

# Surveying the Rust Verification Landscape

Alex Le Blanc  
University of Waterloo  
Canada  
a6leblan@uwaterloo.ca

Patrick Lam  
University of Waterloo  
Canada  
patrick.lam@uwaterloo.ca

## ABSTRACT

Rust aims to be a safe programming language applicable to systems programming applications. In particular, its type system has strong guardrails to prevent a variety of issues, such as memory safety bugs and data races. However, these guardrails can be sidestepped via the `unsafe` keyword. `unsafe` allows certain otherwise-prohibited operations, but shifts the onus of preventing undefined behaviour from the Rust language’s compile-time checks to the developer. We believe that tools have a role to play in ensuring the absence of undefined behaviour in the presence of unsafe code. Moreover, safety aside, programs would also benefit from being verified for functional correctness, ensuring that they meet their specifications.

In this research proposal, we explore what it means to do Rust verification. Specifically, we explore which properties are worth verifying for Rust; what techniques exist to verify them; and which code is worth verifying. In doing so, we motivate an effort to verify safety properties of the Rust standard library, presenting the relevant challenges along with ideas to address them.

## 1 INTRODUCTION

Rust is a modern programming language which aims to be ergonomic and suitable for high-performance systems programming while making it harder to write bugs. As the language designers put it, as quoted in the Rust Book [Klabnik and Nichols 2024]:

The Rust programming language helps you write faster, more reliable software.

Specifically, Rust has a number of features that are novel to widely-deployed programming languages, notably its ownership-based type system used for memory management and to guarantee race freedom. Multi-million line Rust codebases exist; even in 2020, the Servo and TiKV projects exceeded 2 million lines of code [Anderson 2020], in addition to the Rust compiler itself (another million lines). Unlike `C/C++`, Rust aims to limit undefined behaviour to code specifically marked as `unsafe`; and, unlike Java, Rust does not pretend that `unsafe` code does not exist.

Amazon is interested in verifying the Rust standard library<sup>1</sup> and aims to incite community participation. While the Rust authors’ goal is to enable developers to write “faster, more reliable” software, we are specifically interested in contributing techniques to this effort that strengthen specific aspects of Rust’s “more reliable” aspiration and understanding human factors related to carrying out this verification.

In this work, our goal is to sound out the surrounding landscape. We explore a programme for “verifying Rust”—what might it mean

<sup>1</sup><https://model-checking.github.io/verify-rust-std/intro.html>

to verify some nontrivial body of Rust code? Does that mean that Rust code should not crash? Or should it have certain specific functional correctness properties? Are we talking about all Rust code, or just code written in Safe Rust? Is the standard library a reasonable place for us to start? What are the advantages and disadvantages of Amazon’s Kani tool and others in its space? What exists in other languages?

Specifically, in this research proposal, we consider three fundamental questions: what properties should be verified; how should they be verified; and what code should be verified?

We start (Section 2) by exploring which properties are worth verifying. Rust generally aims to prevent memory safety errors and race conditions, but as a systems programming language, it also explicitly allows programmers to use Unsafe Rust, bypassing safeguards. We also discuss properties beyond generic safety properties, namely functional correctness or domain-specific properties. Although our goal is to verify Rust, we also survey (Section 3) the situation of program verification in `C/C++/Java`, as well as existing tools for verifying Rust and the techniques they use.

We also investigate (Section 4) which code is worth verifying. Rust applications rely on libraries, and almost always on the Rust standard library. Applications, libraries, and the standard library have different verification needs: one might expect that verifying domain-specific properties is more relatively valuable for application code than verifying memory safety of `unsafe` code. The standard library, by contrast, is a heavy user of `unsafe` code, and its safety is assumed by the vast amounts of code that depend on it, so that verifying memory safety is especially important there. We then describe (Section 5) our proposed approach to verifying Rust and anticipated challenges. Finally, we discuss (Section 6) additional long-term objectives related to human factors, beyond just the verification of Rust code.

## 2 RUST, PROPERTIES, AND UNSAFE CODE

We first explore the question of what properties we focus on. One could take an empirical approach and focus on techniques that address the most common types of Rust bugs, or perhaps the most important bugs; in a related vein, Shirzad and Lam [2024] investigated Rust bug fix patterns, and specifically those related to borrowing.

However, one of Rust’s primary design goals is to be a “safe” programming language, while also supporting low-level features needed for systems programming. Thus, per the Rustonomicon [Rust team 2016] (part of Rust’s official documentation), the subset of Rust known as Safe Rust is designed to prevent undefined behaviour, providing various safety guarantees. For instance, by design, Safe Rust programs ought to be memory safe and free of data races. The Rustonomicon includes the sentence:

Safe Rust is the true Rust programming language.

Unsafe Rust differs from Safe Rust in that undefined behaviour can happen. Because Safe Rust is free of undefined behaviour, even buggy programs follow Rust semantics. On the other hand, there are no semantics for any Unsafe Rust code that executes undefined behaviour. Thus, one of the primary goals behind the research described in this proposal is to eliminate all undefined behaviour, by verifying Unsafe Rust. This then enables sound reasoning about the behaviour of the entire program. In particular, we plan to start by verifying the Unsafe Rust in the standard library.

As a caveat, to our knowledge, the “no undefined behaviour” guarantee of Safe Rust has not been formally proven. An effort analogous to the CompCert project [Leroy 2009] would be required to formally show the absence of implementation bugs in the specification and compiler. Limiting the scope of such an effort to Safe Rust is likely to be simpler than aiming to verify an optimizing C compiler, as in CompCert’s case. However, there is empirical evidence that suggests that Safe Rust’s guarantees hold in practice. Xu et al. [2021] report, from a study of 186 Rust CVEs, that undefined behaviour in their Rust programs always came from bugs in Unsafe Rust code, rather than Safe Rust.

