



Leveraging Rust Types for Modular Specification and Verification

VYTAUTAS ASTRAUSKAS, ETH Zurich, Switzerland

PETER MÜLLER, ETH Zurich, Switzerland

FEDERICO POLI, ETH Zurich, Switzerland

ALEXANDER J. SUMMERS, ETH Zurich, Switzerland

Rust's type system ensures memory safety: well-typed Rust programs are guaranteed to not exhibit problems such as dangling pointers, data races, and unexpected side effects through aliased references. Ensuring correctness properties beyond memory safety, for instance, the guaranteed absence of assertion failures or more-general functional correctness, requires static program verification. For traditional system programming languages, formal verification is notoriously difficult and requires complex specifications and logics to reason about pointers, aliasing, and side effects on mutable state. This complexity is a major obstacle to the more-widespread verification of system software.

In this paper, we present a novel verification technique that leverages Rust's type system to greatly simplify the specification and verification of system software written in Rust. We analyse information from the Rust compiler and synthesise a corresponding *core proof* for the program in a flavour of separation logic tailored to automation. To verify correctness properties beyond memory safety, users can annotate Rust programs with specifications at the abstraction level of Rust expressions; our technique weaves them into the core proof to verify modularly whether these specifications hold. Crucially, our proofs are constructed and checked automatically without exposing the underlying formal logic, allowing users to work exclusively at the level of abstraction of the programming language. As such, our work enables a new kind of verification tool, with the potential to impact a wide audience and allow the Rust community to benefit from state-of-the-art verification techniques. We have implemented our techniques for a subset of Rust; our evaluation on several thousand functions from widely-used Rust crates demonstrates its effectiveness.

CCS Concepts: • **General and reference** → **Verification**; • **Theory of computation** → **Programming logic**; **Separation logic**; **Program specifications**; **Program verification**; • **Software and its engineering** → **Software verification**; **Formal software verification**.

Additional Key Words and Phrases: Rust, type systems, heap-manipulating programs, concurrency

ACM Reference Format:

Vytautas Astrauskas, Peter Müller, Federico Poli, and Alexander J. Summers. 2019. Leveraging Rust Types for Modular Specification and Verification. *Proc. ACM Program. Lang.* 3, OOPSLA, Article 147 (October 2019), 30 pages. <https://doi.org/10.1145/3360573>

Authors' addresses: Vytautas Astrauskas, Department of Computer Science, ETH Zurich, Switzerland, vytautas.astrauskas@inf.ethz.ch; Peter Müller, Department of Computer Science, ETH Zurich, Switzerland, peter.mueller@inf.ethz.ch; Federico Poli, Department of Computer Science, ETH Zurich, Switzerland, federico.poli@inf.ethz.ch; Alexander J. Summers, Department of Computer Science, ETH Zurich, Switzerland, alexander.summers@inf.ethz.ch.



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

© 2019 Copyright held by the owner/author(s).

2475-1421/2019/10-ART147

<https://doi.org/10.1145/3360573>

1 INTRODUCTION

Producing reliable system software is challenging. Pointer manipulation, mutable heap data, and concurrency are typically employed to achieve high performance, but cause subtle bugs that are notoriously difficult to uncover and reproduce.

The Rust programming language addresses this problem by preventing some errors statically through its type system, which associates an *exclusive capability* [Boyland et al. 2001] with each mutable memory location. At any time, each exclusive capability is held by at most one executing function: only that code may access the memory location. When aliasing is desired, these exclusive capabilities can be exchanged for *shared capabilities*, with which many references can read a location, but none can modify it. Rust’s type system enforces this discipline, ensuring that well-typed Rust programs are guaranteed to not exhibit data races, have dangling pointers or unexpected side effects through aliased references.

Going beyond memory safety, to guarantee absence of assertion failures, or to prove functional correctness, requires static program verification. Despite recent successes [Bhargavan et al. 2017; Hawblitzel et al. 2015, 2014; Klein et al. 2009], formal verification of system software is notoriously difficult. Reasoning about pointers, aliasing, mutable state, and concurrency requires complex program logics, often based on separation logic [O’Hearn 2004; Reynolds 2002], dynamic frames [Kassios 2011; Leino 2010], or object ownership [Cohen et al. 2009; Leino and Müller 2004]. The expressive power of such logics comes at a price: they describe program behaviours via a rich language of custom assertions (e.g. the points-to predicates, separating conjunction, and magic wands of separation logic). Users are forced to understand these logics to write specifications and direct the construction of a suitable proof. Furthermore, these logics typically require a substantial initial specification effort, even to prove simple properties such as crash-freedom or absence of overflow. Consequently, the application of these logics remains the domain of expert researchers, forming a major obstacle to the more-widespread verification of system software.

In this paper, we present a novel verification technique that leverages Rust’s type system to greatly simplify the specification and verification of Rust programs. Our key insight is to combine the rich capability information implicit in Rust’s type system with user-provided assertions that express functional behaviour to automatically construct a proof in an expressive program logic. We analyse information from the Rust compiler and synthesise a *core proof* of memory safety of the program in a flavour of separation logic that facilitates the integration of functional correctness properties. Crucially, our proofs are constructed and checked *automatically*; users of our technique never work with the underlying formal logic. They can add specifications at the abstraction level of Rust expressions; our technique interweaves these specifications automatically into the core proof to verify them modularly. Consequently, our technique shields users completely from the complexity of the underlying logic; assertions and error messages are expressed at the level of Rust expressions, which makes our technique accessible to programmers.

Contributions. The main contributions of our work are:

- (1) We define a specification language for expressing functional properties of Rust programs, suitable for modular verification. Our language is based on Rust expressions, and does not expose the complexity of the underlying verification logic.
- (2) We propose *pledges*: a novel specification construct that enables modular specification and verification of Rust functions that yield borrowed references.
- (3) We define a verification technique that encodes both capability information and user-provided assertions into the implicit dynamic frames logic [Smans et al. 2009], a close relative of separation logic [Parkinson and Summers 2012].

- (4) We automate our verification technique by constructing a translation from the Rust program and specifications into the Viper intermediate verification language [Müller et al. 2016]. Our translation generates correct specifications and, crucially, synthesises all necessary auxiliary annotations needed for proof checking to be completely automatic.
- (5) We provide an implementation of our technique as a plugin for the Rust compiler. We used our implementation to automatically construct core proofs for several thousand unannotated Rust functions, and to verify a range of stronger properties (via our specification language) for selected Rust programs. Our tool is available as an artefact [Astrauskas et al. 2019a].

Unsafe Code and RustBelt. Rust’s type system enforces strong rules, but provides an escape hatch for when these are too restrictive: code blocks and functions can be declared *unsafe*, weakening the compiler’s checks, and correspondingly risking the guarantees they provide. Unsafe code should be encapsulated by libraries such that client code cannot observe its usage [Rust contributors 2019b]. The ongoing RustBelt project [Jung et al. 2018] is aimed at defining formal semantic foundations for making this requirement precise and verifiable. Our work has fundamentally different (and complementary) aims and technical contributions. We do not address unsafe code in this paper, but present a verification technique enabling user-specifications at a high level of abstraction and automatic proofs; RustBelt verification entails directly using an advanced separation logic based on Iris [Jung et al. 2015], for which proofs are interactive (in Coq [Coq Team 2014]), and constructed by experts (see also Sec. 8).

Outline. The rest of this paper is organised as follows. We illustrate our approach on an example in Sec. 2. Sec. 3 presents our specification and verification technique for Rust without references; Secs. 4 and 5 extend our technique to handle mutable and shared references, respectively. Sec. 6 introduces pledges, used to attach functional specifications to functions returning mutable borrows. We describe and evaluate our implementation in Sec. 7, discuss related work in Sec. 8, and conclude in Sec. 9. App. A includes a detailed illustration of our encoding on a simple example.

2 MOTIVATING EXAMPLE

In this section, we illustrate the basics of our approach from a programmer’s perspective. Details are explained in subsequent sections.

Example. Fig. 1 shows a simple Rust program which declares a struct `Point` with two integer fields, and two functions. The function `shift_x`, shifts the x-coordinate of a given `Point` instance. Rust types express capabilities to access memory. Here, the type `&mut Point` expresses that `p` is a mutable reference, also called a *mutable borrow*. When the function is called, the capabilities to access the fields of the passed `Point` instance are *temporarily* transferred from the caller to the callee function, and back when the function terminates. Since the borrow is mutable, `shift_x` is allowed to modify the instance, here, by assigning to its `x` field.

Function `align` takes the capabilities for a mutable pair of *boxed* points. A value of type `Box<T>` represents a pointer (with capabilities) to a value of type `T`; this indirection allows the `Points` to be passed by reference (instead of by copying them). The selectors `.0` and `.1` select elements of the parameter pair.

The `assert!` statement performs a runtime check that its parameter expression evaluates to true. Evaluating `*(segm.1).x` here is allowed although `segm.1` was borrowed on the previous line (`&mut segm.1` creates a mutable reference to `segm.1`), as the compiler infers that the borrow is no longer used after the call to `shift_x`. Therefore, the borrow *expires* after the call, restoring capabilities to the borrowed-from `segm.1`.

```

1 struct Point {
2   x: i32, y: i32
3 }
4
5 #[ensures="p.x == old(p.x) + s"]
6 #[ensures="p.y == old(p.y)"]
7 fn shift_x(p: &mut Point, s: i32) {
8   p.x = p.x + s
9 }
10 fn align(
11   mut segm: (Box<Point>, Box<Point>)
12 ) -> (Box<Point>, Box<Point>)
13 {
14   let diff = (*segm.0).x - (*segm.1).x;
15   shift_x(&mut segm.1, diff);
16   assert!((*segm.0).x == (*segm.1).x);
17   segm
18 }

```

Fig. 1. Points in Rust. Proving modularly that the assertion on line 16 holds requires properties guaranteed by the Rust ownership system as well as a user-provided specification for function `shift_x`. Note that the compiler can infer dereferences of boxed values and so `(*segm.0).x` could be simplified to `segm.0.x`; we make dereferences explicit for clarity.

Correctness Arguments. Consider what is needed to prove that the **assert!** statement can *never* fail at runtime. Showing this property for *all* calls to `align` requires the following properties, in which p_0 and p_1 denote the `Point` instances passed into function `align` as `segm.0` and `segm.1`:

- (1) The call to `shift_x` increases the value of $p_1.x$ by the value of `diff`.
- (2) The call does not modify $p_0.x$. Therefore, right after the call, we have $p_0.x = p_1.x$.
- (3) The call to `shift_x` does not modify the tuple `segm`, that is, we still have $p_0 = \text{segm.0}$ and $p_1 = \text{segm.1}$ and, therefore, $(*\text{segm.0}).x = (*\text{segm.1}).x$.
- (4) The code is data race free and, thus, the values of all memory locations are stable throughout the execution.

Except for property 1, all of these properties are guaranteed by Rust’s type system. In particular, `segm.0` and `segm.1` are guaranteed to reference *different* `Point` instances p_0 and p_1 because the type `Box<Point>` guarantees unique ownership of the boxed value. If `segm.0` and `segm.1` were aliases, `segm.0` (respectively `segm.1`) would not be the only pointer pointing to p_0 (p_1), and the unique ownership property would be violated; hence `segm.0` and `segm.1` cannot be aliases. Since only the capabilities for p_1 are transferred to function `shift_x`, all fields of p_0 are guaranteed to be left unchanged by the call (property 2, and analogously for property 3). Preserving information about mutable state, so-called *framing*, is one of the main difficulties of modular verification [Kassios 2011; Leino and Nelson 2002; Müller 2002; Reynolds 2002]; by leveraging information from Rust’s type system, our technique solves the frame problem without imposing substantial overhead on programmers.

Property 4 is a consequence of the fact that Rust’s type system requires an exclusive capability in order to mutate a memory location. It allows one to verify the assertion without reasoning about thread interleavings [Jones 1983; Owicki and Gries 1976] or explicit proofs of race freedom [O’Hearn 2004], which would increase the specification effort for programmers.

Property 1 follows from the functional behaviour of function `shift_x`, which is expressed as a user-provided postcondition, written as a Rust annotation. Our specification language is based on Rust boolean expressions, extended with few (but powerful) additional constructs; here, the `old` construct [Leavens et al. 2011] is used to refer to the pre-state value of a mutable memory location, which allows one to express relational properties between the pre- and post-state of a call. The second postcondition of `shift_x` is not needed to verify the assertion, but would likely be required by other client code. Since `shift_x` takes a mutable borrow to the `Point` instance `p`, the type system allows it to modify any fields of `p`. The second postcondition tightens framing by guaranteeing that

`p.y` will remain unchanged. Note that our specifications are as simple as traditional contracts [Meyer 1992], but enable the sound verification of concurrent, heap-manipulating programs.

