

Extensible Functional-Correctness Verification  
of Rust Programs by the Technique of Prophecy

預言の技術による Rust プログラムの  
拡張可能な機能正当性検証

by

Yusuke Matsushita

松下 祐介

A Master's Thesis

修士論文

Submitted to

the Graduate School of the University of Tokyo

on February 24, 2021

in Partial Fulfillment of the Requirements

for the Degree of Master of Information Science and Technology

in Computer Science

Thesis Supervisor: Naoki Kobayashi 小林 直樹

Professor of Computer Science

## ABSTRACT

Rust is a systems programming language that gives strong static guarantees on properties like memory and thread safety and has been widely used in recent years for high reliability and performance. In Rust, code is usually under a static check of resource usage under the ownership principle based on lifetimes. Also, we can robustly extend the expressivity of Rust by writing libraries with implementation free from the static check and interfaces following the ownership principle. Therefore, verifying safety and functional correctness of Rust programs in an extensible and scalable way is a key to developing reliable high-performance software. There is some existing work on verifying Rust programs. Jung et al. verified safety of Rust's type system and some Rust libraries in an extensible way using the higher-order separation logic Iris in the Coq Proof Assistant. Matsushita et al. proposed a method to automatically verify functional correctness of programs in a basic subset of Rust. However, it was still unclear how to verify functional correctness of Rust programs in an extensible and scalable way. In this thesis, we propose a novel extensible logical foundation to specify and verify functional correctness (and safety) of Rust programs, by extending the technique of prophecy used by Matsushita et al. We semantically model types and verification conditions in Rust in Iris, combining the work of Jung et al. with the technique of prophecy, to flexibly support new libraries and features of Rust. As a basis for that, we present a new formulation of prophecy that allows flexible operation. We plan to mechanize the results in Coq in the near future.

## 論文要旨

Rust はメモリ安全性やスレッド安全性などについて強い静的保証を与えるシステムプログラミング言語であり、高い信頼性と性能を求めて近年幅広く使われている。Rust では、コードは通常ライフタイムに基づく所有権原理のもとでリソースの使用を静的に検査されている。また、静的検査を受けない実装と所有権原理に従うインターフェースを持つライブラリを書くことで、Rust の表現力を堅牢に拡張することもできる。ゆえに、Rust プログラムの安全性および機能正当性を拡張可能かつスケーラブルな形で検証することが信頼できる高性能ソフトウェアを開発するための鍵となる。Rust プログラムの検証についてはいくつか既存研究がある。Jung らは Rust の型システムおよびいくつかの Rust ライブラリの安全性を定理証明支援系 Coq 上で高階分離論理 Iris を用いて拡張可能な形で検証した。松下らは Rust の基本的なサブセットのプログラムの機能正当性を自動検証する手法を提案した。しかしながら、Rust プログラムの機能正当性を検証するための拡張可能でスケーラブルな手法は知られていなかった。この論文では、松下らの用いた預言の技術を拡張し、Rust プログラムの機能正当性 (と安全性) を記述および検証するための、新しい拡張可能な論理的基盤を提案する。Jung らの成果と預言の技術を組み合わせて、Iris 上で Rust の型や検証条件を意味論的にモデル化し、Rust の新しいライブラリや機能を柔軟に取り入れられるようにする。このための基礎として、柔軟な操作を許す新しい預言の概念を提案する。近い将来に成果を Coq で機械化する予定である。

## Acknowledgments

I express my sincere appreciation to my supervisor, Prof. Naoki Kobayashi. Over years he fostered my skills as a researcher, especially in presentation and communication, with great passion and patience. On this research he gave me great advice, too. I would also like to thank Prof. Takeshi Tsukada for his lively and warm advice on this research. Moreover, they both worked with me on RustHorn, an essential predecessor of this work.

This work emerged during my research internship at the RustBelt team of Max Planck Institute for Software Systems from September to December 2020. I would like to show my gratitude to people who gave me warm assistance during the internship. Prof. Derek Dreyer supervised me with patience to cultivate my research skills and guide the research direction; he encouraged me with a fabulous sense of humor and generously allowed me to use the name RustHornBelt for this research project. Jacques-Henri Jourdan gave me valuable comments in discussion and conceived some important ideas for this work; he also mechanized my new idea about consuming logical steps in Iris just a few days after I proposed it, which excited me a lot. Ralf Jung always encouraged me and gave me insightful advice; it was always my pleasure to interact with him. Xavier Denis collaborated closely with me on this work, with great passion for verifying Rust programs; he gave me advice on this thesis and discussed with me the future direction of this work. Rodolphe Lepigre took care of me continuously. Although the global COVID-19 pandemic prevented me from flying to Europe to meet them, I had amazing collaboration online.

Finally, I would like to thank my mother and father for supporting me from various aspects in my life. Thanks to their warm supports, over years, I have been able to enjoy academic activities.

# Contents

<b>1</b>	<b>Introduction</b>	<b>1</b>
1.1	The Rust Programming Language . . . . .	2
1.2	RustBelt — A Semantic Model of Rust Types . . . . .	10
1.3	RustHorn — Verifying Rust Programs by Prophecy . . . . .	11
1.4	Our Work, RustHornBelt — Unifying RustHorn and RustBelt . . . . .	15
<b>2</b>	<b>Preliminaries on the Higher-Order Separation Logic Iris</b>	<b>17</b>
2.1	Basic Connectives and Deduction Rules of Iris . . . . .	17
2.2	Customizing Iris by Resource Algebras . . . . .	25
2.3	Lifetime Logic — Borrowing Propositions Under Lifetimes . . . . .	28
<b>3</b>	<b>A New Formulation of Prophecy</b>	<b>33</b>
3.1	Motivation . . . . .	34
3.2	Overview of Our Formulation of Prophecy . . . . .	36
3.3	Technical Details . . . . .	37
3.4	Related Work . . . . .	41
<b>4</b>	<b>A Low-level Foundation for Verification</b>	<b>44</b>
4.1	A Lambda Calculus for Imperative Programming . . . . .	44
4.2	Verification in Iris . . . . .	47
4.3	Related Work . . . . .	53
<b>5</b>	<b>A New Design of Semantic Types for Rust</b>	<b>55</b>
5.1	Our Notion of the Semantic Type . . . . .	56
5.2	Basic Semantic Types . . . . .	58
5.3	Modeling the Unique Reference Type with Prophecy . . . . .	61
5.4	Recursive Type . . . . .	66
<b>6</b>	<b>A Refined Type System for Verifying Rust Programs</b>	<b>69</b>
6.1	Judgments for the Refined Type System . . . . .	70
6.1.1	Refined Typing Judgment and Its Variants . . . . .	70
6.1.2	Basic Deduction Rules . . . . .	72
6.1.3	Other Judgments . . . . .	76
6.1.4	Function Type . . . . .	79
6.2	Specifying and Verifying Operations on Basic Types . . . . .	79
6.2.1	Value Types . . . . .	80
6.2.2	Basic Pointer Types . . . . .	81
6.2.3	Constructive Types . . . . .	84
6.2.4	Function Type . . . . .	87
6.2.5	Vector Type . . . . .	89
6.3	Specifying and Verifying Operations on Unique References . . . . .	89
6.3.1	Unique Borrow and Reborrow . . . . .	89