For the work that we are proposing, we believe that it is reasonable to assume that (putting aside language specification and implementation bugs) a Rust programmer must use Unsafe Rust to get undefined behaviour—either directly, or by interacting with Unsafe Rust. In contrast to Safe Rust, Unsafe Rust explicitly includes constructs that can trigger undefined behaviour if used inappropriately.

Contrast Rust’s situation with C and C++, where any line of code may inadvertently trigger undefined behaviour. Rust has a significant advantage here: it is feasible to audit Rust code for the program locations which may contain undefined behaviour, because such locations must be associated with some usage of the `unsafe` keyword.

As an example of undefined behaviour in Rust, the Rustonomicon sketches<sup>2</sup> an implementation for a struct that we call `MyVec`. This struct could include a `make_room` function (in Safe Rust) that modifies the `MyVec`’s capacity, along with unsafe code that relies on the capacity being accurate. A client with access to `make_room` can trigger undefined behaviour in the unsafe code. `make_room` does not have to be accessible by clients; if it is inaccessible, then `MyVec`’s unsafety is encapsulated, all of its clients are sound, and undefined behaviour is impossible. If it is accessible, then only `MyVec` clients that use `make_room` appropriately are sound, and a verifier must check that all calls to that method (1) meet necessary preconditions and (2) do not misuse returned values.

*Encapsulating unsafety.* Within a module that contains Unsafe Rust, undefined behaviour may be triggered by errors in the unsafe code, or by safe code that nevertheless has access to directly break invariants that the unsafe code depends on. As we have seen in the example above, it is possible for an Unsafe Rust-containing module to provide access to functions that break its invariants; clients of that module can—if they break the safety contract—indirectly break the invariants and cause unsafe behaviour, but they cannot break them directly.

Some libraries *encapsulate* unsafety, so that clients can never trigger undefined behaviour. That is, for those libraries, the required safety conditions on the client are nil. As we have seen, though, other libraries (including the Rust standard library), do not. Ozdemir [2019] names the phenomenon of poorly-encapsulated safety properties “Abstraction Escape” and proposes a static analysis to detect simple cases of it, but does not empirically estimate how often it occurs. Astrauskas et al. [2020] report that unsafe Rust code (in their sample set) is typically well-encapsulated. Qin et al. [2020] define the notion of “interior unsafe” (i.e., a safe function that encapsulates an unsafe function) and find that 58% of sampled interior unsafe functions in the Rust standard library require conditions to ensure safety but do not check these conditions, relying on being called correctly. Note that this is not a bug—in Rust, the onus is on the caller to ensure documented safety conditions for calling unsafe `fns`, and the callee is entitled to ensure that they hold.

It is possible that safety encapsulation holds much more often for arbitrary libraries than for the Rust standard library. That is, the standard library may well expose more low-level implementation details than other libraries. This question is an important one in terms of understanding how humans use Rust’s safety mechanisms, but it is one that needs to be empirically investigated.

*Specification language design for unsafety.* Libraries that expose (rather than encapsulate) safety are required to label relevant functions as unsafe `fns`, and to document (in natural language) the required safety conditions. unsafe `fns` must be called from unsafe `{ }` blocks in the caller. The meaning of unsafe `fn` is that the function contract includes (safety-related) requirements that the caller is responsible for ensuring, and furthermore, that neither the compiler nor the callee are responsible for. (In some cases, either the compiler or the callee may fail fast and refuse to compile or execute non-conforming code.) Dually, implementations of functions in traits declared as unsafe `trait` must satisfy requirements that the compiler is not responsible for enforcing. If an Unsafe Rust-using library imposes usage constraints on its client, and the client violates them, then Rust calls the client’s use of the library *unsound*, and the client has triggered undefined behaviour. In turn, the client may encapsulate the unsafety, or it may pass the buck to its own callers, by exposing its own unsafe `fns`.

Verifying that client usages of unsafe `fns` are sound requires designers to state, and verifiers to check, some properties that go beyond the expressive capabilities of the Rust type system. We are not aware of efforts to ensure that C code meets the required postconditions in analogous cases.

An interesting point for specification language design is that specifications for unsafe `fns` may, somewhat unusually, restrict what can happen after the function returns, in addition to documenting (as usual) required preconditions and ensured postconditions. To some extent, these restrictions occur in other languages too. One example is that in C, after returning from `free(x)`, the caller must no longer access `x`. Similarly, in Rust, value `x` that has been through `ManuallyDropped<t>::drop(x)` must no longer be accessed. As another example, somewhat analogous to the C `restrict` qualifier, the `as_mut` implementation on `pointer` returns a value that Rust will assume is unique, but that may actually have aliases; its postcondition requires that there be no accesses to the returned value

<sup>2</sup> *Rustonomicon*, section 1.3, “Working with Unsafe”: <https://doc.rust-lang.org/nomicon/working-with-unsafe.html>

through other pointers (which the type system does *not* prevent from existing) while the returned value remains in scope.

In the absence of safety encapsulation and in the presence of Unsafe Rust-using libraries, even verifying the absence of undefined behaviour may require arbitrary reasoning power and properties; the reason for using `unsafe` is that the library developer believes that the required reasoning power is more than what Rust’s type system provides.

The proposed verification efforts for Rust will require some thought about human aspects of verification, namely a notation that allows library developers to formally specify the needed postconditions, and that client code can be verified against. [Cui et al. 2024] identified 19 classes of safety properties that various parts of standard library API documentation specified as requirements for callers to unsafe APIs. We believe that this notation could well be specialized for the use case of safety properties, and might turn out to be unnecessary for verifying general properties of Rust code.

*Beyond generic safety properties.* So far, we have discussed generic safety properties guaranteed by a lack of undefined behaviour, such as memory safety and data race freedom. These properties ought to hold in all programs, and violations of these properties mean that a program is not guaranteed to satisfy any properties at all. However, a lack of undefined behaviour is not enough.