The assertion could in principle be proved without the postconditions, by inlining the implementation of `shift_x`. However, verifying a call against a specification, instead of an implementation, makes verification *modular*, which is important for scalability, to provide guarantees for library code, and to reduce and localise the re-verification effort when parts of a codebase are changed.

Verification. Our example illustrates that correctness proofs for Rust programs need to combine information about capabilities for memory locations (aliasing, side effects, framing, data race freedom) with information about their values. While the former is provided by the type system, the latter must be supplied as assertions (the *inference* of value information is possible, but orthogonal). To formally integrate both sources of information, our technique encodes capabilities and user-provided functional properties into a program logic that is sufficiently expressive to capture both and reason about their interactions.

Our verification technique encodes the capability information and statement semantics of Rust programs into a formal logic, resulting in a *core proof* that captures information about aliasing and side effects that is essential for verification, in particular, for framing. It is crucial that this encoding, as well as the *checking* of our core proofs, is completely automatic; any required user interaction would expose the complexity of the underlying program logic to the programmer and, thereby, break the abstraction that our work aims to provide. For program logics (such as separation logics) expressive enough to model Rust's type capabilities, this degree of automation is beyond the state of the art.

Our core proofs provide the technical foundations for verifying stronger properties, such as the correctness of user-provided assertions, as well as absence of arithmetic overflows and various kinds of exceptions (called *panics* in Rust), including `assert!` failures. Constructing the core proof is challenging, especially handling complex forms of (re)borrowing and synthesising auxiliary annotations to automate the proof search, as we will explain later.

A major virtue of our approach is that it lowers the barrier to applying verification. The construction of the core proof from a well-typed Rust program is fully automatic, so that programmers can immediately focus on verifying the main properties of interest, such as the validity of a given assertion. They can control the required effort by writing simpler or more comprehensive specifications; if the (optional) checks for built-in properties such as overflow are not enabled, the minimal specification is none at all. This is in stark contrast to most existing verification techniques, which require a substantial initial effort to set up predicates, invariants, or ghost state and to verify memory safety, before programmers can turn to the properties they care about most.

3 RUST'S CAPABILITIES FOR VERIFICATION

In this section, we explain our specification and verification technique for Rust code without borrowing, which we defer until Secs. 4 and 5. We present the capability information that is needed to construct a core proof, explain how we encode this proof to Viper, and then show how to incorporate user-provided assertions.

We present our work for a small but technically-challenging subset of safe Rust. It includes primitive types (`bool`, integers, `char`) and the following compound types: boxes (for heap-allocated data), tuples, structs, enumerations, and generic type parameters. In addition, we support mutable and shared references to those types. The code is organised into functions, methods, and trait methods; we will use *functions* to refer to all of these in this paper. Functions can use generics with trait bounds and lifetime parameters without constraints. Deterministic, side-effect free functions can be marked as *pure*, which allows using them in specifications. Function bodies may contain

```

1 #[ensures="old((*p).x + s) == (*result).x"]
2 #[ensures="old((*p).y) == (*result).y"]
3 fn shift_x(p: Box<Point>, s:i32) -> Box<Point> {
4     box Point { x: (*p).x + s, y: (*p).y }
5 }
6
7 fn align(mut segm: (Box<Point>, Box<Point>)) -> (Box<Point>, Box<Point>) {
8     let mut end = segm.1; // move assignment
9     // segm.1 is now inaccessible
10    let diff = (*segm.0).x - (*end).x;
11    end = shift_x(end, diff);
12    segm.1 = end;
13    // end is now inaccessible
14    assert!((*segm.0).x == (*segm.1).x);
15    segm
16 }

```

Fig. 2. A variation of the example from Fig. 1 that uses move assignments instead of borrowing. The move assignment in line 8 removes the capability for `segm.1` until it is restored in line 12, which prevents accesses in between. In particular, omitting line 12 would cause a compiler error, since it would not be possible to assemble the full capabilities required by the return type.

branching constructs, loops, move and copy assignments, boolean and integer operations, function calls, type constructor invocations, and casts. Functional behaviour can be specified by function pre- and postconditions as well as loop invariants. Commonly-used Rust features which fall outside of our currently-supported subset include: closures, lifetime parameters to struct types, the recently introduced two-phase-borrows [Rust contributors 2019c], as well as unsafe code; these are left for future work. Note that the language supported by our prototype implementation is more restricted; the details are provided in Sec. 7.

3.1 Ownership and Capabilities

The Rust type system enforces a strict discipline governing not only which values can be stored in which locations, but also which *places* (Rust’s terminology for expressions denoting memory locations [Rust community 2018b]) can be used to access those values at each program point.

Ownership. In Rust, every value stored in memory has a unique *owner*, which is a variable (variables always include function parameters) in a currently-active function execution. Ownership is transitive: the owner of a struct value is also the owner of its fields. The scope of a value’s owner implicitly determines the deallocation time of the owned memory. Rust’s type rules guarantee that by the time the owner goes out of scope (or if the owning variable is reassigned), no place will have the capability to access the underlying memory, preventing dangling pointers. Rust types typically convey ownership of the corresponding memory; for instance, `Box<T>` is the type of an *owning* pointer to a location of type `T`.

Fig. 2 shows a variation of the example from Fig. 1 without borrowing (which we will explain shortly). The assignment in line 8 is a *move assignment*, which transfers ownership from `segm.1` to `end`, making `segm.1` unreadable. Similarly, the call to `shift_x` transfers ownership from `end` to parameter `p`. `end` becomes usable again (and owns its contents) on its reassignment when the function terminates, and analogously for `segm.1` in line 12. The subsequent assertion holds for reasons similar to those outlined in the previous section for Fig. 1. In particular, ownership

guarantees that the two points are distinct objects, and is used in our technique to provide framing information for the call to `shift_x`.

Capabilities. Owning a memory location does not necessarily provide the right to access it. For instance, function `align` in Fig. 1 owns both points throughout its execution, but the right to access the point in `segm.1` is temporarily transferred to `shift_x` using a borrow¹. Borrowing affects who may access a location, but not who owns it. To distinguish these concepts, we use the term *place capability* (or *capability* for short) to denote the right to access the value stored in a place.

Precise knowledge of the capabilities at any given program point is crucial for verification, especially framing. For instance, function `align` in Fig. 1 may frame the value of `(*segm.0).x` around the call to `shift_x` because it retains the corresponding capability, whereas `(*segm.1).x` may change because the capability is transferred for the call. Note that Rust source types do not provide complete capability information: for instance, throughout a function body, struct-typed variables retain the same Rust type, but capabilities to their fields vary as they are borrowed or moved. To make this information explicit, we defined an algorithm to compute precise summaries of the capabilities held at each program point, which we call *place capability sets*.

In the following, we define the type of *results* our algorithm computes, but omit the algorithm itself for brevity. At verification boundaries such as function pre- and post-states, we extract these results directly from the Rust compiler: we use the declared types of all definitely-assigned variables in scope at these program points to compute a suitable summary of the capabilities held. However, to elaborate this to an automatable formal proof we need explicit information about how these capabilities evolve at each *intermediate* program point. This is information which, in principle, is internally computed by the Rust type checker, but using representations which are not exposed; our algorithm therefore recovers these intermediate steps to produce a detailed account of the capabilities at every program point.

Definition 3.1 (Place Capability Sets). Places, ranged over by p , are expressions defined by the following grammar: $p ::= x \mid p.f \mid (*p)$. For a place p of the form $p'.f$ and $(*p')$, place p' is called a *sub-place* of p ; this notion is extended transitively in the natural way. A *place capability set* (PCS) is a finite set of places.

The initial PCS for a function contains exactly the capabilities for its parameters, e.g. `segm` for function `align` in Fig. 2. Every subsequent statement may require certain capabilities to be in the PCS and then transform the PCS. For instance, the move assignment on line 8 requires the PCS before the assignment to contain `segm.1`, and transforms the PCS from $\{\text{segm.0}, \text{segm.1}\}$ to $\{\text{segm.0}, \text{end}\}$, reflecting the move of capabilities. These PCS transformations are defined for each primitive statement; we provide details in App. B of our technical report [Astrauskas et al. 2019b].

The evolution of capabilities during Rust type checking can require additional operations on top of the requirements for individual Rust statements. For example, the PCS before the assignment on line 8 is obtained from the initial PCS of the function by exploiting transitivity of capabilities. Since the capabilities of a place also imply capabilities for all its sub-places, the type checker can *unpack* the capability for `segm` into $\{\text{segm.0}, \text{segm.1}\}$. Unpacking is one of several PCS operations the type checker may perform to manipulate PCSs:

Definition 3.2 (PCS Operations). A *PCS operation* is a *remove*, *unpack*, or *pack* of a capability in a PCS. Remove is defined as the corresponding set operation.

Let p be a place of struct type, and let f_1, \dots, f_n be the fields of the struct. For a PCS S such that $p \in S$, the *unpacking* of p in S is the PCS $(S \setminus \{p\}) \cup \{p.f_1, \dots, p.f_n\}$. If p is instead of box type, the *unpacking* is $(S \setminus \{p\}) \cup \{(*p)\}$.

¹In this and the next section, we will discuss only exclusive capabilities; we extend our work to shared capabilities in Sec. 5.

The *packing* of p in S is the inverse operation. It is defined only when the $p.f_i$ (or $*p$) are in S .

The Rust type checker implicitly employs these operations between statements to show that the capabilities required by the next statement are present. `Unpack` is used to enable operations on individual fields of structs (e.g. the `move` on line 8), and `pack` is used when the entire struct is passed as argument or result to check that capabilities for all sub-places can be reassembled, e.g. at the end of `align` in Fig. 2. All three operations can be used at join points in the control flow if the joined paths provide different capabilities: `remove` is needed to drop the capabilities available in one path but not the other (e.g. due to moving out a struct field in a branch), while `pack` and `unpack` are needed to unify the PCSs of the joining paths.

By combining type information extracted from the compiler and our own analysis, we infer automatically, for each statement, the PCS before the statement, a sequence of *necessary* PCS operations to be applied before the statement, and the PCS after the statement, describing the actual flow of capabilities implied by Rust’s type rules. This information is vital for the construction of the core proof, as we explain in the next subsection. Since it provides a detailed account of *why* a Rust program type-checks, we believe that it could also be repurposed as the basis of other analysis, verification, and visualisation tools.

3.2 Constructing the Core Proof

We verify Rust programs by encoding the program, capability information, and user-provided specifications into the intermediate verification language Viper [Müller et al. 2016], and using Viper’s existing verification tools. Viper provides a simple heap-based imperative language, along with a number of reasoning primitives for expressing verification problems; each Viper method is equipped with a precondition and a postcondition; Viper loops are equipped with loop invariants. For each function in the Rust program, we generate a corresponding method in our Viper program, such that successful verification of the Viper method implies correctness of the Rust original.

Viper Resources. Viper’s heap is object-based: heap locations are identified by a pair of a **Ref**-typed value and a field name. Viper’s type system is simple: the built-in **Ref** type is the only type for objects in the Viper heap, and all fields declared in a Viper program are in principle available in all objects. Akin to separation logic, Viper enforces that a field location can be accessed only when *permission* is held to do so. Conversely, so long as the permission to a field location is held, Viper assumes that its value cannot change, which provides framing. Viper field permissions are tracked in the program state as *affine resources*; they can be explicitly added or removed from a program state, or implicitly dropped if not required.

Viper’s logic is based on implicit dynamic frames [Smans et al. 2009], a close relative of separation logic [Parkinson and Summers 2012], but with the important facility to incorporate heap-dependent expressions in logical assertions (including calls to side-effect-free functions) [Smans et al. 2010]. Assertions called *accessibility predicates*, written $\mathbf{acc}(e.f)$ are used to denote the *exclusive field permission* for the field f of the object denoted by e . Viper’s conjunction $\&\&$ acts *multiplicatively* (in the sense of linear logic [Girard 1987]); analogously to separating conjunction in separation logic, it requires the *sum* of the necessary resources in each conjunct. For example, the assertion $\mathbf{acc}(x.f) \ \&\& \ \mathbf{acc}(y.f)$ denotes *two* exclusive permissions, which implies that x and y cannot alias each other. In addition to accessibility predicates, our work makes crucial use of two other kinds of Viper resource assertions adopted from separation logic: predicates and magic wands, which will be explained later.