6.3.2	Access on Unique References . . . . .	91
6.3.3	Resolution of Unique References . . . . .	92
6.3.4	Subdivision of Unique References . . . . .	94
6.3.5	Subtyping on Unique References . . . . .	97
6.3.6	Manipulating Vectors via Unique References . . . . .	98
6.4	Examples of Verification by the Refined Type System . . . . .	100
6.5	Related Work . . . . .	103
<b>7</b>	<b>Conclusion and Future Work</b>	<b>106</b>

# Chapter 1

## Introduction

The *Rust Programming Language* (Matsakis and Klock, 2014; Rust Community, 2021a) is a systems programming language that aims at a high level of performance, reliability and productivity. Like C/C++, Rust allows users to use low-level operations for high performance. At the same time, unlike common mainstream languages like C/C++ and Java, Rust gives a high-level safety guarantees by a strong static check on resource usage by its distinctive *ownership principle* based on *lifetimes*. Rust also allows unsafe code, i.e., code that does not get the static resource-usage check. We can robustly extend the expressivity of Rust by writing libraries with implementation with unsafe code and interfaces that follow the ownership principle. Rust has been widely used in recent years in industry and has attracted various academic interests. Therefore, *verification of safety and functional correctness of Rust programs* is significant for development of reliable high-performance software. In verification of Rust programs, *scalability and extensibility* is important, because Rust is widely used for large-scale software and the expressivity of Rust can be flexibly extended by libraries with unsafe code.

There are some existing studies on verification of Rust programs. *RustBelt* (Jung et al., 2018a) verified memory and thread *safety* of Rust’s type system and some Rust libraries mechanically in the Coq Proof Assistant. Their proof is highly *extensible* thanks to a semantic approach based on the higher-order separation logic *Iris* (Jung et al., 2015, 2018c). *RustHorn* (Matsushita et al., 2020a) proposed a method to *automatically* verify *functional correctness* of Rust programs within a basic subset of features. Their verification method leverages the guarantees of Rust’s ownership principle to give a clean logic model to a Rust program, which allows a fairly scalable verification of Rust programs. The logic model precisely tracks the effect of updates performed through *unique references* by the technique of *prophecy*, while it successfully omits explicit representation of the the heap memory and addresses. However, it was still unclear how to verify functional correctness of Rust programs in an extensible and scalable way.

In this thesis, we propose a novel extensible logical foundation to specify and verify functional correctness including safety, of Rust programs, using a prophecy-based clean logic model in the style of RustHorn and taking a semantic approach in the style of RustBelt. We name this research project *RustHornBelt*. As a basis of that, we present a new framework of prophecy on top of Iris; in this framework, prophecy lives only in the ghost state and allows flexible operations. We plan to mechanize our results in Coq in the near future, taking advantage of the strong Coq support of Iris.

In § 1.1, we introduce the Rust programming language, especially in terms of the type system. In § 1.2, we introduce RustBelt. In § 1.3, we introduce RustHorn and explain its idea, giving some detailed examples of verification in RustHorn. In § 1.4, we give an overview of our work RustHornBelt.

## 1.1 The Rust Programming Language

The *Rust Programming Language* (Matsakis and Klock, 2014; Klabnik et al., 2018; Rust Community, 2021a) is a systems programming language aiming at a high level of performance, reliability and productivity.

As a fairly new language (Rust 1.0 was released in 2015), Rust has been designed on the basis of the lessons learned from existing programming languages and past development experiences. Like C/C++, Rust allows users to use *low-level operations* to achieve high memory and space efficiency. At the same time, Rust provides *high-level safety guarantees* by a strong static check on *resource usage* by the *ownership* principle based on *lifetimes*, which is introduced later in this section. Also, productivity of Rust is enhanced by various modern features, including pattern matching and lambda expressions, and also mature development tools.

Rust has been developed by a large, energetic community and has rapidly gained popularity. Rust is already used by software products such as Firefox (Mozilla, 2021), Dropbox (Dropbox, 2020) and npm (npm, 2019). Also, Rust is officially sponsored by Microsoft, Mozilla, Amazon and Google (Rust Community, 2021b).

Rust has attracted various academic interests. There have been various studies on verification of Rust programs (Toman et al., 2015; Hahn, 2016; Ullrich, 2016; Jung et al., 2018a; Lindner et al., 2018; Baranowski et al., 2018; Astrauskas et al., 2019; Dang et al., 2020; Matsushita et al., 2020a). Some studies worked on aliasing disciplines for Rust (Jung et al., 2020a), feature proposal on Rust (Fallin, 2020), and formalization of Rust (Reed, 2015; Weiss et al., 2019). Some studies investigated real-world Rust programs (Qin et al., 2020; Astrauskas et al., 2020). Some studies applied Rust to systems programming (Levy et al., 2015; Anderson et al., 2016; Balasubramanian et al., 2017; Levy et al., 2017; Lamowski et al., 2017; Ding et al., 2017; Almohri and Evans, 2018; Emmerich et al., 2019), concurrency (Jespersen et al., 2015), and cryptography (Mindermann et al., 2018).

In the remaining part of this section, we introduce and explain Rust's type system, especially in terms of the ownership principle based on lifetimes.

**Possible Dangers of Pointer Use** To understand the motivation of Rust's type system, let us first see how dangerous pointer manipulation can be without static checking on resource usage. In most mainstream programming languages such as C/C++ and Java, each pointer can read, modify and release its target resource without much restriction. This freedom can easily cause, however, unexpected errors.

A programmer can unintentionally use a pointer to a resource that has been implicitly released.

*Example 1.1* (Dangerous Update After Implicit Release). For example, let us consider the following C++ code.

```
1 vector<int>* pv;
2 { vector<int> v { 0, 1, 2 }; pv = &v; } // v is released
3 /* ... some operations ... */ pv->push_back(3); // dangerous!
```

Here, `vector<int>` is a type for an integer vector, i.e., an array of integers that can change in size. In the line 1, we declare an uninitialized pointer `pv` to an integer vector. In the line 2, a local scope is introduced by curly brackets and within it a new vector `v` is created with initial elements `0`, `1`, `2` and the pointer `pv` is set to `&v`. When we leave the local scope, the vector `v` is automatically released. After some operations, in the line 3, we try to append an element to the target vector of `pv`. Because the target of `pv` has actually been *invalidated*, this memory access is highly *dangerous*. If another

object is using the memory cells at the address `pv`, this memory access violates that object.

An internal memory block of an object can be released by update on the object, which may cause a tricky memory error.

*Example 1.2* (Dangerous Update After Reallocation). Let us consider the following example.

```
1 vector<int> v { 0, 1 }; int* p = &v[0];
2 v.push_back(2); // reallocation occurs
3 /* ... some operations ... */ *p = 10; // dangerous!
```

In the line 1, we create a vector of two elements 0 and 1 and store it to the variable `v`, and then we take a pointer `p` to the first element of the vector. In the line 2, we push a new element 2 to (the end of) the vector through `v`. Actually, at this point, the memory block for the elements of the vector gets *reallocated*.<sup>1</sup> After some operations, in the line 3, we try to update the target of the pointer `p`. However, because the memory cell at the address `p` is no longer managed by `v`, it may be now used by another object, so this memory access is highly *dangerous* and can cause *unexpected behavior*.