Another way for Rust programs to go wrong is by panicking. For instance, Rust programs panic when they call `unwrap()` on a `None` value. Prusti [Astrauskas et al. 2022] verifies that Rust programs do not panic due to integer overflows—overflows do not panic on release builds and are explicitly listed in the Rust documentation as not being undefined behaviour, but are likely to be undesired behaviour. Prusti apparently does not aim to verify absence of all panics, and for instance, does not detect unwrapping a `None`.

Classically, program verification efforts consider module or function preconditions and postconditions (Hoare triples), as well as invariants; we have mentioned safety-related postconditions above, and there are also safety-related preconditions. More generally, there can be functional correctness properties, which we refer to as *domain-specific properties*. An example precondition for `MyVec` is that calling `insert()` with an insertion index is only valid if the insertion index is less than the current length of the `MyVec`, promising a panic if the precondition is violated. Such properties encode specifications, which may be implicit or written in natural language. Verifying the Rust standard library could also be construed to require encoding and verifying these domain-specific properties, but writing a complete set of properties is at best an art, and we do not believe that there is consensus on what this set of properties should encompass.

For many programs, total correctness is desirable, requiring proof of termination. Researchers have also developed a vast array of techniques to verify termination of programs [Cook et al. 2006; Cousot and Cousot 2012; Zhu and Kincaid 2021].

### 3 EXTANT VERIFICATION TOOLS AND APPROACHES

We survey parts of the software verification landscape that are relevant to our efforts to verify Rust. We focus on Java and C/C++, imperative languages (like Rust) that have mature verification ecosystems, sharing important language features with Rust, while also having certain key differences we can learn from. We also briefly discuss Dafny to touch on the interesting case of languages that have built-in verification features. Then, we discuss some of the many Rust verification efforts by others, including the techniques that existing tools use.

#### 3.1 Other Languages

*Java.* As defined in the Java Language Specification and Java Virtual Machine Specification, Java guarantees memory safety, though not race freedom (which is highly challenging to ensure statically in the presence of unrestricted aliasing). Featherweight Java [Igarashi et al. 2001] is an influential work in formally reasoning about Java and enables proofs of Java’s type safety.

The real situation for Java is not as straightforward and is closer to Rust than one might think. Evans [2020] discusses Unsafe Java features—which exist and are in the Oracle API documentation but not in the official specifications—and how they are used in practice, ubiquitously, by libraries to achieve better performance. As with Rust, we believe that most developers do stick to “safe Java” and that the unsafe uses are encapsulated within libraries. Mastrangelo et al. [2015] found that 3% of top-ranked Java artifacts (as of the date of their study) directly used the version of Unsafe available then, and 25% of artifacts directly or indirectly depend on unsafe code. One difference with Rust appears to be that Rust libraries impose usage conditions on clients for their uses to be sound, and hence to ensure safety, while it appears to us that Java libraries aim to be unconditionally safe, or in other words, to encapsulate safety.

Nevertheless, almost all Java verifiers focus on the verification of “safe” Java code (examples include JBMC [Cordeiro et al. 2018], KeY [Ahrendt et al. 2014], and OpenJML [Cok 2011]). To our knowledge, the only tool that acknowledges the existence of unsafe Java code is SafeCheck [Huang et al. 2019], which only slightly mitigates the dangers of unsafe Java by adding run-time checks for safety violations and hence improving diagnosability of crashes caused by unsafe code. It does not carry out any static checking of either generic safety properties or domain-specific properties. This is similar to ASan for C/C++ [Serebryany et al. 2012] or ERASan for Rust [Min et al. 2024]; these dynamic approaches can terminate programs that attempt to trigger undefined behaviour, but not proactively prevent it before runtime.

We also draw an analogy between Java native code and Rust intrinsics, which pose the same kinds of challenges for program analysis and verification. The approach in Java is generally to hard-code the behaviour of the intrinsics and rely on a consensus agreement on this behaviour. Sometimes it is possible to provide a mock implementation of the intrinsic in Java or Rust, but there is no guarantee that the real implementation is faithful; we cannot see any static way of providing guarantees.

*C/C++*. Rather than being safe by default, the C language is entirely unsafe, meaning that verifiers, regardless of whether they focus on generic safety or domain-specific properties, must always consider a broader range of undefined behaviour. One of the more popular tools is CBMC [Clarke et al. 2004], a bounded model checker, whose strengths align particularly well with verifying unsafe and concurrent C code—it is less reliant on invariants than symbolic execution approaches.

Other tools take the deductive approach to verification of C code, such as VST [Appel 2011] and VCC [Cohen et al. 2009]. VST (Verified Software Toolchain) is a formal verification tool that allows users to produce machine-checkable proofs for programs, focusing on both memory safety and functional correctness. On the other hand, VCC takes a modular approach and translates annotated C code into BoogiePL, which is used to generate verification conditions that can be solved by Z3 [De Moura and Bjørner 2008]. VCC also verifies both memory safety and functional correctness, but it places a particular focus on concurrent code. Both tools, like effectively any other deductive tool, have the potentially costly requirement of needing function contracts and loop invariants, but in return they support unbounded verification, allowing for stronger guarantees of correctness.

C++ has additional features (e.g., bounds-checking for certain types, more expressive type system including uniqueness). But these features are easy to override—escape hatches are ubiquitous. It therefore seems that most C++ verification tools are extensions of C tools which support C++ features, but which do not trust the stronger type constraints.

*Dafny*. Some languages are designed specifically for verification, such as Dafny [Leino 2023]. Dafny supports a rich specification language and static verification of specifications. While Dafny can compile to existing languages, we consider Dafny and languages like it to be out of scope in our discussion of verification of code in existing language ecosystems.

