$  ends.

To conclude, observe that if the inner type does not depend on the thread identifier  $t$  (which corresponds to saying that it is **Send**), then neither  $\llbracket \mathbf{mutex}(\tau) \rrbracket.\text{own}$  nor  $\llbracket \mathbf{mutex}(\tau) \rrbracket.\text{shr}$  do, so that **mutex**( $\tau$ ) is both **Send** and **Sync**. This exactly corresponds to Rust's behavior.

## 7 PROOF OF SOUNDNESS

Having defined the semantics of types in §4, we can finish up our formal development by defining semantic interpretations of the judgments presented in §3.3. We focus on the two most important ones: typing of instructions and typing of function bodies. Their interpretations use Hoare triples:

$$\Gamma \mid \mathbf{E}; \mathbf{L} \mid \mathbf{T}_1 \models I = x. \mathbf{T}_2 := \\ \forall \gamma, t. \{ \llbracket \mathbf{E} \rrbracket_{\gamma} * \llbracket \mathbf{L} \rrbracket_{\gamma} * [\mathbf{Na} : t] * \llbracket \mathbf{T}_1 \rrbracket_{\gamma}(t) \} I \{ v. \llbracket \mathbf{L} \rrbracket_{\gamma} * [\mathbf{Na} : t] * \llbracket \mathbf{T}_2 \rrbracket_{\gamma[x \leftarrow v]}(t) \} \\ \Gamma \mid \mathbf{E}; \mathbf{L} \mid \mathbf{K}; \mathbf{T} \models F := \forall \gamma, t. \{ \llbracket \mathbf{E} \rrbracket_{\gamma} * \llbracket \mathbf{L} \rrbracket_{\gamma} * [\mathbf{Na} : t] * \llbracket \mathbf{T} \rrbracket_{\gamma}(t) * \llbracket \mathbf{K} \rrbracket_{\gamma}(t) \} F \{ \mathbf{True} \}$$

In the preconditions, we can find the interpretations of the various contexts, together with the thread-local token  $[\mathbf{Na} : t]$  used to access the non-atomic persistent borrows of thread  $t$ . The instruction judgment nicely demonstrates how this token is threaded through, alongside the local lifetime context  $\llbracket \mathbf{L} \rrbracket_{\gamma}$  which contains all the lifetime tokens. The interpretation of the external lifetime context  $\llbracket \mathbf{E} \rrbracket_{\gamma}$  just involves the lifetime inclusion from the lifetime logic; this makes it persistent, so it does not have to be threaded through. Finally,  $\llbracket \mathbf{T} \rrbracket_{\gamma}(t)$  uses the semantic interpretation of types as defined in the previous sections, tying them to the current thread  $t$ .

The function judgment, on the other hand, has a trivial post-condition—remember that functions do not return, they call a continuation. The Hoare triple for that continuation is provided by  $\llbracket \mathbf{K} \rrbracket_{\gamma}(t)$ , and it will require  $[\mathbf{Na} : t]$  as well as  $\llbracket \mathbf{L} \rrbracket_{\gamma}$  in its precondition, which is how the tokens travel between functions. Using **True** as the trivial post-condition may be surprising; the more common choice for a CPS language is certainly **False**. However, for the adequacy theorem below, we want to talk about executing a full program, so we have to give a “halting continuation”. We could

of course make that continuation diverge, but instead we decided to make it return immediately. This has the benefit that the entire program actually terminates; however, it also means that we have to pick True as the post-condition of our continuations.

*Soundness.* With the semantic judgments defined, we can state two core theorems showing the soundness of our type system. The first one shows a deep connection between the semantic judgments and their syntactic counterparts.

**THEOREM 7.1 (FUNDAMENTAL THEOREM OF LOGICAL RELATIONS).** *For any inference rule of the type system, when we replace all  $\vdash$  by  $\models$ , the resulting Iris theorem holds.*

One important corollary of the fundamental theorem is that if a judgment can be derived syntactically, then it also holds semantically. However, **Theorem 7.1** is much stronger than this, because we can use it to glue together safe and `unsafe` code. Given a program that is syntactically well-typed except for certain components that are only semantically (but not syntactically) well-typed, the fundamental theorem tells us that the entire program is semantically well-typed.