To sum up, when a resource is shared by multiple pointers, the programmer can easily cause unexpected behavior by updating the resource from some pointer. In practice, such memory safety errors occur relatively frequently. For example, [Miller \(2019\)](#) reported that about 70% of vulnerabilities in Microsoft in terms of CVE patches were related to memory safety in the years 2006-2018.

**Rust's Ownership Principle** Unlike languages like C/C++ and Java, Rust's type system guarantees *memory safety* by performing a strong static check on *resource usage* under the *ownership principle* based on *lifetimes*. This kind of type system that is aware of resource usage is called *substructural* (in a broad sense). Although substructural type systems with ownership have been actively studied for a long time ([Wadler, 1990](#); [Clarke et al., 1998](#); [DeLine and Fähndrich, 2001](#); [Jim et al., 2002](#); [Fähndrich and DeLine, 2002](#); [Fluet et al., 2006](#); [Mazurak et al., 2010](#); [Haller and Odersky, 2010](#); [Zibin et al., 2010](#); [Tov and Pucella, 2011](#); [Ghica and Smith, 2014](#); [Morris, 2016](#); [Bernardy et al., 2018](#)), the type system of Rust has some notable features. Now, before getting into Rust code examples, we give a high-level overview of Rust's ownership principle.

In Rust, an object can be pointed at by multiple pointers, but at each program point only one pointer can have *ownership* (also called *unique permission*) on the object, which permits writing to and reading from the object. While some pointer has ownership on an object, any other pointers do not have ownership on it. Ownership is structural; that is, ownership on an object includes ownership on any sub-objects of the object.

A key feature of Rust is *borrowing*. When a pointer `a` has ownership to some object, we can create a new pointer `ua` to the object and temporarily move the ownership on the object from `a` to `ua` during some period. The period for borrowing is called a *lifetime*. The newly created pointer `ua` is called a *unique reference*. Also, precisely, this type of borrowing is called a *unique borrow*. The lifetime is statically and globally managed by Rust's type system. While the lifetime is ongoing, the unique reference `ua` has the ownership on the object and the original owner `a` temporarily loses the ownership. At the moment when the lifetime ends, `ua` loses the ownership and `a` retrieves the ownership on the object. The key idea is that the unique reference `ua` and the original owner `a` *don't need to directly communicate with each other when the ownership is*

---

<sup>1</sup> To be precise, whether or not reallocation occurs depends on the capacity of the vector, but in common implementations the capacity of the vector is set exactly to  $n$  when we initialize a vector by an initializer list of the length  $n$ .

*returned*, because the static control based on the lifetime ensures that the ownership is never shared by multiple pointers. So we do not have to strictly keep ua until the lifetime ends. While the lifetime is ongoing, the borrowed ownership on the object can freely be passed around, subdivided, and thrown away. This is a key feature of Rust that realizes *modular* and *scalable* management of the ownership.

Rust also has the notion of *sharing permission*, which can be shared by multiple pointers but permits only reading from the object, not writing to it. Sharing permission is also structural like ownership. While a pointer has sharing permission on an object, any other pointer cannot have ownership on the object. Rust has a *shared borrow*, which is similar to a unique borrow but takes out sharing permission, not ownership (i.e., unique permission). When a pointer a has ownership on some object, we can create a new pointer sa called a *shared reference* and temporarily give it the sharing permission on the object under some lifetime. Because sharing permission is sharable, a shared reference can be *copied*, whereas a unique reference can't be copied. The original owner temporarily loses the ownership while the lifetime is ongoing and retrieves the ownership at the moment the lifetime ends. While the lifetime is ongoing, the borrowed sharing permission can be freely passed around, subdivided, and thrown away, just like a unique borrow.

Notably, a lifetime in Rust is roughly speaking a *set of program points* instead of just a lexical scope, which is important for flexibility of borrowing. This notion of a lifetime is called a *non-lexical lifetime*. The compiler of Rust *automatically infers* non-lexical lifetimes and checks the ownership principle with some elaborate algorithm (Rust Community, 2020). This machinery is dubbed a *borrow checker*. Also, a new algorithm for the borrow checker is under development (Matsakis, 2018, 2020; Rust Community, 2021c).

The ownership principle of Rust can be too restrictive for the purpose of guaranteeing safety. Still, the Rust compiler exploits the guarantees of the ownership principle for *optimization*, which is studied in depth by Jung et al. (2020a). Also, Rust allows users to write *unsafe code*, which is code that does not get the static resource-usage check based on the ownership principle; we can robustly extend the expressivity of Rust by appropriately encapsulating unsafe code. We explain unsafe code more in depth later in this section.

**Examples of Unique Borrows** Let us see how we can use unique borrows. You can try the Rust compiler online at the Rust Playground <https://play.rust-lang.org/>.

For a very simple example, we can use a unique borrow as follows.<sup>2</sup>

```
1 let mut n: i32 = 3;
2 let p: &mut i32 = &mut n;
3 *p = *p + 4;
4 print!("{}", n); // 7
```

The type `i32` represents a (32-bit) integer. In the line 1, we get an integer variable `n` with the initial value 3. The keyword `mut` in `let mut` allows us to modify the data of `n`. The type `&mut T` represents a *unique reference* (also called a mutable reference) to an object typed `T`, which have temporarily acquired the ownership on its target object. In the line 2, by `&mut n`, we perform a *unique borrow* (also called a mutable borrow) to take a unique reference `p` to the integer data of `n`.<sup>3</sup> Then in the line 3, we increment

---

<sup>2</sup> If we want to make it a complete Rust program, we just need to put the code inside the entry point function `fn main() { ... }`.

<sup>3</sup> In Rust, the keyword `mut`, standing for 'mutable', is used for two different contexts. The use of `mut`

the integer data by 4 through the unique reference p. In the line 4, ending the unique borrow, we access the integer data through n and know that the value is now set to 7. In Rust, each reference and borrow is associated with some lifetime but the lifetime is *inferred by the compiler* and the programmer cannot explicitly specify it (except when we handle polymorphism over lifetimes, which is explained later).

Now let us see how Rust's type system prevents the two dangerous pointer usages unregulated in C++ which were discussed above.