Having discussed Java and C/C++, the situation for Rust is closest to Java's, except that we suspect unsafe code to be much more widespread in Rust (given that its “safe” mode is more restrictive than Java's). Mastrangelo et al. [2015] found that 25% of surveyed Java artifacts depended on unsafe code, whereas Evans et al. [2020] found that over half of surveyed Rust crates depended on unsafe code. Rust's unsafe mode is quite similar to C in its power (e.g., it allows manipulating raw pointers), although unsafe Rust code is easy to identify, unlike C code with undefined behaviour, which may be anywhere. Understanding how other languages deal with these challenges can give us insight into how we can do the same for Rust.

### 3.2 Rust Verification Tools and Techniques

There exist a variety of tools for the verification of Rust code, as well as verification frameworks for the implementation and semantics of the Rust language itself, and we survey them here. Many Rust verification tools aim to check properties of safe Rust code, e.g. lack of panics, assertions, and function contracts/invariants. Other tools verify unsafe Rust code and its usage.

*Intermediate representations*. A verification tool's workflow normally consists of taking source code annotated with function contracts according to some specification, partly compiling the code (mostly via `rustc`), generating some verification conditions, and solving them (e.g. via an SMT or SAT solver). When partly compiling, the tool is in fact lowering the code to an intermediate representation (IR), typically a mid-level IR (MIR) or, at a lower level, an IR at the level of the LLVM IR.

The LLVM IR is the last IR before an executable is produced. The main advantage of using a lower-level IR like this is that other languages also use it (e.g., C), allowing for the straightforward reuse of verification backends from other languages for Rust. Indeed, this is the case for several tools, such as SMACK [Baranowski et al. 2018] and SeaHorn [Gurfinkel et al. 2015], which were originally developed for verification of C (or more generally LLVM-based) code, but can be used on Rust code due to this shared IR. Baranowski et al. [2018] show how to verify a lack of overflows in Rust, while an example of Seahorn usage<sup>3</sup> shows the verification of assertions embedded in the Rust code. The downside of the LLVM IR approach is, due to being lower-level, that information about the original structure and semantics of the program is not readily available [VanHattum et al. 2022].

At a higher level, the Kani project by VanHattum et al. [2022] specifically chose to analyze Rust's MIR, which also allowed it to incorporate information encoded in a program's dynamic trait objects. Working at a higher level than the LLVM IR retains more Rust-specific semantic information. MIR is used only by Rust, meaning that reusing backends from other languages becomes a more involved process, but it appears to us that this access to richer semantic information is worthwhile. After all, bridging a Rust-specific IR and a generic backend can be engineered, whereas how to recover lost semantic information from a lower-level IR is unclear, and may be impossible in some cases.

*Survey of related tools*. Prusti [Astrauskas et al. 2022] is a deductive verifier that works at MIR level, though it primarily focuses on the verification of functional correctness of safe code. In particular, it verifies that certain panics do not occur in safe Rust code, as well as user-specified assertions. Creusot [Denis et al. 2022] is a more recent deductive verifier, also focusing on safe code. It proposes a specification language, Pearlite. The main difference between these two tools is that Prusti specifies mutable borrows using pledges, as opposed to the prophecies (via the `final` operator) used by Creusot. This means that Creusot can support a wider range of borrowing patterns than Prusti (which does not support, for example, reborrowing in a loop [Denis et al. 2022]).

Aeneas [Ho and Protzenko 2022] puts forth a fairly novel approach in that it translates Rust programs into a pure-lambda calculus, which allows for much flexibility in terms of the theorem provers that can be used to reason about these programs. Verus [Latuada et al. 2023], on the other hand, has programmers write proofs and specifications using Rust itself, leveraging the language's own type and borrow checking. This way, it can verify even unsafe code, though it has some trouble with more complex borrowing patterns that Aeneas can handle, such as reborrowing. Moreover, it loses out

<sup>3</sup><https://project-oak.github.io/rust-verification-tools/using-seahorn/>

on the flexibility that Aeneas has in terms of being able to reason about programs with almost any theorem prover.

Kani [VanHattum et al. 2022] uses a bounded model checker, translating code from the MIR into Goto-C, which is then passed to CBMC. Besides incorporating higher-level information from Rust, one of the main features of Kani is the verification of unsafe code, which it achieves in part due to its CBMC backend.

Finally, some tools allow for complete formal verification of subsets of the Rust language itself. For instance, RustBelt [Jung et al. 2017] developed an extensible soundness proof for  $\lambda_{\text{Rust}}$ , a formal description of the Rust language, minus certain features deemed superfluous to the proof (e.g. traits). Oxide [Weiss et al. 2019] also models the key aspects of Rust, supporting certain new features like non-lexical lifetimes. Incidentally, Rust does not have a formal memory model, so we would not be able to verify against an official memory model even if we wanted to.

Gillian-Rust [Ayoun et al. 2024] proposes a promising approach to verification of both safe and unsafe Rust. It uses a separation logic combining features of RustBelt and RustHornBelt [Matsushita et al. 2022] to reason about unsafe Rust, and integrates that logic with Creusot to reason about safe Rust (at much lower cost). RefinedRust [Gäher et al. 2024] is similar to Gillian-Rust in that it uses a separation logic to produce proofs of correctness for both safe and unsafe code, but differs in that these proofs are foundational (i.e., they are machine-checkable in a general-purpose proof assistant). We believe that Gillian-Rust and RefinedRust can verify soundness of the usage of unsafe libraries from safe code, which appears to be out of scope for RustBelt. We also speculate that these tools will be more powerful than the Kani-based approach that we propose here, but also harder to use (e.g., due to their use of separation logic).