Modelling Memory. We model Rust’s program states in Viper by mapping every Rust memory location to a corresponding Viper field location. We model any non-primitive type in Rust as a

Viper object (of **Ref** type); each transitive element of the Rust type (struct fields, tuple elements, box contents, reference targets) is modelled as a field of the Viper object. Since Rust references make it possible to take the address at which any value is stored, we also model Rust primitive types with an additional indirection; any primitive type is modelled as a Viper object with a single field that contains the actual value.

As an example, the parameter `segm` in Fig. 2 is modelled as a Viper **Ref** with fields `elt0: Ref`, `elt1: Ref` for the two elements of the tuple; in turn, each of these values is a box, modelled as an object with a single field `val_ref: Ref`. Similarly to pairs, Point struct values are Viper objects with two fields, while their individual `i32` fields (also addressable in Rust) are modelled as objects with a single field `val_i32: Int`. Here, we use Viper’s built-in **Int** type for unbounded integers; if overflow-checking is enabled, we encode bounds using additional assertions.

Modelling Rust Types. As explained above, the core proof requires precise capability information, for instance, to enable sound framing. To provide this information, we model the capabilities represented by Rust types as resource assertions in Viper. Since place capabilities can have unboundedly many sub-places (struct types may recurse via e.g. box types), we cannot enumerate these explicitly. Instead, we translate each Rust type into an instance of a Viper *predicate*. Predicates are a standard means of defining parameterised, possibly recursive assertions [Parkinson and Bierman 2005]; predicate *instances* are tracked as affine resources in Viper.

We define a Viper predicate per Rust type in the source program; each predicate is parameterised by a single **Ref**-typed parameter (the Viper object representing the Rust place): for primitive types, the body contains an accessibility predicate for the single field storing the value, while for structures and tuples, it consists of the conjunction of accessibility predicates for each field, as well as a predicate instance for the translation of the field’s type. As a simple example, we translate a place capability of type `i32` using the following `i32` predicate (the bounds properties within the predicate body are omitted if overflow-checking is disabled):

```
field val_i32 : Int

predicate i32(self: Ref) {
  acc(self.val_i32) && i32MIN <= self.val_i32 && self.val_i32 <= i32MAX
}
```

For polymorphic types such as `Box<T>` where the type parameter `T` is known, we monomorphise, generating a specialised predicate, e.g. for `Box<Point>`, where `Point` is the predicate generated for the Rust `Point` type:

```
predicate BoxPoint(self: Ref) {
  acc(self.val_ref) && Point(self.val_ref)
}
```

When the type parameter of a polymorphic type is not known, e.g. when encoding `Box<T>` in a generic function under a parameter `T`, we encode the type parameter as an abstract predicate, whose body is unspecified. For example, the encoding of `Box<T>` becomes:

```
predicate T(self: Ref);

predicate BoxT(self: Ref) {
  acc(self.val_ref) && T(self.val_ref)
}
```

The predicate instance $T(\mathbf{self.val_ref})$ represents exclusive capability to the boxed value; by making the predicate abstract, we ensure that it conveys no further information about this value.

For enumeration types (a form of tagged unions) we model the discriminant (tag) of the enumeration as an integer field with bounded values. This discriminant is then used on the left-hand-side of implications to guard which accessibility predicates for the variant's fields are actually included in the predicate. An example is provided for the `Option<T>` type in App. A of our technical report [Astrauskas et al. 2019b].

Modelling Place Capabilities. As explained in the previous subsection, Rust types prescribe the available capabilities at function boundaries, but the PCS may change throughout the function. Using the predicates defined above, we can directly translate a place capability set into a corresponding Viper assertion: each element of the PCS gives rise to an instance of the Viper predicate corresponding to its type. We call the Viper assertion consisting of the conjunction of these predicate instances the *Viper embedding of the PCS*.

This embedding allows us to map each Rust function to a corresponding Viper method, along with corresponding preconditions and postconditions at the Viper level. The precondition is the Viper embedding of the PCS representing all input parameters to the Rust function; the postcondition is the Viper embedding of the PCS representing the output parameters. We use the precondition to prescribe the initial state for verifying the Viper method, while the postcondition is checked at the end of the method body. Analogously, the Viper embedding of the PCS at each loop head provides a loop invariant that lets Viper verify loops for our core proof automatically.

As an example, when we generate a Viper method for the Rust function `align` in Fig. 2, the precondition will be `PairBoxPoint(seg)`, where `PairBoxPoint` is a Viper predicate whose body contains permission to the pair's fields and a `BoxPoint` predicate instance for each. The postcondition will be `PairBoxPoint(result)`, where `result` refers to the method's return value.

Many program verifiers, including Viper, prevent indefinite unrolling of recursive predicates by treating predicates *isorecursively* [Abadi and Fiore 1996; Crary et al. 1999; Summers and Drossopoulou 2013]: exchanging a predicate instance for its body is not done automatically, but requires explicit operations in the program, called **fold** and **unfold** in Viper. These statements are needed exactly at those program points at which the Rust type checker performs the packing and unpacking PCS operations. By exploiting the PCS information we summarise for the Rust program, we are therefore able to instrument our generated Viper program with exactly the necessary additional annotations required for Viper to be able to reason about these predicates fully automatically, without any user interaction. Recall that full automation is essential to preserve the abstraction provided by our work and shield programmers from the complexity of the underlying logic employed in Viper.

Modelling Capability Transfer. In order to model Rust function calls, we need to correctly reflect the transfer of place capabilities. Having generated Viper methods, along with appropriate preconditions and postconditions, we want to model a call by *removing* the Viper resources in the precondition, and subsequently *adding* the Viper resources in the postcondition. Viper provides **inhale** and **exhale** statements for such explicit manipulation of resources [Müller et al. 2016]. A statement **exhale** *A* has the effect of checking that the assertion *A* is true in the current state, and removing all resources it requires. Moreover, when permission to a memory location is removed, Viper removes any knowledge about the value stored in the location to reflect that the value could be changed by another function. Dually, **inhale** *A* adds the resources prescribed by *A*.

Using these Viper statements, we encode a Rust-level function call as an **exhale** of the precondition (reflecting that the corresponding capabilities become unavailable to the caller), followed by an **inhale** of the postcondition (reflecting those which are returned). Via the capability information

extracted from the Rust compiler (in PCS form) and our Viper embedding of this information, this handling of Viper resources corresponds precisely to the transfer of capabilities in Rust. More details are illustrated in App. A.

3.3 Functional Specifications

The core proof we have constructed so far, by itself, does not go beyond the guarantees provided by the type system. However, it provides the foundation for verifying stronger properties such as functional correctness. In particular, it provides precise aliasing information and framing, which is essential for sound and modular reasoning about the Rust heap. Due to the design of our modelling of Rust types, and choice of underlying logic, extending our core proof to properties beyond memory safety is surprisingly simple.

We enable optional checking of generic properties such as absence of overflows and absence of panics (e.g. assertion failures) simply, by generating additional assertions in the Viper program. For example, to prove absence of assertion failures, it suffices to insert an `assert false` statement into the branch of the code that raises a panic when the runtime-check fails, to verify that this branch is unreachable.

User-provided assertions such as function pre and postconditions are translated and conjoined to the corresponding assertions of the core proof. This simple treatment is enabled by our choice of implicit dynamic frames for the underlying logic: unlike standard separation logics, implicit dynamic frames separates resource properties from value properties, as in `acc(x.f) && x.f > 0`. Similarly, predicate instances can be combined with applications of heap-dependent mathematical functions to constrain the resources in the predicate. We use this feature to allow user-provided specifications to include calls to side-effect-free Rust functions, similarly to JML's pure methods [Leavens et al. 2011], which is useful to express properties of unbounded data structures and to make use of abstractions already provided in the Rust program.

The technique presented so far supports the specification and verification of programs using only move and copy assignments. The treatment of borrowing is more intricate, both for the core proof and for functional specifications, as explained in the next sections.

4 MUTABLE BORROWS

One of the most important and intricate features of Rust's type system is *borrowing*: creation of references that temporarily take capabilities, but do not change ownership of the referenced value. In this section, we extend the construction of the core proof to *mutable* borrows; shared (immutable) borrows are covered in the next section. The specification and verification of functional properties in the presence of borrows will be discussed in Sec. 6.

4.1 Borrows and Lifetimes

Fig. 3 shows an example built upon the Point example from Fig. 1. Function `shift_nth_x` borrows a route from its caller; that is, the capability for the parameter `r` is transferred to the function. Each borrow has a *lifetime*, computed by the Rust compiler, representing an extent in the program execution for which the borrow needs to remain live; a lifetime always includes at least all program points where the borrow is used. Note that lifetimes are typically not explicit in the program text, but chosen implicitly by the compiler. At the end of a borrow's lifetime, the borrow is said to *expire*, and the capabilities associated with it are restored to the borrowed-from place. Since `r` in `shift_nth_x` is a function argument, its lifetime spans the entire function body.

Reborrowing. It is possible to *reborrow* either the full place or a sub-place of an existing borrow. The call to `nth_point` at line 54 of Fig. 3 implicitly reborrows `r` and transfers its capability to

```

1 // List of Points
2 struct Route {
3     current: Point,
4     rest: Option<Box<Route>>
5 }
6
7 #[pure]
8 #[ensures="result > 0"]
9 fn length(r: &Route) -> i32 {
10     1 + match r.rest {
11         Some(box ref q) => length(q),
12         None => 0
13     }
14 }
15
16 #[pure]
17 #[requires="0 <= n &&
18           n < length(r)"]
19 fn nth_x(r: &Route, n: i32) -> i32 {
20     if n == 0 { r.current.x } else {
21         match r.rest {
22             Some(box ref q) =>
23                 nth_x(q, n-1),
24             None => unreachable!()
25         }
26     }
27 }
28
29 #[requires="0 <= n && n < length(r)"]
30 #[ensures="result.x ==
31           old(nth_x(r, n))"]
32 #[ensures="???"] // See Sec. 6
33 fn nth_point(r:&mut Route, n: i32) ->
34     &mut Point {
35     if n == 0 { &mut r.current } else {
36         match r.rest {
37             Some(box ref mut q) =>
38                 nth_point(q, n-1),
39             None => unreachable!()
40         }
41     }
42 }
43
44 #[requires="0 <= n && n < length(r)"]
45 #[ensures="length(r) ==
46           old(length(r))"]
47 #[ensures="nth_x(r, n) ==
48           old(nth_x(r, n)) + s"]
49 #[ensures="forall i: i32 ::
50           (0<=i && i<length(r) && i != n) ==>
51           nth_x(r, i) == old(nth_x(r, i))"]
52 fn shift_nth_x(r: &mut Route,
53               n: i32, s:i32) {
54     let p = nth_point(r, n);
55     shift_x(p, s);
56 }

```

Fig. 3. An implementation of routes (sequences of points from Fig. 1), illustrating borrows. Function `shift_nth_x` borrows a route from its caller. This reference is reborrowed in the call to `nth_point`, which returns a reference to a point in the route. Both borrows expire after the call to `shift_x` on line 55. Functions annotated with `[pure]` are side-effect-free and can be used in specifications. The missing `???` specification will be explained in Sec. 6.

that function. More interestingly, function `nth_point` creates a borrow to a sub-place of the route, namely its n^{th} point, and returns it to its caller; this reborrow's lifetime persists beyond the call in which it is made. Reflecting this possibility, after the call to `nth_point`, `r` is blocked from being used until `p` expires, since `p`'s capabilities could (and indeed are, in this example) be for a part of the same memory that `r` had capabilities to access; if `r` were usable, this would violate exclusivity of these capabilities.

Reborrowing extends the lifetime of the borrowed-from reference: the original borrow cannot expire until all (transitive) reborrows are known to have expired. In our example, the lifetime of the reborrow created for the call to `nth_point` is extended until the further reborrow `p` expires after the call to `shift_x`².

²The exact rules depend on the version of Rust; in Rust 1.0, only entire explicit scopes (with a few exceptions) can be used as lifetimes. The recently introduced *non-lexical lifetimes* [Rust community 2017] are more fine-grained; we support both.