The second theorem is an adequacy theorem, relating the logical relation to program behavior:

**THEOREM 7.2 (ADEQUACY).** *Let  $f$  be a  $\lambda_{\text{Rust}}$  function such that  $\emptyset \mid \emptyset; \emptyset \mid \emptyset \models f = x. x \triangleleft \text{fn}() \rightarrow \Pi[]$  holds. Then when we execute  $f$  with the default continuation (which is just a no-op), no execution ends in a stuck state.*

In particular, the adequacy theorem guarantees that a semantically well-typed program is memory and thread safe: It will never perform any invalid memory access and will not have data races.

Put together, these theorems establish that, *if the only code in a  $\lambda_{\text{Rust}}$  program that is not syntactically well-typed appears in semantically well-typed libraries, then the program is safe to execute.*

## 8 RELATED WORK

**Substructural type systems for state, and their soundness proofs.** Over the past decades, numerous languages and type systems have been developed that use linear types [Wadler 1990], ownership [Clarke et al. 1998], and/or regions [Fluet et al. 2006] to guarantee safety of heap-manipulating programs. These include Cyclone [Jim et al. 2002], Vault [DeLine and Fähndrich 2001], and Alms [Tov and Pucella 2011]. Much of this work has influenced the design of Rust, but a detailed discussion of that influence is beyond the scope of this paper. The key point for our purposes is that most such systems are *closed-world*, meaning that they are defined by a fixed set of rules and are proven sound using *syntactic* techniques [Wright and Felleisen 1994]. As explained in §1.1, Rust’s extensible type system fundamentally does not fit into this paradigm.

In a related but very different line of work, systems like Ynot [Nanevski et al. 2008], FCSL [Nanevski et al. 2014], and F\* [Swamy et al. 2016] integrate variants of separation logic into dependent type theory. These systems are aimed at full functional verification of low-level imperative code and thus require a significant amount of manual proof and/or type annotations compared to Rust.

Mezzo [Balabonski et al. 2016] can be placed somewhere between these two approaches. It comes with a substructural type system whose expressivity parallels that of a separation logic. Its soundness proof is modular in the sense that the authors start by verifying a core type system, and then add various extensions. This relies on an abstract notion of resources called *monotonic separation algebras*. Nevertheless, Mezzo’s handling of types remains entirely syntactic (e.g., based on the grammar of types); there is no semantic account for types that would permit “adding” new types without revisiting the proofs.

We are only aware of a few substructural type systems for which soundness has been proven semantically (using logical relations). These include  $L^3$  [Ahmed et al. 2007],  $\lambda^{\text{URAL}}$  [Ahmed et al.

2005], and the “superficially substructural” type system of [Krishnaswami et al. \[2012\]](#). Ahmed et al.’s motivations for doing semantic soundness proofs were somewhat different from ours. One of their motivations was to build a foundation for substructural extensions to the Foundational Proof-Carrying Code project [[Appel 2001](#)]. Another was to make it possible to modularly extend soundness proofs when building up the features of a language incrementally (although it is worth noting that Balabonski et al. achieved similarly modular proofs for Mezzo using only syntactic methods). In contrast, following [Krishnaswami et al. \[2012\]](#), we are focused on building a soundness proof that is “extensible” along a different axis, namely the ability to verify soundness of libraries that extend Rust’s core type system through their use of unsafe features. Lastly, all of the prior semantic soundness proofs were done directly using set-theoretic step-indexed models, whereas in the present work, in order to model the complexities of Rust’s lifetimes and borrowing, we found it essential to work at the higher level of abstraction afforded by Iris and our lifetime logic.

Cogent [[Amani et al. 2016](#); [O’Connor et al. 2016](#)] is a purely functional, linearly typed language designed to implement file systems and verify their functional correctness. Its linear type system permits efficient compilation to machine code using in-place updates, while the purely functional semantics enables equational reasoning. Its design is such that missing functionality can be implemented in C functions (much like `unsafe` code in Rust), which are given types to enforce correct usage in the Cogent program. These C functions are then manually verified to implement an equational specification and to follow the guarantees of the type system. However, the language and the type system are much simpler than Rust’s (e.g., there is no support for recursion, iteration, borrowing, or mutable state).

Rust’s concept of lifetimes has appeared before in the form of *regions* [[Fähndrich and DeLine 2002](#); [Grossman et al. 2002](#)]. The work by [Fluet et al. \[2006\]](#) on linear regions bears some similarity to the lifetime logic, with region capabilities corresponding to lifetime tokens and references corresponding to borrows. However, their approach does not rule out combining mutation with aliasing. This is not a problem because they consider neither deep pointers (where writing to one pointer can invalidate an aliasing pointer) nor concurrency. We believe that, to extend linear regions to handle Rust’s unique borrows, one would end up needing something akin to our lifetime logic.