*Approaches: bounded model checking versus symbolic execution.* In practice, Kani’s bounded model checking approach differs from the deductive approach taken by Prusti and Creusot particularly in its handling of loops. The bounded model checking approach explores all behaviours of the program up to a given bound. This approach proved to be particularly fruitful for exploring behaviours of concurrent C programs. From the verification user’s perspective, an advantage of the bounded model checking approach is that it does not require loop invariants to verify loops: the verifier simply explores all executions up to the bound. This approach also does not necessarily sacrifice soundness, at least in cases where a fixed bound can be shown to be sufficient. Of course, cases where this cannot be shown represent a limitation of bounded model checking (e.g., unbounded structures like linked lists which need verification beyond the bounds being checked). Bounded model checking approaches should also be easily modifiable to accept loop invariants, though most tools do not seem to currently support them (but Kani now does). Users of deductive verification approaches can specify loop invariants (which can be difficult) or use automatic loop invariant inference techniques (which do not always work). For these reasons, the user should have the option to choose whether or not to use loop invariants (as is the case with Kani for example)—either approach has non-negligible disadvantages and advantages.

## 4 WHICH CODE TO VERIFY

We discuss which part of the Rust ecosystem we should focus our verification efforts on. We can consider the Rust ecosystem to be composed of three main groups of code:

- (1) Client code, i.e., applications written in Rust;
- (2) Third-party crates, analogous to libraries in other languages; and
- (3) The Rust standard library.

These groups are listed from least to most fundamental to the ecosystem. That is, the standard library can be used by many crates, and these crates can be used by many applications; but not vice-versa. Note that verifying, in isolation, a third-party crate that depends on some part of the standard library without having first verified the standard library decreases our confidence in whatever we have proven about the crate. (While the standard library has the advantage that it is widely used, and thus much less likely to have easily-encountered functional correctness issues, experience has shown that security issues do still crop up in the most widely used codebases.) Thus, even though verification of any part of the ecosystem is important, if one is to prioritize certain parts, it is natural to start with the standard library.

Yet, much of the existing literature focuses on the verification of client code and third-party crates. This was the case, for instance, with Prusti [Astrauskas et al. 2022], Kani [VanHattum et al. 2022] (initially), SMACK [Baranowski et al. 2018], and Creusot [Denis et al. 2022]. The few works that attempted to verify the standard library did so on subsets of the Rust language (e.g., RustBelt [Jung et al. 2017]). Indeed, to our knowledge, the only tool that currently explicitly supports verification of the standard library is Kani, with the recent addition of the `verify-std` feature, though the actual verification effort is just beginning. This further supports a prioritization of the verification of the standard library.

Why have people chosen to verify client code? We speculate that when one’s tool handles only Safe Rust, one is likely to get further verifying client code; crates and the standard library may contain more Unsafe Rust, and that is harder to verify. Also, as we have stated earlier, verifying Safe Rust frees one to focus on assertions and contracts (if one assumes the underlying Unsafe Rust is properly encapsulated), rather than worrying about safety properties; this focus may be attractive to some researchers.

Finally, we note that a verification effort has two main outputs: a direct one in the form of a proof for a particular program, and occasionally an indirect one in the form of a new technique. We have already discussed how verifying the standard library maximizes the direct output (yielding a proof that is useful to more programs), but if we wish to maximize the total output, we should also consider the indirect output. Naturally, the areas of the ecosystem that are more challenging to verify would require more innovation to do so. We believe that the challenges that the verification of the standard library presents (at least in terms of memory safety and functional correctness) are practically a superset of those of regular crates and client code (e.g., writing function contracts for intrinsics, which we elaborate on in Section 5). (This is not completely true, since verifying client code may require techniques seldom-used in library code, and vice versa, but it seems like a defensible starting assumption.) Hence, we believe that focusing verification efforts on

the standard library has the potential to yield more fruitful verification techniques, and these new techniques could then be applied to other parts of the ecosystem.

## 5 CHOOSING A TOOL AND VERIFICATION TARGET

Having discussed some Rust verification tools as well as motivating the verification of the Rust standard library, what remains is to determine which tool would be best as a starting point for this task, as well as specific verification targets within the library.

Our first strict requirement is that the tool must support the verification of unsafe Rust code, as it is highly prevalent in the standard library. We therefore do not consider tools that focus only on safe code, like Creusot [Denis et al. 2022], Prusti [Astrauskas et al. 2022] and Aeneas [Ho and Protzenko 2022]. We are looking for a tool that can handle the full complexity of the standard library (e.g., traits) rather than a subset language; techniques like RustBelt [Jung et al. 2017] and Oxide [Weiss et al. 2019] would need to be extended. Finally, our ideal tool is user-friendly: the standard library is a nontrivially-large project, with around 229K lines of code. The process of writing the specifications and performing the proof would therefore preferably be as lightweight as possible (e.g., with less annotation required and as much automation as possible); in terms of human aspects, trading expressive power for ease-of-use as much as possible.

We believe bounded model checking to be a good choice for our goal. Writing annotations and function contracts can be a labour-intensive process, especially for larger projects. Even if bounded model checking still requires contracts, it can sometimes avoid the challenges of writing invariants, which are needed for approaches using symbolic execution, like Prusti [Astrauskas et al. 2022] and Creusot [Denis et al. 2022]). It also sidesteps the difficulty of using separation logic, required for techniques like Gillian-Rust [Ayoun et al. 2024], Verus [Lattuada et al. 2023] and RefinedRust [Gäher et al. 2024].

To our knowledge, the only tool that satisfies all of our criteria is Kani [VanHattum et al. 2022]. It also has the added advantage of supporting automated verification of some memory safety properties, which very few tools do. However, it is possible that the boundedness fundamental to bounded model checking will prevent some proofs from going through, e.g. for verifying linked list implementations: we are trading off power for ease of use and seeing how far the bounded model checking approach can take us. In any case, we will focus on using and extending Kani in our efforts to verify the standard library. The choice of verification tool, and its power-versus-annotation burden tradeoff, is a key human factor to consider for verification efforts. Our work will help future researchers choose appropriate tools for their work, based on the experiences that we will share.

For verifying the Rust standard library, we consider two main classes of challenges:

- (1) Supporting verification of features exclusive to the standard library; and,
- (2) Ensuring a complete verification of safety properties of unsafe code (heavily used in the standard library).