Borrow Information. As we have explained in the previous section, constructing the core proof for a Rust program requires precise capability information, which we have so far represented via place capability sets and place capability operations at each program point (see Sec. 3.1). This information is insufficient for programs with borrows; in particular, it does not explain how, when borrows expire, the capabilities (and corresponding permissions in our Viper encoding) are restored to where they were borrowed from. For this we need precise information about which borrows are active for which lifetimes, and which reborrow each other.

We obtain information about the lifetimes selected by the Rust compiler from the latest borrow checker implementation [Rust contributors 2019a] and additional compiler analyses. In some cases, we also need to fill in missing information; in general, the Rust compiler tracks *negative* information (sufficient to check whether an error should be raised), but does not always store explicit *positive* information witnessing why type-checking succeeds. We represent the extracted information as follows: we assign identities to each borrow operation in a function, every move assignment of a reference (which we treat as a further reborrow), and every function call which returns a borrow. In terms of these identifiers, we record the set of identifiers of borrows which are alive before each statement. Moreover, we extract a *reborrow relation*: a binary relation on borrow identifiers, indicating which borrows *may* directly reborrow from which.

4.2 Encoding Borrows as Resource Assertions

In this subsection, we explain how we encode the capabilities associated with a mutable borrow as well as those for the remainder of the place from which it was borrowed. The next subsection discusses how the core proof manipulates these capabilities when borrows are created or expire.

The place capabilities associated with a mutable reference are encoded analogously to those for a box type (cf. Sec. 3.2). We define a Viper predicate to represent each reference type used in the program. For instance, predicate `RefMutPoint` for mutable references to points includes an instance of predicate `Point` for the referenced point, stored with an extra indirection through `val_ref`:

```
predicate RefMutPoint(self: Ref) {
  acc(self.val_ref) && Point(self.val_ref)
}
```

When a Rust function returns a borrow, this must always be a reborrow (possibly transitively) of a borrow passed as one of the function’s parameters. Rust does not allow returning borrows to memory owned within the function (for instance, a local variable) because the returned borrow would become a dangling pointer when the function returns and the owned memory is deallocated. We refer to functions which take a mutable borrow as parameter and return a borrow as *lenders*.

Intuitively, when a lender function such as `nth_point` creates and returns a borrow then it takes the capability for a parameter place (here, `r`) and splits it into two parts: the capability for the borrow *returned* by the function and any *remainder* of the original capabilities. Rust does not provide a way of representing such types with missing capabilities, nor are they used at function boundaries in Rust’s type checking; instead, the borrowed-from place (including this remainder) is simply unusable until the returned borrow expires. However, to track information about the *values* stored in the borrowed-from data structure, we need to represent this remainder formally in our core proof and, thus, need a suitable formal model for these remainder capabilities.

Our key insight here is that the separation logic magic wand connective [O’Hearn et al. 2001] lets us formally model the remainder capabilities resulting from a reborrow; it can express *partial* permissions to data structures, such as the `Route` with one `Point` missing. A magic wand assertion $A \multimap B$ represents a resource which can be *combined* with the resource A , and $A \multimap B$ and A together then exchanged for the resource B ; this is called *applying* the magic wand. For our purposes, we

use A to represent the resources that are given up by the expiring borrow, and B to represent those of the borrowed-from place; the magic wand thus represents the remainder. In particular, assertions A and B each encode Rust types for which we already have translations. Magic wands are also supported by Viper [Schwerhoff and Summers 2015].

For lender functions, we generate Viper postconditions to be a conjunction of: (1) the Viper embedding of the PCS for the places returned by the function, (2) the translation of any user postconditions regarding these places, and (3) a magic wand $A \multimap B$, where A is the same assertion as (1), and B is the Viper embedding of the capabilities for the parameter from which the returned reference was borrowed. For example, for the `nth_point` function of Fig. 3, we generate³:

```
method nth_point(r: Ref, n: Ref) returns (res: Ref)
  requires RefMutRoute(r) && i32(n) &&
           0 <= n.val_i32 && n.val_i32 < length(r.val_ref)
  ensures RefMutPoint(res) &&
           res.val_ref.x.val_i32 == old(nth_x(r.val_ref, n.val_i32)) &&
           (RefMutPoint(res)  $\multimap$  RefMutRoute(r))
```

where `RefMutPoint` and `RefMutRoute` are the predicates generated for mutable references to structs `Point` and `Route`, respectively. The magic wand in the example represents both the partial capability for `r` and the promise that this partial capability can be combined with the capability currently associated with `res` to obtain those originally associated with `r`; by applying the magic wand (at call site), we make use of this promise to restore full capabilities for `r`. As we will show in Sec. 6, the ability to connect the capabilities returned on expiry with the capabilities of the borrowed-from data structure is also essential for adding functional specification to lender functions.

4.3 Automating Proofs with Borrows

Viper supports the magic wand connective, but requires annotations in order to reason about it [Schwerhoff and Summers 2015]. We explain here how we generate these annotations using our recorded information from Sec. 4.1.

Restoring Capabilities. In our core proof, we generate operations to formally explain how capabilities are restored when borrows expire. Intuitively, we use the recorded borrow information to undo the borrows in an order *opposite* to that in which they were created, using our extracted reborrow relation. Starting from the PCS describing the capabilities just before the borrows expire, we perform the following steps for each borrow. (1) We synthesise any necessary pack/unpack operations to obtain the place capability for the borrower (for instance, the left-hand side of a borrowing assignment); these are encoded in Viper as **fold/unfold** operations as explained in Sec. 3.2. (2) We replace this capability with the place capability from which it was borrowed (for instance, the right-hand side of a borrowing assignment). For a direct assignment, this is a no-op in Viper (which already knows the equality of the two locations); for reborrows returned from lender functions, the replacement is encoded by *applying* the corresponding magic wand, directed in Viper via an explicit **apply** statement.

Consider for example the program point after the call to `shift_x` at line 55 in Fig. 3. At this point, the borrow `p` expires. Based on the information we record, we know that `p` was returned from the lender function `nth_point`, blocking the function's parameter `r` from being usable, so we apply the wand `RefMutPoint(p) \multimap RefMutRoute(r)`, which was returned by the Viper encoding of the function call.

³Viper also requires us to unfold predicates around expressions which require permissions from their bodies; our work generates these annotations too, but we elide them here for readability.

Creating Reborrows. When verifying the definition of a function returning a borrow (such as `nth_point` above), our core proof needs to create the required magic wand, which is done in Viper using a **package** statement. This Viper statement must be annotated with a proof of *how*, given *any state* satisfying the wand’s left-hand side, one will be able to reassemble the wand’s right-hand-side. As a side-effect, the **package** statement consumes any additional resources needed to obtain this right-hand-side (these reflect the remainder capabilities discussed in Sec. 4.2). We generate these proofs automatically, in an analogous way to the explanation of expiring borrows in the previous paragraph. The annotations required to automate our proofs in Viper can be elaborate, but we demonstrate in Sec. 7 that we are consistently able to generate them fully automatically.

We now have the machinery in place to construct a core proof for Rust code that may include mutable borrows and reborrows. We use similar techniques to handle reborrows inside loops (for example, when a Rust reference is used to traverse a recursive heap data structure) in order to generate the required loop invariants at the Viper level completely automatically. In the next section we will explain how to extend our core proof to shared borrows, and in Sec. 6 how to specify and verify functional properties concerning reborrows, e.g. to add specification to lender functions such as the missing postcondition of function `shift_nth_x` in Fig. 3.

5 SHARED BORROWS

Mutable borrows provide temporary exclusive access to a place, but prevent multiple usable aliases. In contrast, Rust’s *shared references* (or *shared borrows*) permit multiple references to exist to the same place, or to a place and its sub-places, simultaneously. To ensure the absence of data races and unexpected side effects via aliasing, Rust’s type system enforces that the shared parts of the data structure are immutable while at least one such shared reference exists. In this section, we extend the verification technique presented so far to support shared references.

5.1 Read and Write Capabilities

To distinguish between exclusive, mutable access and shared, immutable access, we refine the capabilities associated with a place. We associate *shared capability* with places that store shared borrows, or which are currently borrowed from by a shared borrow, and use *exclusive capability* for all other places (exclusive capabilities correspond to the capabilities used in the previous sections). This refinement also affects the PCS and PCS operations from Sec. 3: we refine place capability sets into partial maps $places \rightarrow \{exclusive, shared\}$ to specify which capability is associated with each place in a PCS. In the initial PCS of a function, parameters of a shared-reference type (such as `&Point`) are mapped to a shared capability, and parameters of other types to an exclusive capability.

Pack and unpack operations (cf. Def. 3.2) are extended accordingly: unpacking assigns to all the sub-places the capability of the unpacked place, while packing requires all the sub-places to have the same kind of capability, which is then assigned to the packed place. Recall that PCS operations are applied between statements to reorganise capabilities. In order to go between exclusive and shared capability, we employ two additional PCS operations, called *downgrade* and *upgrade*. These exchange an exclusive place capability for a shared one and vice versa.

When a shared borrow is created by borrowing a place for which exclusive capability was held, this causes a downgrade to shared capability, which is then duplicated for both the borrowed-from place (to make it immutable while the shared borrow exists) and the newly-created borrow. The original exclusive capability of the borrowed-from place is restored (by an upgrade) only when the borrow checker has determined that all shared borrows have expired, and so the place (and its sub-places) can no longer be aliased via shared references.

```

1 // Count the points of `r` inside the rectangle identified by `a` and `b`
2 fn count_inside(r: &Route, a: &Point, b: &Point) -> u32 { /* ... */ }
3
4 // Move the first point until it is unique in `r`
5 fn make_first_unique(r: &mut Route) {
6     let first = &r.current;
7     // r.current.x += 1;
8     let first2 = first;
9     if count_inside(r, first, first2) > 1 {
10        assert!(first.x == r.current.x);
11        r.current.x += 1;
12        make_first_unique(r);
13    }
14 }

```

Fig. 4. An example using shared references. Function `make_first_unique` keeps incrementing the `x`-coordinate of the first point in route `r` until this point's `x`-coordinate is unique among all points in `r`. By calling function `count_inside` with shared references to the same point for parameters `a` and `b`, it yields how often that point occurs in `r`.

Function `make_first_unique` in Fig. 4 illustrates the use of shared references. Line 6 creates a shared reference, which downgrades `r` to be temporarily immutable for as long as a shared borrow exists. The type system disallows modifying `r` in line 7 because `r` has a shared capability; in contrast, the assignment in line 11 is permitted because the last shared borrow for `r` expires at line 10, causing an upgrade to exclusive capability for `r`. In the meantime, the assignment in line 8 creates a second shared borrow, which expires after the call to `count_inside`. Since this function takes only shared references, it cannot modify `r` and, therefore, the assertion in line 10 can be proved to hold.

5.2 Encoding Read Capabilities

Many separation logics support *fractional permissions* [Boyland 2003] to distinguish between read and write access to memory locations. In these logics, a permission can be split into several fractions. A full permission allows (exclusive) write access, whereas any non-zero fraction permits read access. After a full permission has been split into fractions, those fractions can be re-combined to get back the full permission and, thus, write access.

Viper supports fractional permissions as its standard means of expressing read-only access to the heap; we therefore use these to model shared capabilities. However, in order to accurately reflect Rust's type-checking, and fully exploit the information we can extract from the Rust compiler, we use fractional permissions in a non-standard way. Constructing a standard proof in a fractional-permission logic would require elaborate specifications to keep track of the fractional amounts of permission associated with each shared borrow (which changes as new borrows are created) and to describe how fractional permissions are transferred between different function executions. In particular, we would need to precisely reassemble these fractional amounts to justify restoring write access to a Viper heap location. However, the Rust compiler does not use such accounting to justify write access; instead, it restores an exclusive capability once it can prove that all involved shared references have expired.

With this in mind, we construct our core proof relying on Rust's borrow checker to indicate when a first shared borrow for a place is created (and a downgrade is performed to make the

borrowed-from place temporarily immutable) and when the last shared borrow expires (and a corresponding upgrade is performed, so that the place becomes mutable again). In particular, our core proof does not need to add up fractions in order to reassemble a full permission and, thus, the precise fractional amounts associated with shared borrows are irrelevant, as long as we can distinguish between read and write access.

Our encoding uses full permissions for write access (as in the previous sections) and a so-called symbolic read permission [Heule et al. 2013] for read access. A *symbolic read permission* uses a fractional amount that is unspecified, but constrained to satisfy the following properties: (1) it is greater than zero and, thus, permits read access, and (2) the sum of all symbolic amounts for any given resource is less than a full permission. The latter property allows one to create additional symbolic read permissions without the risk of ever obtaining (or exceeding) a full permission.