*Example 1.3* (Update After Release Prevented by Rust). The first C++ code snippet of [Example 1.1](#) can be translated into the following Rust code.

```
1 let pv: &mut Vec<i32>;
2 { let mut v: Vec<i32> = vec![0, 1, 2]; pv = &mut v; }
3 /* ... some operations ... */ pv.push(3); // dangerous!
```

The type `Vec<T>` represents a vector of the element type T, which corresponds to C++'s `vector<T>`. In the line 2, we first construct a vector by `vec![0, 1, 2]` and let the variable v own it. Later in the line 2, we uniquely borrow the vector `&mut v` to get a unique reference, which is named pv. The Rust compiler correctly judges this Rust code dangerous, emitting an error message like below.

```
error: `v` does not live long enough
* | { ... pv = &mut v; } ...
  |           ^^^^^ - `v` dropped here while still borrowed
  |           borrowed value does not live long enough
* | ... pv.push(3);
  |     -- borrow later used here
```

While the lifetime is ongoing, the alias v temporarily loses the ownership. Because pv is used in the line 3 by the push `pv.push(3)`, Rust judges that the lifetime is alive until the end of the line 3. When the variable v gets out of the lexical scope of curly braces, Rust tries to release (or *drop*) the vector through v but notices that this violates the ownership principle with regard to the lifetime. This is the reason why the Rust compiler emit the error message like shown above. Note that we remove the curly brackets of the line 2 as follows, the code passes the type check.

```
1 let pv: &mut Vec<i32>;
2 let mut v: Vec<i32> = vec![0, 1, 2]; pv = &mut v;
3 pv.push(3);
```

*Example 1.4* (Update After Reallocation Prevented by Rust). The second C++ code snippet of [Example 1.2](#) is translated into the following Rust code.

```
1 let mut v: Vec<i32> = vec![0, 1]; let p: &mut i32 = &mut v[0];
2 v.push(2); // reallocation occurs
3 /* ... some operations ... */ *p = 10; // dangerous!
```

Here, the operation `&mut v[0]` is syntax sugar of `Vec::index_mut(&mut v, 0)`<sup>4</sup> and `v.push(2)` is syntax sugar of `Vec::push(&mut v, 2)`. Again the Rust compiler correctly judges this program dangerous and emits an error message like below.

```
error: cannot borrow `v` as mutable more than once at a time
* | ... let p: &mut i32 = &mut v[0];
  |                               - first mutable borrow occurs here
* | v.push(2); ...
  |   ^ second mutable borrow occurs here
* | ... *p = 10; ...
  |     ----- first borrow later used here
```

in `let mut` allows the object directly owned by the variable to be modified. In the meanwhile, the use of `mut` in `&mut` indicates a unique borrow/reference rather than a shared borrow/reference.

<sup>4</sup> We need a declaration such as `use std::ops::IndexMut`; to use the name `index_mut`.

In the line 1, the operation `&mut v[0]`, which stands for `Vec::index_mut(&mut v, 0)`, uniquely borrows `v` to take a unique reference and then *subdivide* it into a unique reference to the head element of the vector by calling the function `Vec::index_mut`. The resulting unique reference, typed `&mut i32`, is passed to the variable `p`. Since we use `p` in the line 3 by `*p = 10`, the lifetime of this unique borrow should be continued until the line 3. However, in the line 2, the operation `v.push(2)`, which stands for `Vec::push(&mut v, 2)`, uniquely borrows `v` and performs the push operation by calling the function `Vec::push`. To perform this second unique borrow, we need the ownership of `v` here, but it should be temporarily lost until the lifetime of the first borrow ends. This is why the compiler emits the error message like above. If we swap the order of the updates like shown below, the Rust code becomes safe and passes the static check of Rust.

```
1 let mut v: Vec<i32> = vec![0, 1]; let p: &mut i32 = &mut v[0];
2 *p = 10; // the lifetime of p ends here
3 v.push(2); // reallocation occurs but it's ok
```

Here, the Rust compiler infers that the lifetime of `p` or the unique borrow of `&mut v[0]` ends at the line 2. Note that the Rust code passes the type check even if we change `&mut v[0]` into `&mut v[10]`; the method `index_mut` performs dynamic bounds checking and we simply get a dynamic error by this change.

In Rust, we can also create a new reference by borrowing from an existing unique reference. This operation is called a *reborrow*.<sup>5</sup>

*Example 1.5* (Reborrowing). For example, let us consider the following Rust code.

```
1 let mut v = vec![0]; let pv = &mut v;
2 pv.push(1); pv.push(2);
```

It passes Rust's type check and works correctly. It seems fairly natural but actually is made possible by reborrowing. The push operation `pv.push(1)` actually calls the function `Vec::push`, which *consumes* the input unique reference. In order to perform the second push `pv.push(2)` after the first push `pv.push(1)`, we cannot consume `pv` by the first push. Therefore, Rust implicitly performs *reborrowing*. For each push on `pv`, Rust takes a new unique reference `&mut *pv` that lives only during the function call of `Vec::push`. We can rewrite the Rust code above into the following code with explicit reborrowing.<sup>6</sup>

```
1 let mut v = vec![0]; let pv = &mut v;
2 Vec::push(&mut *pv, 1); Vec::push(&mut *pv, 2);
```

**Examples of Shared Borrows** Let us see how we can use shared borrows.

We can copy shared references and use multiple shared references at the same time. For a simple example, the following Rust code is valid.

```
1 let n: i32 = 7;
2 let p: &i32 = &n; let q: &i32 = p;
3 print!("{}", and, *p); print!("{}", *q); // 7 and 7
```

<sup>5</sup> We do not need to think about reborrowing from a shared reference because shared references can be copied.

<sup>6</sup> We can also write `Vec::push(pv, 1)`; instead of `Vec::push(&mut *pv, 1)`; but still Rust implicitly performs reborrowing of `pv`.

The type `&T` (without the keyword `mut`) represents a *shared reference* (also called an immutable reference) to an object typed `T`, which have temporarily acquired *sharing permission* on the target object. In the line 2, we first perform a *shared borrow* (also called an immutable borrow) `&n` to get a shared reference `p` to the integer data of `n`; then we *copy* the address of `p` to get a shared reference `q`. In the line 3, we read the integer value from `p` and then from `q`. This makes the lifetimes of `p` and `q` *overlap* but it is still *valid* because sharing permission can be shared. At a low level (without optimization), unique and shared borrows/references are the same operation/object, but for simplicity of the type system Rust expects that programmers *explicitly annotate* whether the borrow/reference is unique or shared by the notations `&mut` and `&`.

Rust forbids release of an object while we are using a shared reference to the object.

*Example 1.6* (Read After Release Prevented by Rust). For example, the following Rust code is correctly rejected by the Rust compiler.

```
1 let pv: &Vec<i32>;
2 { let v: Vec<i32> = vec![0, 1, 2]; pv = &v; }
3 /* ... */ print!("{}", pv[0]); // dangerous!
```

The type `Vec<i32>` represents a vector of shared references to integer data. In the line 2, we construct the vector `v` and then take a shared reference `pv` to the vector `v`. When we leave from the inner scope, the vector `v` is automatically released. In the line 3, by `pv[0]`, we try to access the head element of the vector through `pv`, where `pv[0]` is syntax sugar for `*Vec::index(&pv, 0)`.<sup>7</sup> However, it is dangerous because it can cause invalid memory access, including *dereference of a null pointer*. When we perform `pv[0]`, we first read the address data of the vector and then dereference it; however, since the vector has already been released, the memory cell that used to store the address now may have an invalid address, possibly the null value. This Rust code does violate Rust's ownership principle and the compiler emits an error message like the following.

```
error: `v` does not live long enough
* | { ... pv = &v; } ...
  |             ^ - `v` dropped here while still borrowed
  |             borrowed value does not live long enough
* | ... print!("{}", pv[0]); ...
  |                   -- borrow later used here
```

Rust also forbids update of an object while we are using a shared reference to the object.