For (1), consider intrinsics (e.g., `transmute`, which reinterprets the bits of one type as another type). These functions have specifications given in natural language in the Rust documentation, but their implementations are directly woven into the code by the Rust compiler.

From a verification point of view, intrinsics appear as external functions, whose verification Kani does not currently support. Under `assume/guarantee` reasoning, to verify these functions, we need to (1) have contracts for them (`assume`) and (2) establish that the functions indeed implement their contracts (`guarantee`). Natural-language contracts for intrinsics exist in the Rust documentation; we will translate them into a machine-readable form, seek community consensus on the correctness of our translations, and make them available to the verification tool. To our knowledge, techniques for verifying the implementation of intrinsics do not exist for any language, and that part of the problem is likely out of scope of our effort; we would be satisfied with formal specifications for intrinsics, agreed-on by the Rust community.

For (2), we need to improve the verification of unsafe code, which is highly prevalent in the standard library. While Kani is one of the few tools that does support the verification of unsafe code, there is still a wide range of undefined behaviour it does not detect (e.g., data races, mutation of immutable data, breaking lifetime-related aliasing rules). As stated in Rust's documentation<sup>4</sup>, any program containing undefined behaviour has no semantics, meaning that guarantees given by Kani may not necessarily hold. Thus, it is clear that extending the range of undefined behaviour that Kani detects is crucial to gain better confidence in Kani's proofs.

As mentioned in Section 1, Amazon has initiated an effort to verify the Rust standard library, and invites community participation in this effort. They recently released a set of challenges to be solved by the Rust verification community<sup>5</sup>, consisting of specific high-priority parts of the standard library that they would like to see verified, including `transmute` and its uses, inductive data types, raw pointer arithmetic operations, and use of raw pointers in intrinsics. We have chosen to start with the `transmute` challenge, as it seems both reasonably achievable and is a high-value target. Indeed, `transmute` is heavily used throughout the standard library, while also being highly prone to causing undefined behaviour, such as creating instances of types with invalid states. As per the *Rustonomicon*<sup>6</sup>, "This is really, truly, the most horribly unsafe thing you can do in Rust". Specifications for the `transmute` function itself would likely include things like checking that the output of `transmute` is valid with respect to the type the input value has been reinterpreted as.

## 6 HUMAN-RELATED OUTCOMES

By this point, we have motivated the verification of unsafe code in Rust's standard library, as well as the use of Kani for this task. Here, we discuss outcomes desired from this verification effort, beyond the inherent value of the verification itself. In particular, we focus

<sup>4</sup>*Rustonomicon*, section 16.2, "Behaviour considered undefined": <https://doc.rust-lang.org/reference/behavior-considered-undefined.html>

<sup>5</sup>*Verify Rust Std Lib*, section 3, "Challenges": <https://model-checking.github.io/verify-rust-std/challenges.html>

<sup>6</sup>*Rustonomicon*, section 4.4, "Transmutes": <https://doc.rust-lang.org/nomicon/transmutes.html>

on the interaction between human behaviour and Rust verification. Part of our proposal includes documenting our verification efforts, reflecting on the resulting learnings about Rust verification, and sharing them as an experience report. We plan on answering questions related to specification languages, developer use of safe and unsafe Rust, as well as Kani and its use.

*Specification languages.* We hope to learn more about how specification languages come into play, as they are a key part of the verification process that humans interact with. For instance, we will investigate the relationship between safety properties in plain text and formalized specification languages, particularly with respect to the completeness of the plain text properties we encounter (i.e., how often are there missing safety properties that need to be added?). Moreover, in the case of the standard library, what kind of specification language is needed to express functional correctness of safe code? In other words, from a human aspects point of view, what improvements to the specification language can be made to bring the verification of safe code closer to full-functional correctness? On the other hand, what does a specification language for safety properties look like, and how specific must it be to safety properties?

*Use of safe/unsafe.* By examining lots of Rust code, we hope to understand the conditions under which developers mark safe code as unsafe, and to what extent. This might happen, for instance, when the developer makes the boundaries of the unsafe block much larger than they need to be, encompassing some safe code in the process. This naturally has ramifications in terms of verification, as verification of safe code can benefit from more assumptions that can be made about the code. After all, that safe code will have passed the rigorous compiler checks for safe code.

On the subject of safety properties, we also seek to understand how often clients fail to respect the explicit safety properties that they are bound to uphold. In other words, how often does API misuse occur, and are there patterns in the nature of the misuse? If some stated safety properties are violated a lot more often by clients than others, then understanding the underlying reason could help in establishing some solution. Perhaps stronger checks should be implemented in the APIs for these properties, or maybe the property is generally poorly communicated in documentation. We fundamentally believe that optimizing the affordances of the specification language—an extremely human aspect of enabling reasoning about the code—will contribute to developers writing code that respects necessary safety properties.

As we are primarily focusing on verification of the standard library, we also seek to understand if there is a meaningful difference in the way unsafe code is used in the standard library as opposed to regular libraries. This is important not only because it sheds some light on a potential threat to generalizability for this project, but also because characterizing these differences (if they exist) can help future verification efforts on the standard library to adjust their approaches accordingly.

*Kani.* As for Kani specifically, by extensively using and eventually extending Kani, we will gain an understanding of which properties can be verified by a model-checking approach such as

Kani, and which ones cannot. We plan to augment its existing documentation with our learnings about its abilities and shortcomings, as well as adding capabilities to Kani as needed.

Finally, it is well-understood that usability (or lack thereof) is a major obstacle for the widespread use of software verification in general. If the tools that verify code are easier to use, then more code will be verified. Following our work, we will have a better understanding of which properties are infeasible to verify via Kani. For each of these difficult properties, we also more generally aim to evaluate the cost of verifying them (i.e., the usability of other approaches that allow their verification), as well as the benefit (i.e., how much the lack of verification of this property results in bugs). After we publicize this information, researchers and practitioners will be able to make more informed decisions about which approaches to use, and which properties are worth verifying.