Our encoding maintains the following invariant for each place p whose type is not a shared-reference type: if p is not borrowed from, there exists a full permission for the associated Viper resource; otherwise, if there exist shared borrows for p then p and each of the shared references is associated with a symbolic read permission. To maintain this invariant, we encode operations on shared borrows as follows. When the capability for p is downgraded, we replace the full permission for the corresponding Viper resource with a symbolic read permission (through a Viper **exhale** operation) and create a symbolic read permission for the new shared reference (through an **inhale** operation). Conversely, when the last shared borrow expires (and a downgrade is performed), we remove its read permission and restore the full permission of the borrowed-from place. In between, when additional shared borrows are created, we simply give them an additional symbolic read permission. This forging of symbolic read permissions is sound because, as we explained above, the sum of all symbolic read permissions is constrained to be less than a full permission. Analogously, we forge a new symbolic read permission when a shared reference p is passed to a function call. This encoding allows the caller to retain its read permission to p and to use it for framing, that is, conclude that the function call cannot change the referenced value.

In the example of Fig. 4, function `make_first_unique` starts out with a full permission to `r` because it is a *mutable* reference. The assignment on line 6 creates the first shared borrow for `r` and, thus, replaces `r`'s full permission with a symbolic read permission. Line 8 creates another shared borrow with another read permission. When calling the function in line 9, the caller retains a read permission to `r` and is therefore able to conclude that the call does not affect the equality `first == r.current`, which was established on line 6 and which implies the assertion in line 10. After this assertion, the last shared borrow of `r` expires, which restores full permission and, thus, enables the assignment on line 11 as well as the subsequent recursive call.

6 PLEDGES FOR MUTABLE BORROWS

It is common for Rust functions to be passed a reference (borrow) as parameter, create a further reference from it (via reborrowing) and to return the new reference to their callers. This idiom is for instance used in getters such as function `nth_point` in Fig. 3. Callers of such lender functions require information about the new borrow as well as the borrowed-from place in order to determine properties of the data structure when the borrow expires. For instance, verifying the three postconditions of `shift_nth_x` in Fig. 3 relies on the fact that the call to `nth_point` does not modify the route `r`; if the call, for instance, removed the first point from the route, none of the postconditions would hold.

Since lender functions (such as `nth_point`) may modify the borrowed-from place before the borrow is created, information about the presence or absence of such side effects needs to be conveyed to callers via the function's postcondition. For shared borrows, this is possible because the borrowed-from place remains usable and, thus, may be accessed in the function's postcondition.

```

1 #[ensures="after_expiry<result>(
2   length(r) == old(length(r)) &&
3   nth_x(r, n) == before_expiry(result.x) &&
4   forall i: i32 :: (0<=i && i<length(r) && i != n) ==>
5     nth_x(r, i) == old(nth_x(r, i))
6 )"]
7 fn nth_point(r: &mut Route, n: i32) -> &mut Point {
8   // ...
9 }

```

Fig. 5. An example of our pledge specification feature. The postcondition shown here is the missing third postcondition of `nth_point` from Fig. 3. The first conjunct states that the function does not change the length of the route before creating the borrow, by relating the prestate of a call to `nth_point` to the state in which the `result` borrow expires. Similarly, the third conjunct states that all points other than the n^{th} have unchanged `x`-values. The second conjunct relates the `result` borrow right before the expiry to the rest of the borrowed-from place.

However, as we discussed in Sec. 4, returning a *mutable* borrow renders the borrowed-from parameter unusable until the borrow expires. Consequently, referring to the borrowed-from place in the postcondition would violate Rust’s type rules and have an unclear semantics as it would introduce aliasing among mutable references. In this section, we introduce pledges, a novel specification construct that lets us specify lender functions. For instance, pledges let us express that `nth_point` does not modify the route `r` (despite having the capabilities to do so, according to its signature).

Pledges. Pledges are assertions that can be used in postconditions of lenders to specify properties of borrowed-from places guaranteed to hold *at the time when the borrow expires*, that is, when the borrowed-from place becomes usable again. This design is compliant with the Rust rules as it avoids referring to unusable places; as we will explain towards the end of the section, it is also essential for modular reasoning.

When a lender returns a mutable borrow, this effectively separates the place for the borrowed-from parameter into a part that can still be accessed through the returned borrow and an unusable remainder. For example, for `nth_point` from Fig. 3, these are the returned point `p` and the remainder of the route `r`. In general (though not in our example), a lender could have modified both parts of the borrowed-from parameter before returning. After returning, the unusable remainder is known to be unchanged until the new borrow expires. However, the part that can be accessed through the returned borrow may get modified by client code, after calling the lender. For instance, `shift_nth_x` modifies the borrowed point `p`. In a modular setting, lenders cannot anticipate how the returned borrow will be modified. Therefore, pledges specify their guarantees *parametrically with* the state of the returned borrow when it expires. For example, for `nth_point` we need a way to explain how `r` will look after `p` expires, which partially depends on how the client manipulates the borrowed memory.

Fig. 5 shows the third postcondition for `nth_point` from Fig. 3. It contains a pledge, expressed as argument to the `after_expiry` construct, which is parameterised by the borrow it specifies (here, `result`). The pledge itself expresses three guarantees: (1) the length of the borrowed-from Route `r` will be unchanged since the prestate of the call to `nth_point`, (2) the `x`-coordinate of `r`’s n^{th} Point will be equal to that of the returned borrow at the time the borrow expires, and (3) all other `x`-coordinates of the Route will be unchanged. Note that `nth_point` can guarantee these properties because the remainder of `r` cannot be modified until the borrow expires and because

they hold irrespective of the changes made to `result.x` until then. The `before_expiry` notation lets a pledge refer to the borrow right before it expires; it is needed because the borrow is no longer usable after it expires.

This pledge is sufficiently strong to verify the postcondition of `shift_nth_x` in Fig. 3. It provides strong guarantees about the borrowed-from place without constraining how clients may modify the borrow and, thus, enables modular specifications.

Modularity. Our pledges feature respects Rust’s typing discipline by expressing properties of the borrowed-from place (`r` in our example) only in states in which Rust allows it to be used. One might be tempted to consider a simpler alternative design which *does allow* such violations for the sake of writing standard postconditions: e.g. allowing `nth_point` to express directly that `r` is unchanged using `forall i: i32 :: (0 <= i && i < length(r)) ==> nth_x(r, i) == old(nth_x(r, i))`. Yet, such a postcondition is useless to a caller that subsequently makes modifications via the returned borrow `p`. Since `p` aliases a sub-place of `r`, any such change may affect properties of `r`, for instance, the value of `nth_x(r, i)`. Consequently, we no longer obtain an automatic way to precisely frame such properties. Clients would need to know precisely how to reach `p` from `r` and how the functions `length` and `nth_x` are defined, which would violate information hiding. Moreover, clients would have to prove inductive lemmas to show how changes via `p` affect properties of `r`, losing automation. Our pledges avoid all of these problems.

Encoding. The encoding of our verification technique to Viper extends naturally to pledges; a pledge is translated as an additional conjunct on the right-hand-side of the magic wand that is created when a mutable borrow is created. Recall from Sec. 4.2 that the right-hand-side indeed represents the state as it will look once the borrow expires; pledges let us complement the resources there with properties of the values stored in the borrowed-from place. The borrow together with the pledge from Fig. 5 would be encoded to the following magic wand (we omitted the necessary unfold operations for readability):

```
RefMutPoint(res) --* (RefMutRoute(r) &&
  length(r) == old(length(r.val_ref)) &&
  nth_x(r.val_ref, n.val_i32) == old[lhs](res.val_ref.x.val_i32) &&
  (forall i: Int :: 0 <= i && i < length(r.val_ref) && i != n.val_i32 ==>
    nth_x(r.val_ref, i) == old(nth_x(r.val_ref, i))))
```

When such a magic wand is packaged (*cf.* Sec. 4.3), Viper proves that the claimed pledge will indeed hold for *any* future state of the borrowed memory, which is necessary to soundly account for any possible changes to the borrow. Viper permits expressions `old[lhs](e)` on the right-hand-side of a magic wand, which evaluate to `e`’s value in the state immediately before the wand is applied (e.g. just before `res` expires, in our example); this feature enables us to encode our `before_expiry` construct in a straightforward way.

7 IMPLEMENTATION AND EVALUATION

We have implemented our work as a plugin for the Rust compiler, and evaluated it on a wide variety of crates (Rust packages) from the Rust package repository. The evaluation shows that our technique can generate core proofs fully automatically and verify interesting correctness properties without the need for complicated specifications.

7.1 Implementation

We implemented a tool called Prusti as a Rust compiler plugin, usable with Cargo, the official package manager for Rust. Working with Prusti provides a similar experience to existing tools used by Rust

developers, such as the Rust linter Clippy [Clippy contributors 2019]. Prusti performs its main work after the type checking pass of the Rust compiler. We extract the compiler’s CFG representation (MIR) along with type and borrow-checker information, construct the corresponding Viper program, and verify it with Viper’s symbolic execution verifier; verification results are translated back from Viper to Rust and reported using the Rust compiler’s error reporting mechanisms. In addition to proving user specifications, Prusti optionally checks absence of panics and overflows.

Our current tool works on an expressive subset of safe Rust, including for instance heap-allocated data and **Box** types, shared and mutable borrows, traits, generics, and common use cases of lifetime parameters to functions. We have not yet implemented support for reborrowing inside loops, pure functions returning non-primitive types, and abrupt termination of loop bodies; these restrictions will be lifted in the future (and can typically be worked around by rewriting the program). During the development of our tool, we built a test suite of more than 300 correct and incorrect Rust programs (annotated with expected verification errors) to check that we model corner cases of Rust’s semantics correctly; these are provided with our implementation.

To support libraries, our tool provides a `#[trusted]` annotation, allowing us to equip functions with contracts used by callers but *not* checked against the function’s implementation.

7.2 Evaluation

We evaluate our work in three ways: (1) we evaluate the construction of core proofs on all functions from the top 500 Rust crates that fall within our supported language subset; (2) we evaluate the ability to verify panic-freedom by proving the absence of overflows in examples that check for overflow at runtime, to determine whether these runtime checks may ever fail (without any user-provided specifications); (3) we evaluate the use of user-provided specifications by verifying panic freedom and richer functional correctness properties of existing implementations of well-known algorithms. All timings were performed using a clean Ubuntu 18.04.1 installation, on a desktop with a 4-core (8 hyper-threads) Intel i7-2600K 3.40GHz CPU, 32GB of RAM and an SSD disk.

(1) *Core Proofs.* To test the automation of our core proof construction, we took the 500 most popular Rust community crates [Rust community 2018a], and applied three simple filters: firstly, we discarded any crates (148) which did not compile successfully within 15 minutes using the standard compiler⁴ (without our tool); secondly, we filtered all remaining 56,257 functions (top-level, `impl` and trait functions) with a simple syntactic check for unsupported language features; thirdly, we manually discarded ten unusually large functions that would have taken more than one minute just for the encoding, due to the large number of local variables used (five implement 4×4 matrix operations; the other five contain huge match expressions with up to 2,000 cases). This left us with 11,791 functions (21% of the total) to evaluate our work on. We re-ran the compiler with Prusti on the unmodified source code of these functions to generate and verify core proofs.

The verification of these 11,791 functions succeeded as expected, without any need for manual intervention. This shows that we generate sufficient annotations to automate the core proofs in Viper. These annotations are non-trivial: we generated a total of 1,140,384 lines of Viper code, of which 138,499 are **fold**, **unfold**, **package** or **apply** statements necessary to automate the proofs.

We measured how much time is required by Viper to verify each function, reporting results (averaged across three runs) in Fig. 6 (left). We observe that the average verification time per function is 1.2 seconds, that only 0.16 seconds are enough to verify 50% of the functions, and that almost all of the functions (98.6%) are verified in less than 10 seconds each. A small fraction of the functions take more than 10 seconds to verify, for the same reasons as the ten unusually-large

⁴`rustc nightly-2018-06-27; flags -Zborrowck=mir -Zpolonius -Znll-facts` and using the reference Polonius algorithm (“Naive”).