**Formal results for Rust.** Patina [[Reed 2015](#)] is a formalization of the Rust type system, with accompanying partial proofs of progress and preservation. Being syntactic, these proofs do not scale to account for `unsafe` code. To keep our formalization feasible, we did not reuse the syntax and type system of Patina, but rather designed  $\lambda_{\text{Rust}}$  from scratch in a way that better fits Iris.

CRUST [[Toman et al. 2015](#)] is a bounded model checker designed to verify the safety of Rust libraries implemented using `unsafe` code. It checks that all clients calling up to  $n$  library methods do not trigger memory safety faults. This provides an easy-to-use, automated way of checking `unsafe` code, before attempting a full formal proof. Their approach has successfully re-discovered some soundness bugs that had already been fixed in Rust’s standard library. However, by only considering one library at a time, it cannot find bugs that arise from the interaction of multiple libraries [[Ben-Yehuda 2015c](#)].

**Concurrent separation logics.** RustBelt builds on the Iris framework [[Jung et al. 2017b](#)], which in turn incorporates several great advances made in the past decade in the area of concurrent separation logics [[Appel 2014](#); [Dinsdale-Young et al. 2013, 2010](#); [Dodds et al. 2009](#); [Nanevski et al. 2014](#); [O’Hearn 2007](#); [Svendsen and Birkedal 2014](#)]. In particular, RustBelt depends crucially on Iris’s support for: (1) custom notions of logical resource (i.e., “fictional separation” [[Jensen and Birkedal 2012](#)]), which we use to model novel abstract predicates like the various forms of borrow propositions; (2) impredicative invariants [[Svendsen and Birkedal 2014](#)], which we use to model higher-order state; and (3) support for tactical proofs in Coq [[Krebbers et al. 2017b](#)], without which a verification of the scale and complexity of RustBelt would not be possible.

One recent innovation in separation logics is *temporary read-only permissions* [Charguéraud and Pottier 2017]. The authors introduce a duplicable “read-only” modality with rules that resemble ours for shared references at “simple” types like `i32`. However, since shared references permit interior mutability, the read-only permission is not suited to directly modeling shared references. Nevertheless, it would be interesting to explore whether this approach can facilitate the tracking of lifetime tokens, just like read-only permissions eliminate the bookkeeping involved in fractional permissions. One challenge here is that  $\lambda_{\text{Rust}}$  supports non-lexical lifetimes [Matsakis 2016b], whereas read-only permissions are strictly lexical.

## 9 CONCLUSION

We have described  $\lambda_{\text{Rust}}$ , a formal version of the Rust type system that we used to study Rust’s ownership discipline in the presence of `unsafe` code. We have shown that various important Rust libraries with unsafe implementations, many of them involving *interior mutability*, are safely encapsulated by their type. We had to make some concessions in our modeling: We do not model (1) more relaxed forms of atomic accesses, which Rust uses for efficiency in libraries like `Arc`; (2) Rust’s trait objects (comparable to interfaces in Java), which can pose safety issues due to their interactions with lifetimes; or (3) stack unwinding when a panic occurs, which causes issues similar to exception safety in C++ [Abrahams 1998]. We proved safety of the destructors of the verified libraries, but do not handle automatic destruction, which has already caused problems [Ben-Yehuda 2015b] for which the Rust community still does not have a modular solution [Rust team 2016]. The remaining omissions are mostly unrelated to ownership, like proper support for type-polymorphic functions, and “unsized” types whose size is not statically known<sup>12</sup>.

Despite these limitations, we believe we have captured the essence of Rust’s ownership discipline. The framework provided by the lifetime logic proved flexible enough to handle functions that are correct for subtle reasons, like `Ref::map` and `RefMut::map`, part of `RefCell`, which had to have their signature changed from the initial design to ensure soundness [Sapin 2015]. In fact, our verification work resulted in uncovering and fixing a bug in Rust’s standard library [Jung 2017], demonstrating that our model of Rust is realistic enough to be useful. Furthermore, our type system already handles features that are still being sketched for Rust itself, like non-lexical lifetimes [Matsakis 2016b], and we are in active discussion with the Rust community on these topics.