*Example 1.7* (Read After Reallocation Prevented by Rust). For example, the following Rust code is correctly rejected by the Rust compiler.

```
1 let mut v: Vec<&i32> = vec! [&0, &1]; let p: &&i32 = &v[0];
2 v.push(&2); // reallocation occurs
3 /* ... */ print!("{}", **p); // dangerous!
```

In the line 1, we first crate a vector `v` of shared references to integer data, of the type `Vec<&i32>` and the initial value `vec! [&0, &1]`, and then take a shared reference `p` to the head element of the vector, which itself is a shared reference to integer data. Here, by performing `&0`, we statically allocate integer data `0` and take an address of it. In the line 2, we push a new element `&2` to the vector `v`, which causes reallocation. In the line 3, after some operations, we try to access the inner integer data of `**p`. However it is dangerous, because the data at `p` is no longer managed by the vector and thus `*p` can be an invalid address, possibly the null value. This code violates Rust's ownership principle and the compiler emits an error message like the following.

<sup>7</sup> We need declaration like `use std::ops::Index`; to use the method `index`, just like `index_mut`.

```

error: cannot borrow `v` as mutable because it is also borrowed as
↳ immutable
* |   ... let p: &&i32 = &v[0];
  |                                     - immutable borrow occurs here
* |   v.push(&2); ...
  |   ~~~~~~ mutable borrow occurs here
* |   ... print!("{}", **p); ...
  |                                     --- immutable borrow later used here

```

In Rust, a unique reference under a shared reference can only have the sharing permission instead of the ownership (unique permission).

*Example 1.8* (A Unique Reference Under a Shared Reference). For example, the following Rust code is correctly rejected by the Rust compiler.

```

1 let mut v: Vec<&i32> = vec! [&0, &1];
2 let pv: &mut Vec<&i32> = &mut v;
3 let ppv: & &mut Vec<&i32> = &pv; let q: &&i32 = &ppv[0];
4 ppv.push(2); // reallocation occurs
5 /* ... */ print!("{}", **q); // dangerous!

```

In the line 1, we create a vector `v` of shared references to integer data. Then in the line 2, we take a unique reference `pv` to the vector. In the line 3, we take a shared reference `ppv` to the unique reference `pv` and then take a shared reference `qv` to the head element of the vector using `ppv`. In the line 4, we perform the push operation `ppv.push(2)` through `ppv`, which causes reallocation. In the line 5, after some operation, we try to read the inner integer data of `q`. However, it is dangerous because the address at `*q` may be now invalid due to the reallocation. This code violates Rust's ownership principle and the compiler emits an error message like the following.

```

error: cannot borrow `**ppv` as mutable, as it is behind a `&`
↳ reference
* |   let ppv: ... = &pv; ...
  |               --- help: consider changing this to be a mutable
  |               ↳ reference: `&mut pv`
* |   ppv.push(2); ...
  |   ^^^ `ppv` is a `&` reference, so the data it refers to cannot be
  |   ↳ borrowed as mutable

```

**Polymorphism over Lifetimes** In Rust, a function can be *polymorphic over lifetimes*, which is an important feature for modularity of borrowing.

For example, the method `Vec::push` used above has the following function signature.<sup>8</sup>

```

1 fn Vec::push<'a, T>(pv: &'a mut Vec<T>, a: T) -> ()

```

The keyword `fn` stands for 'function'. The method `Vec::push` inputs a unique reference to a vector `pv: &'a mut Vec<T>` and a value `a: T` and appends the value to the vector, possibly performing reallocation. The output type is the unit type `()`. The function is polymorphic over the *lifetime* `'a` of the reference `pv`, as well as the element type `T`.

For another example, the method `Vec::index_mut`, which is used for index access on a vector `&mut pv[i]`, has the following signature.

```

1 fn Vec::index_mut<'a, T>(pv: &'a mut Vec<T>, i: usize) -> &'a mut T

```

<sup>8</sup> Although Rust uses a special notation for methods, we use here a notation for usual functions.

This method outputs the unique reference to the *i*-th element of the vector of the type `&'a mut T` by *subdividing* the input unique reference to a vector `pv: &'a mut Vec<T>`. The type `usize` represents an unsigned 64-bit integer (if we are using the 64-bit architecture). We subdivide the input reference instead of performing a new borrow. Therefore, the lifetime of the output reference `&'a mut T` has the same lifetime `'a` as the input reference `&'a mut Vec<T>`.

**Variant Types** Rust supports *variant types*, unlike common imperative languages like C/C++ and Java.

For example, the following type `Result<T, E>` is commonly used.

```
1 enum Result<T, E> { Ok(T), Err(E) }
```

This type is parametrized over the main result type `T` and the error type `E`. A value of the type `Result<T, E>` is either a successful value `Ok(v)` with a body value `v: T` or an error value `Err(e)` with an error value `e: E`.

We can also make recursive types using variant types. For example, the singly linked list type can be defined as follows.

```
1 enum List<T> { Nil, Cons(T, Box<List<T>>) }
```

Here, the *box pointer type* `Box<U>` represents a pointer to an object typed `T` with the ownership on the object and the right to free (or deallocate) the memory block for the object. A value of the type `List<T>` is either the nil `Nil` or the cons value `Cons(v, l)` of the head value `v: T` and the boxed tail list `l: Box<List<T>>`. Using recursion, we can write various functions over the list type. In particular, we can write interesting functions by combining lists with unique references.

For example, we can write the following function `inc_all`.<sup>9</sup>

```
1 fn inc_all<'a>(pla: &'a mut List<i32>) -> () {
2   match pla {
3     Nil => {},
4     Cons(pa, pla2) => { *pa += 1; inc_all(pla2); }
5   }
6 }
```

It inputs a unique reference to a list `pla` and destructively increments all elements of the list. The pattern `Cons(pa, pla2)` *splits* the unique reference `pla: &'a mut List<i32>` into a unique reference to the head `pa: &'a mut i32` and a unique reference to the tail `pla2: &'a mut List<i32>`.

For another example, we can write the following function `split_unq`.

```
1 fn split_unq<'a>(pla: &'a mut List<i32>) -> List<&'a mut i32> {
2   match pla {
3     Nil => Nil,
4     Cons(px, plx2) => Cons(px, Box::new(split_unq(plx2)))
5   }
6 }
```

It inputs a unique reference to an integer list `pla: &'a mut List<i32>` and outputs a list of unique references to each element of the integer list.

---

<sup>9</sup> We need a declaration like `use List::*`; to use the constructors `Nil` and `Cons` (without the qualifier `List::`).