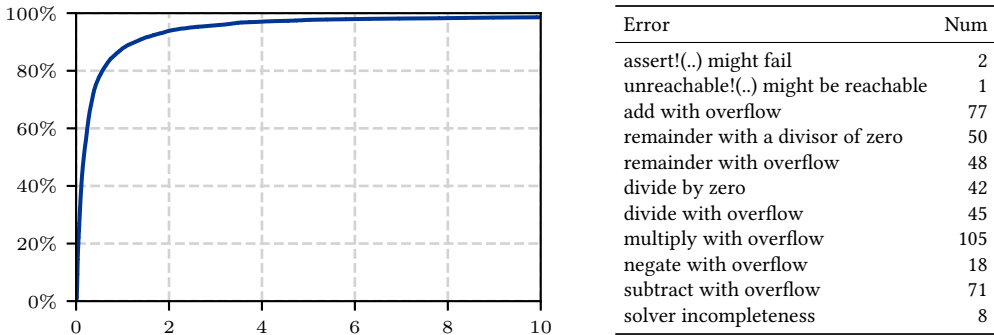


Fig. 6. Left: cumulative distribution of the verification time (horizontal axis, in seconds) required for the core proof verification of each of the 11,791 supported functions (177 functions required between 10 and 120 seconds; 11 required between 120 and 888 seconds). Right: distribution of error messages for the overflow freedom evaluation on 519 functions.

functions we discarded (see above). Since the MIR representation used for the encoding is highly unoptimised and uses significantly more (temporary) local variables than the source program, we can reduce this overhead in the future by enabling simple optimisations that the compiler runs in later stages.

(2) *Overflow Freedom.* We automatically identified all 519 supported functions in our evaluation crates which contain runtime checks for integer overflows or divisions by zero. We re-ran Prusti on these, enabling checks for panics and overflows (again, without specifications). Interestingly, 52 of these functions verified; on manual inspection, this was due to expressions that cannot overflow (e.g. $x-x/4$), or that were guarded by range checks. Since our tool proves soundly that these checks can never fail, one could eliminate them to improve performance without compromising safety.

For each of the remaining 467 functions, Prusti reported a verification error, listed in Fig. 6 (right). Manual inspection showed that these are mostly due to implicit assumptions on argument ranges; our technique makes it possible for developers to make these assumptions explicit as preconditions, and verify them at each call site. In eight cases, Prusti failed to prove that Rust’s dynamic overflow checks actually imply that an operation does not overflow. Our inspection revealed that these verification failures are due to the handling of non-linear arithmetic in the underlying Z3 SMT solver. Increasing Viper’s timeout for each Z3 query from 10 to 60 seconds results in “divide by zero” verification errors in all eight cases, which is the expected result.

(3) *Specifications and Functional Behaviour.* In the third part of the evaluation, we investigated the specification and verification of both absence of panics and richer functional properties, using examples from the programming chrestomathy site Rosetta Code [Rosetta Code contributors 2018], a Rust tutorial on linked lists [Rust community 2019], and from Matsakis’ blog posts on Rust’s language design [Matsakis 2018a,b]. From Rosetta Code, we manually selected a diverse list of eleven examples that either fall into the supported subset of the language or can be adapted without major changes. In order to handle examples using standard library types, we wrote wrappers marked with `#[trusted]` for these types (as explained above); we also rewrote `for` loops as `while` loops, and restructured some code to avoid `return` and `break` statements.

Table 1 gives an overview of the verified examples (we provide the code including specifications in the accompanying artefact). Before any manual modifications, the Rosetta Code examples had between 10 to 89 lines of code (excluding blank lines and comments) and between 1 and 6 functions. The average total verification time (averaged over 3 runs) is typically less than 30 seconds, which we

Table 1. An overview of the examples verified in the third part of the evaluation. The column “LOC” indicates the number of lines in the unmodified example; “#Fns” is the number of functions; “Spec. LOC” is the number of lines used for specification and ghost code; “All Time” indicates the time in seconds required to encode and verify the example; “Viper Time” is just the time needed by the Viper symbolic execution back-end verifier to verify the encoding. “No Panic”/“No Overflow” shows whether we verified absence of panics/overflows (“-” means that the example contains no operations that could panic/unchecked arithmetic). The first two groups of examples are taken from the Rosetta Code website [Rosetta Code contributors 2018], except the “Linked List Stack” example which we took from [Rust community 2019] because it is more complete than the one in Rosetta. The second group differs from the first one in that we verified some functional properties. For example, for the “Ackermann Func.” and “Fibonacci Seq.” examples, we showed that multiple implementations all compute the correct result. We had to monomorphise “Binary Search” and “Selection Sort” for proving stronger functional properties because Prusti does not yet support intrinsic trait properties such as transitivity of the equivalence operator. The preconditions we chose for the Ackermann functions do not prevent overflow and, thus, this aspect could not be verified (indicated by “x”). In “Knapsack Problem/0-1”, we verified correctness of all intermediate computations; correctness of the result and absence of overflow would require sum comprehensions, an advanced specification feature not yet supported in Prusti. The two examples in the last group are from Matsakis’ blog posts about non-lexical lifetimes in Rust [Matsakis 2018a,b]. For one of them, proving panic freedom failed because the program does not handle all IO errors.

Example	LOC	#Fns	Spec. LOC	Time (s)		No Panic	No Overflow	Verified Additional Properties
				All	Viper			
100 doors	19	2	7	10.9	7.4	✓	✓	
Binary Search (generic)	16	1	2	16.2	12.9	✓	✓	
Heapsort	39	3	18	30.6	26.2	✓	✓	
Knight’s tour	89	6	71	127.6	120.2	✓	✓	
Knuth Shuffle	16	2	3	9.5	6.2	✓	✓	
Langton’s Ant	58	4	22	16.7	11.8	✓	✓	
Selection Sort (generic)	20	2	8	19.2	15.2	✓	✓	
Ackermann Func.	16	2	17	7.4	4.4	-	x	Correct result
Binary Search (monomorphic)	16	1	29	25.5	21.4	✓	✓	Correct result
Fibonacci Seq.	46	6	26	9.1	5.7	-	-	Correct result
Knapsack Problem/0-1	27	1	86	139.4	131.6	✓	x	Correct computation
Linked List Stack	59	5	60	21.4	16.9	✓	-	Correct behaviour
Selection Sort (monomorphic)	20	2	34	29.6	24.2	✓	✓	Sorted result
Towers of Hanoi	10	2	5	5.9	3.2	-	✓	Correct param. range
Borrow First	7	1	1	6.6	3.6	✓	✓	
Message	13	1	0	7.2	4.2	x	-	

consider reasonable for our so-far unoptimised encoding and tool. The slowest examples “Knight’s tour” and “Knapsack Problem/0-1” take less than two and a half minutes (each of them contains one large function that takes most of the time). In all cases, standard deviations were around 1 second.

For most examples, we verified the absence of panics and overflows, by adding specifications where necessary. In some cases, for example for “Binary Search”, this required adding only a simple invariant that the indices are no larger than the vector’s length, which allowed the verifier to prove not only the absence of out-of-bounds accesses, but also the absence of overflows. In other cases, for example for “Knight’s tour”, we had to add ghost code to encode object invariants. The most interesting specification for proving panic-freedom is for “Langton’s Ant”, which required not only quantifiers to specify an invariant of the grid on which the ant walks, but also a pledge to specify how changes made via borrows affect the invariant of the grid. Via our evaluation, we found a bug in the source code of this example, which causes an integer overflow during execution. We fixed this error by correcting existing boundary checks and types.

For seven examples, we also verified properties that go beyond basic safety. For two of them, we had to monomorphise the generic parameters to integers in order to use integer comparisons instead of a trait function. Functional correctness of the binary search example initially failed to verify; closer inspection revealed an off-by-one bug in the source code (a fixed version verifies with our tool). We encode other properties such as sortedness (“Selection Sort”), functional correctness of recursive and iterative implementations (“Fibonacci Seq.” and “Ackermann Func.”), functional correctness of a data structure (“Linked List Stack”), correctness of intermediate computations (“Knapsack Problem/0-1”), and validity of parameter values in function calls (“Towers of Hanoi”).

These seven examples require on average 1.3 lines of annotation per line of code. While this overhead is not negligible, it is lower than the overhead required by existing verifiers for heap-manipulating programs. Moreover, our annotations are conceptually much simpler since they are expressed in terms of Rust expressions rather than complex program logics. Another core advantage of our approach is that the user is not forced to provide all of them from the beginning, but can add them gradually to strengthen the verified properties. For instance, proving safety for “Binary Search” requires only two lines of annotations. To additionally prove that the returned index is correct if `Some` is returned, the user needs to add two additional straightforward assertions. Finally, proving correctness for the case that `None` is returned is slightly more involved because it requires writing a quantifier that expresses that the vector is sorted. Nevertheless, none of these assertions expose the complexity of program logics for concurrent, heap-manipulating programs.

We also evaluated our tool on two examples from Matsakis’s blog [Matsakis 2018a,b], designed to illustrate difficult borrowing patterns. The support for the first example was added to stable Rust only recently, while the second one still requires a nightly-build version of Rust. Both examples are already supported by our tool (using the corresponding new borrow checker implementation).

8 RELATED WORK

Capability-Based Type Systems. Many other type systems can also be understood to associate capabilities with reference types [Boyland et al. 2001]. Some extend pre-existing languages (e.g. Sing# [Barnett et al. 2011], C# [Gordon et al. 2012] and Scala actors [Haller and Odersky 2010]); more recently, several programming languages have built these in (e.g. Pony [Clebsch et al. 2015], AEminium [Stork et al. 2014], and Rust itself [Matsakis and Klock II 2014]). Such built-in type systems are exploitable by the compiler: e.g. for memory management in Rust, or to enable the distributed garbage collection in Pony. While these systems provide programmers with stronger guarantees than traditional type safety, *functional correctness* of programs cannot be expressed: our work shows how to layer such verification concerns on top, while exploiting the benefits provided by the type system.

Type Systems for Verification. Liquid Types [Rondon et al. 2008] equip types with logical qualifiers prescribing value properties; their extension to Alias Refinement Types [Bakst and Jhala 2016] applies to mutable heap data structures. Type checking is decidable, and loop invariants can be inferred. Unlike our work (*cf.* Secs. 4 and 6), there is no support for references (reborrows) which persist beyond the function calls or loops they are created in.

SYMPAR [Bierhoff 2011] targets formal verification for Java, employing a notion of permissions to separate reasoning about aliases from verification conditions concerning values. Like our work, user-specification is at the level of the programming language. A planned addition to SPARK (a subset of Ada designed for formal verification) will add pointer support [Maalej et al. 2018], using a type system similar to Rust. In both systems, returning reborrowed references is not supported.

Rust Verification Tools. CRUST [Toman et al. 2015], a recent adaptation of SMACK [Baranowski et al. 2018], and Lindner et al. [2018]’s work provide *bounded* verification tools for Rust (including

unsafe code); these tools allow user checks to be added as Rust expressions. These tools work on C/LLVM code where Rust's type information is absent. By contrast, we exploit this information for modular *unbounded* (sound) verification, and support richer functional specifications via old expressions and pledges.

Ullrich [Ullrich 2016] encodes safe Rust programs into functional programs, to be interactively verified in Lean [de Moura et al. 2015]. Reborrows are supported via lenses [Foster et al. 2005]. Recent work at Galois similarly reduces reasoning about a subset of safe Rust to proofs about functional programs in Saw [Dockins et al. 2016]. In contrast to these works, our technique does not require the manual construction of proofs or verifier directives; in addition, our underlying separation logic formalism will provide a suitable (imperative-style) model for an extension to unsafe code in the future.

As a general point, we believe our implementation to be the first verification technology so far to operate directly on the Rust compiler's analysis results and representations of source programs; there is no gap between the Rust programs and notions and the starting point for our work.

Rust Semantics and Formalisations. A number of formalisms for subsets of Rust have been designed, focusing on type soundness results [Kan et al. 2018; Reed 2015; Wang et al. 2018; Weiss et al. 2018]. It would be interesting to compare these formal models with the PCS/borrow summaries that our work produces from the compiler.

RustBelt [Jung et al. 2018] provides a formalisation aimed at proving that unsafe library implementations encapsulate their unsafe behaviour, and defining formally what this notion should mean for Rust. As explained in the introduction, the goals of our work are very different, and this led to different technical choices and contributions. Whereas RustBelt aims to formalise metatheory for the language and to discharge proofs using the Coq proof assistant, we do not address Rust semantics, and aim primarily at the functional (automatic) verification of safe Rust code, and to equip programmers with specification mechanisms at the level of abstraction of such code, shielding them from the complexity of formal logics.