## 7 CONCLUSION

In this research proposal, we explored the current Rust verification landscape. We first identified two main classes of properties to verify—generic safety and domain-specific properties—and discussed specific properties from these two classes that are worth verifying. Next, we surveyed verification approaches in C/C++ and Java, and then categorized and compared existing Rust techniques. We discussed which parts of the Rust ecosystem are highest-priority for verification, motivating a focus on the standard library. Finally, we motivated the use of Kani for this verification, exploring the challenges inherent to the verification of the standard library and suggesting ways to solve them.

Our exploration has yielded a set of questions and challenges, including:

- verifying sound use of unsafe code in Rust (i.e. do clients respect safety properties required to use the unsafe code);
- how far bounded model checking can go in verifying unsafe Rust; and
- what specifications should be given to extern functions in the standard library.

We look forward to discussing these challenges and their solutions with the community.

## ACKNOWLEDGMENTS

This material is based upon work supported in part by an Amazon Research Award. We thank the HATRA reviewers and Michael Coblenz for detailed reviews and discussions which helped improve this work.

## REFERENCES

- Wolfgang Ahrendt, Bernhard Beckert, Daniel Bruns, Richard Bubel, Christoph Gladisch, Sarah Grebing, Reiner Hähnle, Martin Hentschel, Mihai Herda, Vladimir Klebanov, et al. 2014. The KeY platform for verification and analysis of Java programs. In *Verified Software: Theories, Tools and Experiments: 6th International Conference, VSTTE 2014, Vienna, Austria, July 17-18, 2014, Revised Selected Papers 6*. Springer, 55–71.
- Brian Anderson. 2020. The Rust Compilation Model Calamity. <https://pingcap.medium.com/the-rust-compilation-model-calamity-1a8ce781cf6c>. (January 2020). Accessed on 14 July 2024.
- Andrew W Appel. 2011. Verified Software Toolchain: (Invited Talk). In *European Symposium on Programming*. Springer, 1–17.