In ongoing and future work, we plan to fill some of the gaps mentioned above and to bring  $\lambda_{\text{Rust}}$  closer to MIR, the most important intermediate language in the Rust compiler. Concretely, we would like to make the fact that all local variables are heap-allocated more implicit, and to extend paths to include dereferencing a pointer. That should permit us to reduce the number of primitive instructions, making each of them correspond to exactly one construct in MIR.

## ACKNOWLEDGMENTS

We wish to thank the Rust community in general, and Aaron Turon and Niko Matsakis in particular, for their feedback and countless helpful discussions.

This research was supported in part by a European Research Council (ERC) Consolidator Grant for the project “RustBelt”, funded under the European Union’s Horizon 2020 Framework Programme (grant agreement no. 683289).

<sup>12</sup>Notice that the `Vec` type, providing dynamically resizable arrays, is supported – though we have not implemented it. The size of `Vec` itself is statically known, the unknown part is behind a pointer indirection, so this is fine. Rust also considers “unsized” types like `[i32]` (an array of integers) independent of any pointer indirection; those types are not currently supported by our model.



## REFERENCES

- David Abrahams. 1998. Exception-safety in generic components. In *International Seminar on Generic Programming (LNCS)*. [https://doi.org/10.1007/3-540-39953-4\\_6](https://doi.org/10.1007/3-540-39953-4_6)
- Amal Ahmed, Andrew W. Appel, Christopher D. Richards, Kedar N. Swadi, Gang Tan, and Daniel C. Wang. 2010. Semantic foundations for typed assembly languages. *TOPLAS* 32, 3 (2010), 1–67.
- Amal Ahmed, Matthew Fluet, and Greg Morrisett. 2005. A step-indexed model of substructural state. In *ICFP*. <https://doi.org/10.1145/1086365.1086376>
- Amal Ahmed, Matthew Fluet, and Greg Morrisett. 2007.  $L^3$ : A linear language with locations. *Fundamenta Informaticae* 77, 4 (2007), 397–449.
- Amal Jamil Ahmed. 2004. *Semantics of types for mutable state*. Ph.D. Dissertation. Princeton University.
- Sidney Amani, Alex Hixon, Zilin Chen, Christine Rizkallah, Peter Chubb, Liam O’Connor, Joel Beeren, Yutaka Nagashima, Japheth Lim, Thomas Sewell, Joseph Tuong, Gabriele Keller, Toby Murray, Gerwin Klein, and Gernot Heiser. 2016. Cogent: Verifying high-assurance file system implementations. In *ASPLOS*. <https://doi.org/10.1145/2872362.2872404>
- Andrew W. Appel. 2001. Foundational proof-carrying code. In *LICS*.
- Andrew W. Appel. 2007. *Compiling with Continuations*. Cambridge University Press.
- Andrew W. Appel (Ed.). 2014. *Program Logics for Certified Compilers*. Cambridge University Press.
- Andrew W. Appel and David McAllester. 2001. An indexed model of recursive types for foundational proof-carrying code. *TOPLAS* 5 (2001). <https://doi.org/10.1145/504709.504712>
- Andrew W. Appel, Paul-André Melliès, Christopher Richards, and Jérôme Vouillon. 2007. A very modal model of a modern, major, general type system. In *POPL*. <https://doi.org/10.1145/1190216.1190235>
- Thibaut Balabonski, François Pottier, and Jonathan Protzenko. 2016. The design and formalization of Mezzo, a permission-based programming language. *TOPLAS* 38, 4 (2016). <https://doi.org/10.1145/2837022>