RustBelt handles the expiry of borrowed references with a combination of a *lifetime logic* for determining when lifetimes can be known to safely end, and rules which restore full capabilities to a borrowed-from place once the corresponding lifetime is over. Restoring capabilities using this indirection via the lifetime logic has the advantage that this solution works consistently for mutable and for shared borrows, and for both safe and unsafe code (in the latter case, more work is required in the lifetime logic itself). Our handling of shared borrows (and the corresponding upgrade to full permissions) was inspired by this approach, but we rely on Rust's borrow checker to determine when borrows expire.

One key difference in our technical approach is that our logical encoding of borrows (using magic wands) reflects the flow of capabilities from the re-borrowed reference to the place borrowed from. It is this modelling which enables our pledges feature to be layered on top, since we can relate the two states of the accessible memory before and after the borrow's expiry in one assertion. In this way, we can directly express how changes made via a borrow affect the borrowed data; this is a fundamental difference in the two models. It would be interesting future work to integrate the two approaches, using RustBelt's lifetime logic to justify expiry of references in unsafe code in place of the (currently trusted) borrow checker we use for safe code.

There are also other technical differences that were motivated by the differing goals of the projects. Since RustBelt focuses on a formal program model usable even in the presence of unsafe Rust, and on Coq-based proof, their logical foundation is a powerful and complex separation logic based on Iris, and their language formulation is a continuation-passing-style intermediate representation, convenient to work with in Coq; translation from Rust to this representation is

performed manually by experts. For our goals, it is essential that the translation to the language in which the proof is performed is automated, and that the input specifications written by users match the features and abstraction level of Rust's source code.

9 CONCLUSIONS

We have presented a new specification and verification technique for Rust, leveraging the guarantees provided by the language's type system, and synthesising from these automatic core proofs in a logic akin to separation logic. By providing specifications at the level of abstraction of the Rust language, programmers can extend this core proof to verify rich functional properties. Verification is performed via an automatic translation to the Viper infrastructure. A key virtue of our technique is that it does not expose the complexity of the underlying verification logic; programmers work exclusively on the level of Rust programs, which facilitates adoption. Our work is implemented and freely available. Our evaluation shows that we can reliably automatically construct core proofs for real-world Rust code and verify functional correctness properties.

Our main goal for future work is to extend the subset of Rust supported by our technique and tool, in particular, to add iterators, closures, and structures with lifetime parameters; these extensions will broaden the applicability of our technique to a wider range of real-world code. We also plan to add support for certain classes of *unsafe* code.

A ENCODING TO VIPER

At a high level, Prusti encodes each Rust function to a Viper method, each Rust *pure* function to a Viper function, and each Rust type to a Viper predicate. In this section we will show step-by-step how the function `force_inc` provided in Fig. 7 is encoded to Viper. Some technical details are simplified, because the compiler internally translates the source program to MIR: a CFG-based representation, much more verbose than the original program. An explanation of the encoding of Rust *pure* functions and types is available in App. A of our technical report [Astrauskas et al. 2019b].

A.1 Encoding of functions

The encoding of functions is structured into three parts:

- (1) generation of **pack** and **unpack** operations;
- (2) encoding of the function signature;
- (3) encoding of function's statements.

Pack and unpack. The technique described in Sec. 3 is used to annotate the MIR representation with **pack** and **unpack** operations. In the case of `force_inc`, the generated operations are shown as comments in the source program of Fig. 7. The first three **unpack** operations are required to access `r.current.x`; then at lines 22–23, further **unpack** operations are needed to access `*r.rest.0`, due to the (implicit) initialization of `rest` at line 25. Note that in order to track precisely the capabilities associated to `r.rest`, the **unpack** operation at line 22 needs to know the variant of that place, which inside the branch the compiler knows to be **Some**. At the end of the branch (lines 29–30), the **pack** operations are needed to unify the PCS with the one coming from the (empty) **else** branch, in which no **unpack** was ever performed. Note that in this case, joining the two PCSs is only possible by using **pack** operations, because line 22 used branch-specific information. Finally, at lines 33–35, the **pack** operations are used to restore the initial fully-packed state of `r`: the only argument of reference type, whose capabilities need to be transferred back to the caller.



Fig. 7. An encoding of `force_inc` to Viper. The function is augmented with `pack` and `unpack` operations, and is then translated to the Viper method shown on the right. The encoding of the postcondition `<post>` is provided in App. A of our technical report [Astrauskas et al. 2019b]. Note that Rust implicitly dereferences `r` when accessing its fields.

Pre- and postcondition. The signature of the function is then used to generate the signature of the Viper method, in which each Rust argument is encoded as a Viper reference. The type of the arguments is used to encode the first **inhale** in the Viper method, which encodes the permissions of the precondition — that is, the capabilities transferred to the function. Similarly, the last **exhale** statement encodes the permissions of the postcondition — that is, the capabilities that the function gives back to the caller. In Fig. 7, only the capabilities of `r` go back to the caller, because it is the only argument of type reference. When specified, the functional specification of the precondition is conjoined to the expression of the first **inhale**, and the postcondition to the expression of the last **exhale**, as in the case of `force_inc`.

Statements. Each statement is then encoded independently. Considering Fig. 7:

- (1) At *unpack-1* and *unpack-2*, the **unpack** statements are encoded as Viper’s **unfold** statements. Similarly, at *pack-1* and *pack-2*, the **pack** statements are encoded as Viper’s **fold** statements. Since which fields belong to the enumeration `r.rest` depends on the value of the discriminant, the related **pack** and **unpack** PCS operations also depend on the variant of the enumeration — in this case **Some**. The **pack** operations at *unpack-2* are needed to join with the PCS of the (empty) else branch.
- (2) At *pure-call* an assignment and call of a function marked as **pure** is encoded with a corresponding Viper assignment and function call. Since the Viper encoding uses one more indirection than Rust, any statement that needs to access the value of a Rust type is wrapped in pairs of **unfold-fold** statements; in this case to access the last field of `min_x.val_i32` and `r.val_ref.current.x.val_i32`.
- (3) The **if let** construct is encoded as a lookup of the value of the discriminant, used immediately afterwards as a branch condition. The **unfolding** expression acts similarly to a pair of **unfold-fold** statements, temporarily unfolding `r.val_ref.rest` to access the discriminant field.
- (4) At *assign*, the allocation and initialisation of the local variable `q` is encoded using a Viper local variable and an assignment that sets its referenced object. The predicate instance that encodes its type, `RefMutRoute`, is obtained in three steps: first the variable obtains at its allocation the access predicate for the `val_ref` field; then, the predicate instance for the referenced object is obtained from the assignment; finally, a **fold** statement packs the wanted predicate.
- (5) At *copy* and *call*, a call of a non-pure function is encoded. Initially, a temporary local variable `arg2` is introduced to encode the copy of the second argument. Then, at *call*, the precondition of the called function is exhaled and its postcondition inhaled. The label `call_pre` is used in the encoding of expressions of the postcondition that refer to the state of the precondition.

ACKNOWLEDGEMENTS

We would like to warmly thank Nicholas D. Matsakis, Nick Cameron, Derek Dreyer, Ralf Jung, Klaus Havelund and the anonymous reviewers for their extensive feedback on this work. We are grateful to Alexandra Bugariu, Jérôme Dohrau, Marco Eilers, and Arshavir Ter-Gabrielyan for feedback on paper drafts and our artefact. We are also very grateful to Florian Hahn for his work on a precursor to this project [Hahn 2015], and to Malte Schwerhoff and Nicolas Trüssel for their generous help with our evaluation.

This work was partially funded by the Swiss National Science Foundation (SNSF) under Grant No. 200021_169503.

REFERENCES

Martin Abadi and Marcelo Fiore. 1996. Syntactic Considerations on Recursive Types. In *Logic in Computer Science (LICS)*. IEEE Computer Society, 242–252.

- Vytautas Astrauskas, Peter Müller, Federico Poli, and Alexander J. Summers. 2019a. Artefact containing the prototype implementation. <https://doi.org/10.5281/zenodo.3363914>
- Vytautas Astrauskas, Peter Müller, Federico Poli, and Alexander J. Summers. 2019b. *Leveraging Rust Types for Modular Specification and Verification*. Technical Report. ETH Zurich.
- Alexander Bakst and Ranjit Jhala. 2016. Predicate Abstraction for Linked Data Structures. In *Verification, Model Checking, and Abstract Interpretation (VMCAI) (Lecture Notes in Computer Science)*, Barbara Jobstmann and K. Rustan M. Leino (Eds.), Vol. 9583. Springer, 65–84.
- Marek Baranowski, Shaobo He, and Zvonimir Rakamarić. 2018. Verifying Rust Programs with SMACK. In *Automated Technology for Verification and Analysis (ATVA) (Lecture Notes in Computer Science)*, Shuvendu K. Lahiri and Chao Wang (Eds.), Vol. 11138. Springer, 528–535.
- Mike Barnett, Manuel Fähndrich, K. Rustan M. Leino, Peter Müller, Wolfram Schulte, and Herman Venter. 2011. Specification and Verification: The Spec# Experience. *Commun. ACM* 54, 6 (June 2011), 81–91.
- Karthikeyan Bhargavan, Barry Bond, Antoine Delignat-Lavaud, Cédric Fournet, Chris Hawblitzel, Catalin Hritcu, Samin Ishtiaq, Markulf Kohlweiss, K. Rustan M. Leino, Jay R. Lorch, Kenji Maillard, Jianyang Pan, Bryan Parno, Jonathan Protzenko, Tahina Ramananandro, Ashay Rane, Aseem Rastogi, Nikhil Swamy, Laure Thompson, Peng Wang, Santiago Zanella Béguelin, and Jean Karim Zinzindohoue. 2017. Everest: Towards a Verified, Drop-in Replacement of HTTPS. In *Summit on Advances in Programming Languages (SNAPL) (LIPICs)*, Benjamin S. Lerner, Rastislav Bodík, and Shriram Krishnamurthi (Eds.), Vol. 71. Schloss Dagstuhl, 1:1–1:12.
- Kevin Bierhoff. 2011. Automated Program Verification Made SYMPLAR: Symbolic Permissions for Lightweight Automated Reasoning. In *Symposium on New Ideas, New Paradigms, and Reflections on Programming and Software (Onward! 2011)*. ACM, 19–32.
- John Boyland. 2003. Checking Interference with Fractional Permissions. In *Static Analysis Symposium (SAS) (Lecture Notes in Computer Science)*, Radhia Cousot (Ed.), Vol. 2694. Springer, 55–72.
- John Boyland, James Noble, and William Retert. 2001. Capabilities for Sharing: A Generalisation of Uniqueness and Read-Only. In *European Conference on Object-Oriented Programming (ECOOP) (Lecture Notes in Computer Science)*, Jørgen Lindskov Knudsen (Ed.), Vol. 2072. Springer, 2–27.
- Sylvan Clebsch, Sophia Drossopoulou, Sebastian Blessing, and Andy McNeil. 2015. Deny Capabilities for Safe, Fast Actors. In *International Workshop on Programming Based on Actors, Agents, and Decentralized Control (AGERE! 2015)*. ACM, 1–12.
- Clippy contributors. 2019. Clippy. <https://github.com/rust-lang/rust-clippy> Accessed April 4, 2019.
- Ernie Cohen, Markus Dahlweid, Mark A. Hillebrand, Dirk Leinenbach, Michal Moskal, Thomas Santen, Wolfram Schulte, and Stephan Tobies. 2009. VCC: A Practical System for Verifying Concurrent C. In *Theorem Proving in Higher Order Logics (TPHOLs) (Lecture Notes in Computer Science)*, Stefan Berghofer, Tobias Nipkow, Christian Urban, and Makarius Wenzel (Eds.), Vol. 5674. Springer, 23–42.
- Coq Team. 2014. The Coq Proof Assistant Reference Manual. <http://coq.inria.fr>.
- Karl Cray, Robert Harper, and Sidd Puri. 1999. What is a Recursive Module?. In *Programming Language Design and Implementation (PLDI)*, Barbara G. Ryder and Benjamin G. Zorn (Eds.). ACM, 50–63.
- Leonardo Mendonça de Moura, Soonho Kong, Jeremy Avigad, Floris van Doorn, and Jakob von Raumer. 2015. The Lean Theorem Prover (System Description). In *Automated Deduction (CADE) (Lecture Notes in Computer Science)*, Amy P. Felty and Aart Middeldorp (Eds.), Vol. 9195. Springer, 378–388.
- Robert Dockins, Adam Foltzer, Joe Hendrix, Brian Huffman, Dylan McNamee, and Aaron Tomb. 2016. Constructing Semantic Models of Programs with the Software Analysis Workbench. *Theories, Tools, and Experiments (VSTTE) (Lecture Notes in Computer Science)*, Sandrine Blazy and Marsha Chechik (Eds.), Vol. 9971. 56–72.
- J. Nathan Foster, Michael B. Greenwald, Jonathan T. Moore, Benjamin C. Pierce, and Alan Schmitt. 2005. Combinators for Bi-directional Tree Transformations: A Linguistic Approach to the View Update Problem. *SIGPLAN Not.* 40, 1 (Jan. 2005), 233–246.
- Jean-Yves Girard. 1987. Linear Logic. *Theor. Comput. Sci.* 50, 1 (Jan. 1987), 1–102.
- Colin S. Gordon, Matthew J. Parkinson, Jared Parsons, Aleks Bromfield, and Joe Duffy. 2012. Uniqueness and Reference Immutability for Safe Parallelism. *SIGPLAN Not.* 47, 10 (Oct. 2012), 21–40.
- Florian Hahn. 2015. *Rust2Viper: Building a static verifier for Rust*. Master’s thesis. ETH Zurich.
- Philipp Haller and Martin Odersky. 2010. Capabilities for Uniqueness and Borrowing. In *European Conference on Object-Oriented Programming (ECOOP) (Lecture Notes in Computer Science)*, Theo D’Hondt (Ed.), Vol. 6183. Springer, 354–378.
- Chris Hawblitzel, Jon Howell, Manos Kapritsos, Jacob R. Lorch, Bryan Parno, Michael L. Roberts, Srinath Setty, and Brian Zill. 2015. IronFleet: Proving Practical Distributed Systems Correct. In *Symposium on Operating Systems Principles (SOSP)*, Ethan L. Miller and Steven Hand (Eds.). ACM, 1–17.
- Chris Hawblitzel, Jon Howell, Jacob R. Lorch, Arjun Narayan, Bryan Parno, Danfeng Zhang, and Brian Zill. 2014. Ironclad Apps: End-to-End Security via Automated Full-System Verification. In *Operating Systems Design and Implementation (OSDI)*, Jason Flinn and Hank Levy (Eds.). USENIX Association, 165–181.