**Unsafe Code** In Rust, we can also write *unsafe code* (Klabnik et al., 2018, §19.1), which is code that does not take the static check on resource usage. If we want to make a piece of code unsafe, we can just put it into the unsafe code block `unsafe { ... }`. In unsafe code, we can use the *mutable raw pointer type* `*mut T` and the *immutable raw pointer type* `*const T`, which represent a pointer to an object typed `T` with or without the ability to update its target object. In unsafe code, we can actively use any number of mutable and immutable raw pointers to one object *without restriction*. Also, in unsafe code, we can freely perform conversion between (unique or shared) references and raw pointers.

We can robustly extend the expressivity of Rust using libraries that have *internal implementation with unsafe code* and *interfaces that follow the ownership principle*. For example, the internal implementation of the vector type `Vec<T>` depends on unsafe code but we can use the vector type with *safety guarantees* as long as we access vectors through the safe interfaces like `Vec::push` and `Vec::index_mut`. Although Rust's type system is quite restrictive by the strict enforcement of the ownership principle, we can robustly extend the expressivity of Rust in this way.

**Interior Mutability** Using unsafe code, we can write a type whose internal data structure can be modified through *shared references*. Such a property on a type is called *interior mutability* (Klabnik et al., 2018, §15.5) (Jung, 2020, §8.6). Rust provides a bunch of robust libraries that offer interior mutability.

For a simple example, Rust has the cell type `Cell<T>`. It can be used as follows.

```
1 let a: &Cell<i32> = & Cell::new(3); let b: &Cell<i32> = a;
2 a.set(7); print!("{}", b.get()); // 7
```

Here, we make a new integer cell with the initial value 3 by `Cell::new(3)`. We take a shared reference `a` to the cell, which can be *copied* to `b`. By `a.set(7)`, which is syntax sugar for `Cell::set(a, 7)`, we can set the internal value of the cell to 7. By `b.get()`, which is syntax sugar for `Cell::get(b)`, we can get the internal value of the cell, which is now set to 7. The methods `Cell::set` and `Cell::get` have the following function signatures.

```
1 fn Cell::set<'a, T>(c: &'a Cell<T>, d: T) -> ()
2 fn Cell::get<'a, T>(c: &'a Cell<T>) -> T
```

Any shared reference to a cell has the right to call `set` and `get`. Actually, the call of `Cell::get` requires that the type `T` is *copyable*, i.e., an object typed `T` can be copied. Types like the integer type `i32` and the shared reference type `&T` are copyable but types like the vector type `Vec<T>` and the mutable reference type `&mut T` are not copyable. Rust manages a property like copyability on a type using the machinery called a *trait* (Klabnik et al., 2018, §10.2), which is comparable to Java's interface and Haskell's type class.

For an advanced example, Rust also supports the mutex type `Mutex<T>`. A mutex can be shared by multiple threads and be modified by exclusively taking the lock.

## 1.2 RustBelt – A Semantic Model of Rust Types

As explained in §1.1, Rust guarantees memory and thread safety using a static resource-usage check by its distinctive lifetime-based ownership principle. In fact, many common errors in imperative programming (e.g., Example 1.1 and Example 1.2) are prevented by Rust's type system. Also, the expressivity of Rust can be robustly extended using libraries have internal implementation with unsafe code and interfaces that follow the ownership principle.

The Rust community has been making a lot of efforts to enhance and guarantee the safety of Rust’s type system and libraries through tests, bug reports and human thoughts. Still, safety of Rust is highly tricky. In particular, when we have *multiple* libraries with unsafe internal implementation, it is challenging to guarantee that *any combination of method calls* on these libraries does not cause unexpected behaviors. In fact, Rust has suffered from a number of bugs caused by interaction of multiple libraries (Ben-Yehuda, 2015; Borowski, 2019).

Therefore, it is significant to give *formal safety guarantees* to Rust’s type system and libraries in an *extensible* way, being open to new libraries and features of Rust. This need gave rise to the work *RustBelt* (Jung et al., 2018a,b). They gave the first *mechanized* proof of *memory and thread safety* of a core of Rust’s type system and some key standard libraries of Rust, including ones with interior mutability such as `Cell`, `RefCell`, `Mutex` and `RwLock`. They found a tricky bug of a real-world Rust library (Jung, 2017) in the process of the formal verification. Their proof is mechanized in the Coq Proof Assistant (Coq Community, 2021), making full use of the higher-order separation logic *Iris* (Jung et al., 2015, 2018c), which is provided as a library in Coq (Krebbers et al., 2017, 2018). *Iris* is a general-purpose separation logic that attains outstanding extensibility by *impredicativity* or higher-order ghost states (*Iris* is explained more in depth in Chapter 2).

The core of their proof is construction of *semantic models of Rust types*, which consist of *predicates in Iris*. Each syntactic typing rule is semantically interpreted in *Iris* and gets a separate soundness proof. Also, each library method with unsafe implementation and typed interfaces is likewise semantically interpreted in *Iris* and separately verified. In this way, the safety proof of *RustBelt* is highly *extensible*. This semantic approach stands in contrast to common *syntactic* type soundness proof based on progress and preservation (Wright and Felleisen, 1994); although the syntactic approach seems handy, it is in general non-trivial how we can modify the syntactic soundness proof when we add new execution rules and typing rules.

As an essential basis for their semantic approach, they designed the *lifetime logic* on top of *Iris* in order to flexibly discuss borrowing of resources based on lifetimes. In the lifetime logic, we can directly borrow *Iris* propositions instead of some fixed set of resources, which is made possible by the impredicativity of *Iris*.

### 1.3 RustHorn – Verifying Rust Programs by Prophecy

How can we verify *functional correctness* of Rust programs leveraging the guarantees given by Rust’s ownership principle? Let us first focus on a basic subset of Rust without unsafe code. *Without unique references*, we can model Rust programs just as functional programs. But what about *with unique references*? When we perform a unique borrow and create a unique reference in Rust, in the verification model we can pass the current value of the target object to the unique reference. When we update the target object of the unique reference, in the model we can just accordingly update the current target value of the unique reference. The problem is, after the lifetime of the borrow ends, how we can *tell* the original owner the *new value* of the object, reflecting updates that have been performed by the unique reference.

*RustHorn* (Matsushita et al., 2020a) proposed a very simple solution to this question: we use a *prophecy variable* on the value at the end of the lifetime of the unique borrow.<sup>10</sup> Roughly speaking a *prophecy variable* is a variable that peeks out some value that will

---

<sup>10</sup> The idea of using prophecy for verifying Rust programs was first proposed by the senior thesis of the author (Matsushita, 2019). The author completed the *RustHorn* paper (Matsushita et al., 2020a) in collaboration with Takeshi Tsukada and Naoki Kobayashi, the supervisor of this thesis.

be observed in the *future*. More specifically, the verification model handles a unique borrow as follows. First we have some object  $a: T$ , which is modeled as a value  $v$ . When we create a unique reference  $ua: \&mut T$  by uniquely borrowing  $a$ , we create a *prophecy variable*  $x$ , which prophesy the value of the borrowed object at the end of the lifetime of the unique borrow. We model the unique reference  $ua$  as  $(v, x)$ , a pair of the current value  $v$  of the target object and the prophecy variable  $x$ . Also, until the lifetime of the borrow ends, we model  $a$  as the prophecy variable  $x$ . When we update the target object of  $ua$ , in the model we update the current target value  $v$  of the pair  $(v, x)$  into the new target value  $v'$ , obtaining the model  $(v', x)$ . And importantly, at the point we throw away (or release) the unique reference  $ua$  modeled  $(v, x)$ , we *resolve* the prophecy variable  $x$  into the target value  $v$  (i.e., assign  $v$  to  $x$ ), obtaining the equality information  $x = v$  as a postcondition. This method also supports advanced usage of unique references like *reborrowing* and *subdivision* (see [Example 1.10](#)). After the lifetime of the borrow ends, the original owner  $a$  retrieves the ownership, now being able to read and update its object. We know that  $a$  has the value  $x$ , because  $x$  has been resolved to the value  $v$ . In this way, we can translate any Rust program (within a basic subset of Rust) into a logic model *without* explicit representation of the heap memory and addresses, just like programs of functional programming languages. In a sense, this prophecy-based technique gives us a general way to view a basic subset of Rust as a functional programming language.