- Ariel Ben-Yehuda. 2015a. Can mutate in match-arm using a closure. Rust issue #27282. <https://github.com/rust-lang/rust/issues/27282>.
- Ariel Ben-Yehuda. 2015b. dropck can be bypassed via a trait object method. Rust issue #26656. <https://github.com/rust-lang/rust/issues/26656>.
- Ariel Ben-Yehuda. 2015c. std::thread::JoinGuard (and scoped) are unsound because of reference cycles. Rust issue #24292. <https://github.com/rust-lang/rust/issues/24292>.
- Christophe Biocca. 2017. std::vec::IntoIter::as\_mut\_slice borrows &self, returns &mut of contents. Rust issue #39465. <https://github.com/rust-lang/rust/issues/39465>.
- John Boyland. 2003. Checking interference with fractional permissions. In *SAS (LNCS)*. [https://doi.org/10.1007/3-540-44898-5\\_4](https://doi.org/10.1007/3-540-44898-5_4)
- Arthur Charguéraud and François Pottier. 2017. Temporary read-only permissions for separation logic. In *ESOP (LNCS)*. [https://doi.org/10.1007/978-3-662-54434-1\\_10](https://doi.org/10.1007/978-3-662-54434-1_10)
- David G. Clarke, John M. Potter, and James Noble. 1998. Ownership types for flexible alias protection. In *OOPSLA*. <https://doi.org/10.1145/286936.286947>
- The Coq team. 2017. The Coq proof assistant. <https://coq.inria.fr/>.
- Robert DeLine and Manuel Fähndrich. 2001. Enforcing high-level protocols in low-level software. In *PLDI*. <https://doi.org/10.1145/381694.378811>
- Thomas Dinsdale-Young, Lars Birkedal, Philippa Gardner, Matthew J. Parkinson, and Hongseok Yang. 2013. Views: Compositional reasoning for concurrent programs. In *POPL*. <https://doi.org/10.1145/2429069.2429104>
- Thomas Dinsdale-Young, Mike Dodds, Philippa Gardner, Matthew J. Parkinson, and Viktor Vafeiadis. 2010. Concurrent abstract predicates. In *ECOOP (LNCS)*. [https://doi.org/10.1007/978-3-642-14107-2\\_24](https://doi.org/10.1007/978-3-642-14107-2_24)
- Mike Dodds, Xinyu Feng, Matthew J. Parkinson, and Viktor Vafeiadis. 2009. Deny-guarantee reasoning. In *ESOP (LNCS)*. [https://doi.org/10.1007/978-3-642-00590-9\\_26](https://doi.org/10.1007/978-3-642-00590-9_26)
- Derek Dreyer, Amal Ahmed, and Lars Birkedal. 2011. Logical step-indexed logical relations. *LMCS* 7, 2:16 (2011). [https://doi.org/10.2168/LMCS-7\(2:16\)2011](https://doi.org/10.2168/LMCS-7(2:16)2011)
- Derek Dreyer, Georg Neis, Andreas Rossberg, and Lars Birkedal. 2010. A relational modal logic for higher-order stateful ADTs. In *POPL*. <https://doi.org/10.1145/1706299.1706323>
- Manuel Fähndrich and Robert DeLine. 2002. Adoption and Focus: Practical Linear Types for Imperative Programming. In *PLDI*. <https://doi.org/10.1145/512529.512532>
- Matthew Fluet, Greg Morrisett, and Amal J. Ahmed. 2006. Linear regions are all you need. In *ESOP (LNCS)*. [https://doi.org/10.1007/11693024\\_2](https://doi.org/10.1007/11693024_2)
- Douglas Gregor and Sibylle Schupp. 2003. Making the usage of STL safe. In *IFIP TC2 / WG2.1 Working Conference on Generic Programming*. [https://doi.org/10.1007/978-0-387-35672-3\\_7](https://doi.org/10.1007/978-0-387-35672-3_7)
- Dan Grossman, J. Gregory Morrisett, Trevor Jim, Michael W. Hicks, Yanling Wang, and James Cheney. 2002. Region-Based Memory Management in Cyclone. In *PLDI*. <https://doi.org/10.1145/512529.512563>