- Vytautas Astrauskas, Aurel Bily, Jonáš Fiala, Zachary Grannan, Christoph Matheja, Peter Müller, Federico Poli, and Alexander J Summers. 2022. The Prusti project: Formal verification for Rust. In *NASA Formal Methods Symposium*. Springer, 88–108.
- Vytautas Astrauskas, Christoph Matheja, Federico Poli, Peter Müller, and Alexander J. Summers. 2020. How do programmers use unsafe Rust? *Proc. ACM Program. Lang.* 4, OOPSLA, Article 136 (nov 2020), 27 pages. <https://doi.org/10.1145/3428204>
- Sacha-Élie Ayoun, Xavier Denis, Petar Maksimović, and Philippa Gardner. 2024. A hybrid approach to semi-automated Rust verification. <https://arxiv.org/pdf/2403.15122>. (March 2024).
- Marek Baranowski, Shaobo He, and Zvonimir Rakamarić. 2018. Verifying Rust programs with SMACK. In *Automated Technology for Verification and Analysis: 16th International Symposium, ATVA 2018, Los Angeles, CA, USA, October 7–10, 2018, Proceedings 16*. Springer, 528–535.
- Edmund Clarke, Daniel Kroening, and Flavio Lerda. 2004. A tool for checking ANSI-C programs. In *Tools and Algorithms for the Construction and Analysis of Systems: 10th International Conference, TACAS 2004, Held as Part of the Joint European Conferences on Theory and Practice of Software, ETAPS 2004, Barcelona, Spain, March 29–April 2, 2004. Proceedings 10*. Springer, 168–176.
- Ernie Cohen, Markus Dahlweid, Mark 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: 22nd International Conference, TPHOLs 2009, Munich, Germany, August 17–20, 2009. Proceedings 22*. Springer, 23–42.
- David R Cok. 2011. OpenJML: JML for Java 7 by extending OpenJDK. In *NASA Formal Methods: Third International Symposium, NFM 2011, Pasadena, CA, USA, April 18–20, 2011. Proceedings 3*. Springer, 472–479.
- Byron Cook, Andreas Podelski, and Andrey Rybalchenko. 2006. Termination Proofs for Systems Code. In *PLDI*. 415–426.
- Lucas Cordeiro, Pascal Kesseli, Daniel Kroening, Peter Schrammel, and Marek Trtik. 2018. J BMC: A bounded model checking tool for verifying Java bytecode. In *International Conference on Computer Aided Verification*. Springer, 183–190.
- Patrick Cousot and Radhia Cousot. 2012. An Abstract Interpretation Framework for Termination. In *Proceedings of the 39th Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL '12)*. Philadelphia, PA, USA, 245–258.
- Mohan Cui, Shuran Sun, Hui Xu, and Yangfan Zhou. 2024. Is unsafe an Achilles' Heel? A Comprehensive Study of Safety Requirements in Unsafe Rust Programming. In *Proceedings of the IEEE/ACM 46th International Conference on Software Engineering*. 1–13.
- Leonardo De Moura and Nikolaj Bjørner. 2008. Z3: An efficient SMT solver. In *International conference on Tools and Algorithms for the Construction and Analysis of Systems*. Springer, 337–340.
- Xavier Denis, Jacques-Henri Jourdan, and Claude Marché. 2022. Creusot: a foundry for the deductive verification of Rust programs. In *International Conference on Formal Engineering Methods*. Springer, 90–105.
- Ana Nora Evans, Bradford Campbell, and Mary Lou Soffa. 2020. Is Rust used safely by software developers?. In *Proceedings of the ACM/IEEE 42nd International Conference on Software Engineering*. 246–257.
- Ben Evans. 2020. The Unsafe Class: Unsafe at Any Speed. <https://blogs.oracle.com/javamagazine/post/the-unsafe-class-unsafe-at-any-speed>. (May 2020).
- Lennard Gäher, Michael Sammler, Ralf Jung, Robbert Krebbers, and Derek Dreyer. 2024. Refinedrust: A type system for high-assurance verification of Rust programs. *Proceedings of the ACM on Programming Languages* 8, PLDI (2024), 1115–1139.
- Arie Gurfinkel, Temesghen Kahsai, Anvesh Komuravelli, and Jorge A Navas. 2015. The SeaHorn verification framework. In *International Conference on Computer Aided Verification*. Springer, 343–361.
- Son Ho and Jonathan Protzenko. 2022. Aeneas: Rust verification by functional translation. *Proceedings of the ACM on Programming Languages* 6, ICFP (2022), 711–741.
- Shiyu Huang, Jianmei Guo, Sanhong Li, Xiang Li, Yumin Qi, Kingsum Chow, and Jeff Huang. 2019. SafeCheck: Safety Enhancement of Java Unsafe API. In *2019 IEEE/ACM 41st International Conference on Software Engineering (ICSE)*. 889–899. <https://doi.org/10.1109/ICSE.2019.00095>
- Atsushi Igarashi, Benjamin C Pierce, and Philip Wadler. 2001. Featherweight Java: a minimal core calculus for Java and GJ. *ACM Transactions on Programming Languages and Systems (TOPLAS)* 23, 3 (2001), 396–450.
- Ralf Jung, Jacques-Henri Jourdan, Robbert Krebbers, and Derek Dreyer. 2017. RustBelt: Securing the foundations of the Rust programming language. *Proceedings of the ACM on Programming Languages* 2, POPL (2017), 1–34.
- Steve Klabnik and Carol Nichols. 2024. *The Rust Programming Language*. No Starch Press. Accessed at <https://doc.rust-lang.org/book/> on 14 July 2024.
- Andrea Lattuada, Travis Hance, Chanhee Cho, Matthias Brun, Isitha Subasinghe, Yi Zhou, Jon Howell, Bryan Parno, and Chris Hawblitzel. 2023. Verus: Verifying rust programs using linear ghost types. *Proceedings of the ACM on Programming Languages* 7, OOPSLA1 (2023), 286–315.
- K. Rustan M. Leino. 2023. *Program Proofs*. MIT Press.
- Xavier Leroy. 2009. Formal verification of a realistic compiler. *Commun. ACM* 52, 7 (2009), 107–115. <http://xavierleroy.org/publi/compcert-CACM.pdf>
- Luis Mastrangelo, Luca Ponzanelli, Andrea Mocchi, Michele Lanza, Matthias Hauswirth, and Nathaniel Nystrom. 2015. Use at Your Own Risk: The Java Unsafe API in the Wild. In *OOPSLA 2015: Proceedings of the 2015 ACM SIGPLAN International Conference on Object-Oriented Programming, Systems, Languages, and Applications*. 695–710.
- Yusuke Matsushita, Xavier Denis, Jacques-Henri Jourdan, and Derek Dreyer. 2022. RustHornBelt: a semantic foundation for functional verification of Rust programs with unsafe code. In *Proceedings of the 43rd ACM SIGPLAN International Conference on Programming Language Design and Implementation (PLDI 2022)*. Association for Computing Machinery, New York, NY, USA, 841–856. <https://doi.org/10.1145/3519939.3523704>
- Jiun Min, Dongyeon Yu, Seongyun Jeong, Dokyung Song, and Yuseok Jeon. 2024. ERASAN : Efficient Rust Address Sanitizer. In *2024 IEEE Symposium on Security and Privacy (SP)*. IEEE Computer Society, Los Alamitos, CA, USA, 239–239. <https://doi.org/10.1109/SP54263.2024.00258>
- Alex Ozdemir. 2019. Unsafe in Rust: The Abstraction Safety Contract and Public Escape. <https://cs.stanford.edu/~aozdemir/blog/unsafe-rust-escape>. (2019).
- Boqin Qin, Yilun Chen, Zeming Yu, Linhai Song, and Yiyang Zhang. 2020. Understanding memory and thread safety practices and issues in real-world Rust programs. In *Proceedings of the 41st ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI 2020)*. Association for Computing Machinery, New York, NY, USA, 763–779. <https://doi.org/10.1145/3385412.3386036>
- Rust team. 2016. The Rustonomicon: What Unsafe Means. <https://doc.rust-lang.org/nomicon/safe-unsafe-meaning.html>. (2016). Accessed: 2024-07-12.
- Konstantin Serebryany, Derek Bruening, Alexander Potapenko, and Dmitriy Vyukov. 2012. AddressSanitizer: A fast address sanity checker. In *2012 USENIX annual technical conference (USENIX ATC 12)*. 309–318.
- Mohammad Robati Shirzad and Patrick Lam. 2024. A study of common bug fix patterns in Rust. *Empirical Software Engineering* 29, 2 (February 2024).
- Alexa VanHattum, Daniel Schwartz-Narbonne, Nathan Chong, and Adrian Sampson. 2022. Verifying dynamic trait objects in Rust. In *Proceedings of the 44th International Conference on Software Engineering: Software Engineering in Practice*. 321–330.
- Aaron Weiss, Olek Gierczak, Daniel Patterson, and Amal Ahmed. 2019. Oxide: The essence of Rust. *arXiv preprint arXiv:1903.00982* (2019).
- Hui Xu, Zhuangbin Chen, Mingshen Sun, Yangfan Zhou, and Michael R Lyu. 2021. Memory-safety challenge considered solved? An in-depth study with all Rust CVEs. *ACM Transactions on Software Engineering and Methodology (TOSEM)* 31, 1 (September 2021), 1–25. <https://doi.org/10.1145/3466642>
- Shaowei Zhu and Zachary Kincaid. 2021. Termination Analysis without the Tears. In *PLDI*.