- Stefan Heule, K. Rustan M. Leino, Peter Müller, and Alexander J. Summers. 2013. Abstract Read Permissions: Fractional Permissions without the Fractions. In *Verification, Model Checking, and Abstract Interpretation (VMCAI) (Lecture Notes in Computer Science)*, Vol. 7737. Springer, 315–334.
- Cliff B. Jones. 1983. Specification and design of (parallel) programs. In *IFIP Congress*. 321–332.
- Ralf Jung, Jacques-Henri Jourdan, Robbert Krebbers, and Derek Dreyer. 2018. RustBelt: Securing the Foundations of the Rust Programming Language. *PACMPL* 2, POPL (2018), 66:1–66:34.
- 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. *ACM SIGPLAN Notices* 50, 1 (2015), 637–650.
- Shuanglong Kan, David Sanán, Shang-Wei Lin, and Yang Liu. 2018. K-Rust: An Executable Formal Semantics for Rust. *CoRR* abs/1804.07608 (2018). arXiv:1804.07608 <http://arxiv.org/abs/1804.07608>
- Ioannis T. Kassios. 2011. The Dynamic Frames Theory. *Formal Aspects of Computing* 23, 3 (2011), 267–289.
- Gerwin Klein, Kevin Elphinstone, Gernot Heiser, June Andronick, David Cock, Philip Derrin, Dhammika Elkaduwe, Kai Engelhardt, Rafal Kolanski, Michael Norrish, Thomas Sewell, Harvey Tuch, and Simon Winwood. 2009. seL4: Formal Verification of an OS Kernel. In *Symposium on Operating Systems Principles (SOSP)*, Jeanna Neefe Matthews and Thomas E. Anderson (Eds.). ACM, 207–220.
- Gary T. Leavens, Erik Poll, Curtis Clifton, Yoonsik Cheon, Clyde Ruby, David Cok, Peter Müller, Joseph Kiniry, Patrice Chalin, Daniel M. Zimmerman, and Werner Dietl. 2011. *JML Reference Manual*. <http://www.jmlspecs.org/>.
- K. Rustan M. Leino. 2010. Dafny: An Automatic Program Verifier for Functional Correctness. In *Logic for Programming, Artificial Intelligence, and Reasoning (LPAR) (Lecture Notes in Computer Science)*, Edmund M. Clarke and Andrei Voronkov (Eds.), Vol. 6355. Springer, 348–370.
- K. Rustan M. Leino and Peter Müller. 2004. Object Invariants in Dynamic Contexts. In *European Conference on Object-Oriented Programming (ECOOP) (Lecture Notes in Computer Science)*, M. Odersky (Ed.), Vol. 3086. Springer, 491–516.
- K. Rustan M. Leino and Greg Nelson. 2002. Data Abstraction and Information Hiding. *ACM Trans. Program. Lang. Syst.* 24, 5 (2002), 491–553.
- Marcus Lindner, Jorge Aparicius, and Per Lindgren. 2018. No Panic! Verification of Rust Programs by Symbolic Execution. In *Industrial Informatics (INDIN)*. IEEE, 108–114.
- Maroua Maalej, Tucker Taft, and Yannick Moy. 2018. Safe Dynamic Memory Management in Ada and SPARK. (2018).
- Nicholas D. Matsakis. 2018a. MIR-based borrow check (NLL) status update. <http://smallcultfollowing.com/babysteps/blog/2018/06/15/mir-based-borrow-check-nll-status-update> Accessed April 4, 2019.
- Nicholas D. Matsakis. 2018b. MIR-based borrowck is almost here. <http://smallcultfollowing.com/babysteps/blog/2018/10/31/mir-based-borrowck-is-almost-here/> Accessed April 4, 2019.
- Nicholas D. Matsakis and Felix S. Klock II. 2014. The Rust language. In *ACM SIGAda Ada Letters*, Vol. 34. ACM, 103–104.
- Bertrand Meyer. 1992. Design by Contract. In *Advances in object-oriented software engineering*, Dino Mandrioli and Bertrand Meyer (Eds.). Prentice Hall, 1–50.
- P. Müller. 2002. *Modular Specification and Verification of Object-Oriented Programs*. Lecture Notes in Computer Science, Vol. 2262. Springer.
- Peter Müller, Malte Schwerhoff, and Alexander J. Summers. 2016. Viper: A Verification Infrastructure for Permission-Based Reasoning. In *VMCAI (Lecture Notes in Computer Science)*, Barbara Jobstmann and K. Rustan M. Leino (Eds.), Vol. 9583. Springer, 41–62.
- Peter W. O’Hearn. 2004. Resources, Concurrency and Local Reasoning. In *Concurrency Theory (CONCUR) (Lecture Notes in Computer Science)*, Philippa Gardner and Nobuko Yoshida (Eds.), Vol. 3170. Springer, 49–67.
- Peter W. O’Hearn, John C. Reynolds, and Hongseok Yang. 2001. Local Reasoning about Programs that Alter Data Structures. In *Computer Science Logic (CSL) (Lecture Notes in Computer Science)*, Laurent Fribourg (Ed.), Vol. 2142. Springer, 1–19.
- Susan Owicki and David Gries. 1976. Verifying Properties of Parallel Programs: An Axiomatic Approach. *Commun. ACM* 19, 5 (May 1976), 279–285.
- Matthew J. Parkinson and Gavin M. Bierman. 2005. Separation logic and abstraction. In *Principles of Programming Languages (POPL)*, Jens Palsberg and Martin Abadi (Eds.). ACM, 247–258.
- Matthew J. Parkinson and Alexander J. Summers. 2012. The Relationship Between Separation Logic and Implicit Dynamic Frames. *Logical Methods in Computer Science* 8, 3:01 (2012), 1–54.
- Eric W. Reed. 2015. *Patina: A Formalization of the Rust Programming Language*. Technical Report UW-CSE-15-03-02. University of Washington.
- John C. Reynolds. 2002. Separation Logic: A Logic for Shared Mutable Data Structures. In *Logic in Computer Science (LICS)*. IEEE Computer Society, 55–74.
- Patrick M. Rondon, Ming Kawaguci, and Ranjit Jhala. 2008. Liquid Types. *SIGPLAN Not.* 43, 6 (June 2008), 159–169.
- Rosetta Code contributors. 2018. Rosetta Code. <https://rosettacode.org/wiki/Category:Rust> Accessed November 5, 2018.
- Rust community. 2017. Non-Lexical Lifetimes RFC. <https://github.com/rust-lang/rfcs/blob/master/text/2094-nll.md> Accessed November 4, 2018.

- Rust community. 2018a. The Rust community's crate registry. <https://crates.io> Downloaded on November 2, 2018.
- Rust community. 2018b. Rust: The Reference — Place Expressions and Value Expressions. <https://doc.rust-lang.org/reference/expressions.html#place-expressions-and-value-expressions> Accessed November 4, 2018.
- Rust community. 2019. Learn Rust by writing Entirely Too Many Linked Lists. <https://rust-unofficial.github.io/too-many-lists/first-final.html> Accessed April 4, 2019.
- Rust contributors. 2019a. The Polonius Reference Implementation for the Rust Borrow-Checker. <https://github.com/rust-lang/polonius> Accessed April 4, 2019.
- Rust contributors. 2019b. The Rustonomicon: Working with Unsafe. <https://doc.rust-lang.org/nomicon/working-with-unsafe.html> Accessed April 4, 2019.
- Rust contributors. 2019c. Tracking issue for generalized two-phase borrows. <https://github.com/rust-lang/rust/issues/49434> Accessed April 4, 2019.
- Malte Schwerhoff and Alexander J. Summers. 2015. Lightweight Support for Magic Wands in an Automatic Verifier. In *European Conference on Object-Oriented Programming (ECOOP) (LIPICs)*, J. T. Boyland (Ed.), Vol. 37. Schloss Dagstuhl, 614–638.
- Jan Smans, Bart Jacobs, and Frank Piessens. 2009. Implicit Dynamic Frames: Combining Dynamic Frames and Separation Logic. In *European Conference on Object-Oriented Programming (ECOOP) (Lecture Notes in Computer Science)*, Sophia Drossopoulou (Ed.), Vol. 5653. Springer, 148–172.
- Jan Smans, Bart Jacobs, and Frank Piessens. 2010. Heap-Dependent Expressions in Separation Logic. In *Formal Techniques for Distributed Systems (FMOODS/FORTE) (Lecture Notes in Computer Science)*, John Hatcliff and Elena Zucca (Eds.), Vol. 6117. Springer, 170–185.
- Sven Stork, Karl Naden, Joshua Sunshine, Manuel Mohr, Alcides Fonseca, Paulo Marques, and Jonathan Aldrich. 2014. AEMinium: A Permission-Based Concurrent-by-Default Programming Language Approach. *ACM Trans. Program. Lang. Syst.* 36, 1 (March 2014), 2:1–2:42.
- Alexander J. Summers and Sophia Drossopoulou. 2013. A Formal Semantics for Isorecursive and Equirecursive State Abstractions. In *European Conference on Object-Oriented Programming (ECOOP) (Lecture Notes in Computer Science)*, Giuseppe Castagna (Ed.), Vol. 7920. Springer, 129–153.
- John Toman, Stuart Pernsteiner, and Emina Torlak. 2015. Crust: A Bounded Verifier for Rust (N). In *Automated Software Engineering (ASE)*, Myra B. Cohen, Lars Grunske, and Michael Whalen (Eds.). IEEE, 75–80.
- Sebastian Ullrich. 2016. *Simple Verification of Rust Programs via Functional Purification*. Master's thesis. Karlsruhe Institute of Technology.
- Feng Wang, Fu Song, Min Zhang, Xiaoran Zhu, and Jun Zhang. 2018. KRust: A Formal Executable Semantics of Rust. *CoRR* abs/1804.10806 (2018). arXiv:1804.10806 <http://arxiv.org/abs/1804.10806>
- Aaron Weiss, Daniel Patterson, and Amal Ahmed. 2018. Rust Distilled: An Expressive Tower of Languages. *arXiv preprint arXiv:1806.02693* (2018).