- ISO Working Group 21. 2011. Programming languages – C++.
- Jonas Braband Jensen and Lars Birkedal. 2012. Fictional separation logic. In *ESOP (LNCS)*. [https://doi.org/10.1007/978-3-642-28869-2\\_19](https://doi.org/10.1007/978-3-642-28869-2_19)
- Trevor Jim, J. Gregory Morrisett, Dan Grossman, Michael W. Hicks, James Cheney, and Yanling Wang. 2002. Cyclone: A safe dialect of C. In *USENIX ATC*.
- Ralf Jung. 2017. MutexGuard<Cell<i32>> must not be Sync. Rust issue #41622. <https://github.com/rust-lang/rust/issues/41622>.
- Ralf Jung, Jacques-Henri Jourdan, Robbert Krebbers, and Derek Dreyer. 2017a. RustBelt: Securing the foundations of the Rust programming language – Technical appendix and Coq development. <https://plv.mpi-sws.org/rustbelt/pop118/>.
- Ralf Jung, Robbert Krebbers, Lars Birkedal, and Derek Dreyer. 2016. Higher-order ghost state. In *ICFP*. <https://doi.org/10.1145/2951913.2951943>
- Ralf Jung, Robbert Krebbers, Jacques-Henri Jourdan, Aleš Bizjak, Lars Birkedal, and Derek Dreyer. 2017b. Iris from the ground up: A modular foundation for higher-order concurrent separation logic. (2017). Conditionally accepted to Journal of Functional Programming.
- Ralf Jung, David Swasey, Filip Sieczkowski, Kasper Svendsen, Aaron Turon, Lars Birkedal, and Derek Dreyer. 2015. Iris: Monoids and invariants as an orthogonal basis for concurrent reasoning. In *POPL*. <https://doi.org/10.1145/2676726.2676980>
- Jan-Oliver Kaiser, Hoang-Hai Dang, Derek Dreyer, Ori Lahav, and Viktor Vafeiadis. 2017. Strong logic for weak memory: Reasoning about release-acquire consistency in Iris. In *ECOOP*.
- Robbert Krebbers. 2015. *The C standard formalized in Coq*. Ph.D. Dissertation. Radboud University.
- Robbert Krebbers, Ralf Jung, Aleš Bizjak, Jacques-Henri Jourdan, Derek Dreyer, and Lars Birkedal. 2017a. The essence of higher-order concurrent separation logic. In *ESOP (LNCS)*. [https://doi.org/10.1007/978-3-662-54434-1\\_26](https://doi.org/10.1007/978-3-662-54434-1_26)
- Robbert Krebbers, Amin Timany, and Lars Birkedal. 2017b. Interactive proofs in higher-order concurrent separation logic. In *POPL*. <https://doi.org/10.1145/3009837.3009855>
- Neelakantan R. Krishnaswami, Aaron Turon, Derek Dreyer, and Deepak Garg. 2012. Superficially substructural types. In *ICFP*. <https://doi.org/10.1145/2364527.2364536>
- Morten Krogh-Jespersen, Kasper Svendsen, and Lars Birkedal. 2017. A relational model of types-and-effects in higher-order concurrent separation logic. In *POPL*. <https://doi.org/10.1145/3009837.3009877>
- Juneyoung Lee, Yoonseung Kim, Youngju Song, Chung-Kil Hur, Sanjoy Das, David Majnemer, John Regehr, and Nuno P. Lopes. 2017. Taming undefined behavior in LLVM. In *PLDI*. <https://doi.org/10.1145/3062341.3062343>
- Xavier Leroy, Andrew Appel, Sandrine Blazy, and Gordon Stewart. 2012. *The CompCert memory model, version 2*. Technical Report RR-7987. Inria.
- Nicholas D. Matsakis. 2016a. Introducing MIR. <https://blog.rust-lang.org/2016/04/19/MIR.html>.
- Nicholas D. Matsakis. 2016b. Non-lexical lifetimes: Introduction. <http://smallcultfollowing.com/babysteps/blog/2016/04/27/non-lexical-lifetimes-introduction/>.
- Nicholas D. Matsakis and Felix S. Klock II. 2014. The Rust language. In *SIGAda Ada Letters*, Vol. 34. <https://doi.org/10.1145/2663171.2663188>
- Robin Milner. 1978. A theory of type polymorphism in programming. *J. Comput. System Sci.* 17, 3 (1978), 348–375. [https://doi.org/10.1016/0022-0000\(78\)90014-4](https://doi.org/10.1016/0022-0000(78)90014-4)
- Aleksandar Nanevski, Ruy Ley-Wild, Ilya Sergey, and Germán Andrés Delbianco. 2014. Communicating state transition systems for fine-grained concurrent resources. In *ESOP (LNCS)*. [https://doi.org/10.1007/978-3-642-54833-8\\_16](https://doi.org/10.1007/978-3-642-54833-8_16)
- Aleksandar Nanevski, Greg Morrisett, Avraham Shinnar, Paul Govereau, and Lars Birkedal. 2008. Ynot: Dependent types for imperative programs. In *ICFP*. <https://doi.org/10.1145/1411204.1411237>
- Liam O'Connor, Zilin Chen, Christine Rizkallah, Sidney Amani, Japheth Lim, Toby Murray, Yutaka Nagashima, Thomas Sewell, and Gerwin Klein. 2016. Refinement through restraint: Bringing down the cost of verification. In *ICFP*. <https://doi.org/10.1145/2951913.2951940>
- Peter W. O'Hearn. 2007. Resources, concurrency, and local reasoning. *Theoretical computer science* 375, 1-3 (2007). <https://doi.org/10.1016/j.tcs.2006.12.035>
- Gordon Plotkin. 1973. Lambda-definability and logical relations. Unpublished manuscript.
- Eric Reed. 2015. *Patina: A formalization of the Rust programming language*. Master's thesis. University of Washington.
- John C. Reynolds. 2002. Separation logic: A logic for shared mutable data structures. In *LICS*. <https://doi.org/10.1109/LICS.2002.1029817>
- The Rust team. 2016. Drop Check. In *The Rustonomicon*. <https://doc.rust-lang.org/nomicon/dropck.html>.
- The Rust team. 2017. The Rust programming language. <http://rust-lang.org/>.
- Simon Sapin. 2015. Add std::cell::Ref::map and friends. Rust PR #25747. <https://github.com/rust-lang/rust/pull/25747#issuecomment-105175235>.
- Josh Stone and Nicholas D. Matsakis. 2017. The Rayon library. <https://crates.io/crates/rayon>.
- Bjarne Stroustrup. 2012. Foundations of C++. In *ESOP*. [https://doi.org/10.1007/978-3-642-28869-2\\_1](https://doi.org/10.1007/978-3-642-28869-2_1)