RustHorn ([Matsushita et al., 2020a](#)) demonstrated this prophecy-based method particularly in the context of automated verification based on *constrained Horn clauses* (CHCs) ([Grebenshchikov et al., 2012](#); [Bjørner et al., 2015](#)). Roughly speaking, a CHC representation of a program is a set of logical formulas of some format called CHCs, which describe constraints on *predicate variables* that represent the input-output relation of some function or continuation. The program verification problem is reduced to the satisfiability problem on CHCs, which can be automatically solved by existing CHC solvers ([Komuravelli et al., 2014](#); [Fedyukovich et al., 2017](#); [Hojjat and Rümmer, 2018](#); [Champion et al., 2018](#)) in a highly *scalable* way. RustHorn ([Matsushita et al., 2020a](#)) implemented an automated verifier that translates Rust programs into CHCs using the idea of prophecy for unique borrows and then calls the backend CHC solvers Spacer ([Komuravelli et al., 2014](#)) and HoIce ([Champion et al., 2018](#)). They also performed experiments on benchmarks using the verifier, including comparison with SeaHorn ([Gurfinkel et al., 2015](#)), an existing CHC-based verification platform for the C Programming Language. RustHorn succeeded in automated verification of fairly complicated Rust programs.

In principle, the prophecy-based translation of RustHorn can also be applied to many Rust libraries including vectors `Vec`. Still it cannot handle libraries with *interior mutability* in a clean way in general, because mutable states are actually *shared* under interior mutability, unlike ownership by unique references.

Below we show two examples of verification in RustHorn with detailed explanation. *Example 1.9* (Dynamic Decision of the Address of a Unique Reference). For example, let us consider the following Rust code.

```

1 fn take_max<'a>(ua: &'a mut int, ub: &'a mut int)-> &'a mut int {
2   if *ua >= *ub { ua } else { ub }
3 }
4 fn test(mut a: int, mut b: int) {
5   let uc: &mut int = take_max(&mut a, &mut b); *uc += 1;
6   assert!(a != b);
7 }

```

For simplicity of verification, we assume here that we have an unbounded integer type

**int.** The function `take_max` takes two unique references to integer `ua` and `ub` and returns the one with the greater target value. This function is quite simple, but the point is that the *address* of the unique reference returned by it is determined by a *dynamic* condition `*ua >= *ub`. The function `test` performs a simple test of `take_max`. It inputs two integer objects `a` and `b`. It uniquely borrows the two and passes them to `take_max` to take a unique reference `uc`. Then it increments the target value of `uc`. Finally it checks that the values of `a` and `b` are different. We want to verify that the assertion of `test` always succeeds, regardless of the inputs. That property holds because the increment on `uc` makes the difference between `a` and `b` exactly one bigger. However, it is tricky because we need to ensure that `uc` points to the greater of `a` and `b`, not only either `a` or `b`.

RustHorn translates the Rust program (or verification problem) above into the following two CHCs.

$$\begin{aligned} & \text{take\_max}((a, a_o), (b, b_o), r) \iff \\ & (a \geq b \wedge b_o = b \wedge r = (a, a_o)) \vee (a < b \wedge a_o = a \wedge r = (b, b_o)) \\ a_o \neq b_o & \iff \text{take\_max}((a, a_o), (b, b_o), (c, c_o)) \wedge c_o = c + 1 \end{aligned}$$

Each CHC is universally quantified over the free variables. The predicate logic for CHCs is implicitly multi-sorted and each variable has some sort.

The first CHC characterizes the function `take_max` through the predicate variable `take_max`, which represents the input-output relation of `take_max`. The unique reference `ua` is modeled as  $(a, a_o)$ , the current target integer  $a$  and the *prophecy* target integer  $a_o$ . We do similarly for `ub`. Suppose  $a \geq b$  holds. We *release* the unique reference  $(b, b_o)$  and thus *resolve* the prophecy variable  $b_o$  to  $b$ , obtaining the equality  $b_o = b$ . We also set the return value  $r$  to  $(a, a_o)$ . When  $a < b$  holds, we do a similar thing swapping  $(a, a_o)$  and  $(b, b_o)$ .

The second CHC describes the function `test`. The inputs `a` and `b` correspond to the variables  $a$  and  $b$ . When we perform a unique borrow `&mut a`, we take a new *prophecy variable*  $a_o$ , which represents the value of `a` just after the end of the borrow, and model the created unique reference as  $(a, a_o)$  and `a` now as  $a_o$ . We do similarly for `&mut b`. The call of `take_max` is modeled as the use of the predicate variable `take_max`. The return value `uc` is modeled again as  $(c, c_o)$ . We increment the target integer of `uc`, which updates the model of `uc` into  $(c + 1, c_o)$ . Then we *release* `uc` and thus *resolve* the prophecy variable  $c_o$  into  $c + 1$ , obtaining the equality  $c_o = c + 1$ . Finally, the assertion `assert!(a != b)` is modeled as the condition  $a_o \neq b_o$ . The variables  $a_o$  and  $b_o$  perfectly reflect the values of `a` and `b` at this point.

The system of two CHCs can be satisfied by the following solution.

$$\text{take\_max}((a, a_o), (b, b_o), (c, c_o)) := c_o = c + 1 \Rightarrow a_o \neq b_o$$

In the experiment of RustHorn (Matsushita et al., 2020a, §4), this problem was automatically and instantly solved.

We can also represent the Rust program as the following functional program (we use OCaml here).

```
let take_max (a, a') (b, b') = if a >= b
  then ( assume (b' = b); (a, a') )
  else ( assume (a' = a); (b, b') )
let test a b =
  let a' = Random.int(0) in let b' = Random.int(0) in
  let (c, c') = take_max (a, a') (b, b') in
```

```

assume (c' = c + 1);
assert (a <> b)

```

Here, the function `assume` is defined as follows.

```

let assume b = if b then () else assume b

```

This function falls into an infinite loop when `b` is false. Since we are solving the *safety problem*, where non-termination is considered success, `assume b` simply introduces the postcondition that `b` is true.

Combining the technique of prophecy with recursive types, we can perform interesting verification.