- Kasper Svendsen and Lars Birkedal. 2014. Impredicative concurrent abstract predicates. In *ESOP (LNCS)*. [https://doi.org/10.1007/978-3-642-54833-8\\_9](https://doi.org/10.1007/978-3-642-54833-8_9)
- Nikhil Swamy, Cătălin Hrițcu, Chantal Keller, Aseem Rastogi, Antoine Delignat-Lavaud, Simon Forest, Karthikeyan Bhargavan, Cédric Fournet, Pierre-Yves Strub, Markulf Kohlweiss, Jean-Karim Zinzindohoue, and Santiago Zanella-Béguelin. 2016. Dependent types and multi-monadic effects in F\*. In *POPL*. <https://doi.org/10.1145/2837614.2837655>
- David Swasey, Deepak Garg, and Derek Dreyer. 2017. Robust and compositional verification of object capability patterns. In *OOPSLA (PACMPL)*. <https://doi.org/10.1145/3133913>
- W. W. Tait. 1967. Intensional interpretations of functionals of finite type I. *Journal of Symbolic Logic* 32, 2 (1967). <https://doi.org/10.2307/2271658>
- Joseph Tassarotti, Ralf Jung, and Robert Harper. 2017. A higher-order logic for concurrent termination-preserving refinement. In *ESOP (LNCS)*. [https://doi.org/10.1007/978-3-662-54434-1\\_10](https://doi.org/10.1007/978-3-662-54434-1_10)
- Amin Timany, Léo Stefanescu, Morten Krogh-Jespersen, and Lars Birkedal. 2018. A Logical Relation for Monadic Encapsulation of State: Proving contextual equivalences in the presence of runST. Accepted to POPL.
- John Toman, Stuart Pernsteiner, and Emina Torlak. 2015. CRUST: A bounded verifier for Rust. In *ASE*. <https://doi.org/10.1109/ASE.2015.77>
- Jesse A. Tov and Riccardo Pucella. 2011. Practical affine types. In *POPL*. <https://doi.org/10.1145/1926385.1926436>
- Aaron Turon. 2015a. Abstraction without overhead: Traits in Rust. <https://blog.rust-lang.org/2015/05/11/traits.html>.
- Aaron Turon. 2015b. Implied bounds on nested references + variance = soundness hole. Rust issue #25860. <https://github.com/rust-lang/rust/issues/25860>.
- Aaron Turon, Derek Dreyer, and Lars Birkedal. 2013. Unifying refinement and Hoare-style reasoning in a logic for higher-order concurrency. In *ICFP*. <https://doi.org/10.1145/2500365.2500600>
- Philip Wadler. 1990. Linear types can change the world! In *Programming Concepts and Methods*.
- Andrew K Wright and Matthias Felleisen. 1994. A syntactic approach to type soundness. *Information and computation* 115, 1 (1994). <https://doi.org/10.1006/inco.1994.1093>