*Example 1.10* (Update of a List via a Unique Reference). Let us consider the following recursive function on the list type `List<T>` (defined in §1.1).

```

1 fn take_some<'a>(ula: &'a mut List<int>) -> &'a mut int {
2   match ula {
3     Nil => take_some(ula),
4     Cons(ua, ula2) => if rand() { ua } else { take_some(ula2) }
5   }
6 }

```

The function `take_some` inputs a unique reference to an integer list `ula` and outputs a unique reference to some non-deterministically chosen element of the list. Here, we assume some non-deterministic function `fn rand() -> bool`. When the target list of `ula` is `cons`, we split `ula` into unique references on the head `ua: &'a mut int` and the tail `ula2: &'a mut List<int>` and either return `ua` or recursively call `take_some` on `ula2`. When the target list is `nil`, we fall into an infinite loop. We can write the following test function on `take_some`.

```

1 fn test(mut la: List<int>) {
2   let n = sum(&la); let ua = take_some(&mut la);
3   *ua += 1; assert!(sum(&la) == n + 1);
4 }

```

The function `test` first saves the initial sum of the list `la`, then uniquely borrows `la` and passes it to `take_some` to get a unique reference `ua`, increments the target of `ua`, and finally asserts that the sum of the list has increased by one. Here, the function `sum` is defined as follows.

```

1 fn sum<'a>(sla: &'a List<int>) -> int {
2   match sla { Nil => 0, Cons(sa, sla2) => *sa + sum(sla2) }
3 }

```

RustHorn translates this Rust program into the following CHCs.

$$\begin{aligned}
&take\_some((nil, la_o), r) \iff take\_some((nil, la_o), r) \\
&take\_some((a :: la', la_o), r) \iff la_o = a_o :: la'_o \wedge \\
&\quad ((la'_o = la' \wedge r = (a, a_o)) \vee (a'_o = a' \wedge take\_some((la', la'_o), r))) \\
&sum(nil, r) \iff r = 0 \\
&sum(a :: la, r) \iff sum(la, r') \wedge r = a + r' \\
&n' = n + 1 \iff sum(la, n) \wedge take\_some((la, la_o), (a, a_o)) \wedge \\
&\quad a_o = a + 1 \wedge sum(la_o, n')
\end{aligned}$$

The first and second CHCs characterize the function `take_some`, respectively for the cases where the target list of `u1a` is `nil` or `cons`. The first CHC is a tautology. The second CHC is interesting. Because we destruct `u1a` into `Cons(ua, u1a2)`, we accordingly *subdivide* the prophecy variable  $la_o$  into two prophecy variables  $a_o, la_o$ , constraining  $la_o$  to  $a_o :: la'_o$ . Now `ua` is modeled as  $(a, a_o)$  and `u1a2` is modeled as  $(la, la_o)$ . In general, when we subdivide a unique reference, we also accordingly subdivide its prophecy variable, which is an essential technique in this prophecy-based approach. Now what we do is simple. Depending on the non-deterministic boolean `rand()`, we either (i) release `u1a2`, obtaining  $la'_o = la'$ , and return `ua` or (ii) release `ua`, obtaining  $a'_o = a'$ , and recursively call `take_some(u1a2)`. The third and fourth CHCs characterize the function `sum`. A shared reference is represented simply as the value of its target object. The fifth CHC characterizes the test function. This system of CHCs has the following simple solution.

$$\begin{aligned} take\_some((la, la_o), (a, a_o)) &:= a_o = a + 1 \Rightarrow \text{sum } la_o = \text{sum } la + 1 \\ sum(la, r) &:= r = \text{sum } la \end{aligned}$$

Here, the function  $\text{sum} : \text{List } \mathbb{Z} \rightarrow \mathbb{Z}$  is defined inductively as follows.

$$\text{sum nil} := 0 \qquad \text{sum}(a :: la) := a + \text{sum } la$$

In the experiment (Matsushita et al., 2020a, §4), the verifier of RustHorn succeeded in *automated verification* of this Rust program. This is because HoIce (Champion et al., 2018), one of the CHC backend solvers, can find and handle an inductive function like `sum` from the CHCs of `sum`. Once we have `sum`, the solution to `take_some` presented above can be easily found from the last CHC. The validity of this solution can be easily checked using the equality  $\text{sum}(a :: la) = a + \text{sum } la$ .

As these examples show, using the idea of prophecy, we can give a clean logic model to a fairly wide class of Rust programs, which makes verification highly *scalable*. Although the paper of RustHorn focused on CHC-based automated verification, by the technique of prophecy, we can also verify Rust programs using existing semi-automated verification platforms for functional programming languages such as F\* (Swamy et al., 2016) and Why3 (Filliâtre and Paskevich, 2013).

The correctness of this prophecy-based method is non-trivial. The RustHorn paper (Matsushita et al., 2020a,b) presented the proof of *soundness and completeness* of the reduction from Rust programs to CHCs. It was achieved mainly by constructing a bisimulation between the execution of Rust programs and some deduction algorithm on CHCs. On the side of CHCs, prophecy variables of unique references are represented simply as syntactic variables, which can be specialized later in the deduction algorithm.

## 1.4 Our Work, RustHornBelt – Unifying RustHorn and RustBelt

As introduced in §1.2, *RustBelt* (Jung et al., 2018a) verified memory and thread *safety* of the core type and some basic libraries of Rust in an *extensible* way. Still, RustBelt did not verified functional correctness. In particular, RustBelt models unique reference as a kind of invariant saying that *some* object of the type is stored at the target address and does not track precise information about the object. As introduced in §1.3, *RustHorn* (Matsushita et al., 2020a) presented a new method of verifying *functional correctness* of a wide class of Rust programs in a fairly scalable way leveraging the guarantees of Rust’s ownership principle. The point is that for each unique borrow we can take a *prophecy variable* on the value of the target object just after the end of the borrow and model the unique reference as a pair of the current target value and the prophecy variable

to track the final target value when the unique reference is released. Although this prophecy-based modeling is expected to be applicable to a wide class of Rust libraries including `Vec` and `HashMap`, we did not have a logical foundation for proving that in an *extensible* way. Although Matsushita et al. (2020a,b) presented a proof of soundness and completeness of their prophecy-based reduction, their proof depends heavily on syntactic structures and has rather little extensibility (this aspect is discussed more in depth in §3.1 and §6.5). It was unclear how we can unify the approaches of RustHorn and RustBelt.

In this thesis, we propose a novel *extensible* logical foundation to specify and verify functional correctness of Rust programs, using a prophecy-based logic model in the style of RustHorn and taking a semantic approach in the style of RustBelt. We name this research project *RustHornBelt*, as it combines the approaches of RustHorn and RustBelt. We build a new formalization of prophecy in Iris (Chapter 3), where the information about prophecy lives only in the ghost state and we can perform an operation that we call dependent resolution. We also develop a low-level foundation for verifying imperative programs in Iris (Chapter 4), with a new technique for spending an unbounded number of logical steps at one physical step. We give new semantic models or *semantic types* to program types in Rust (Chapter 5), extending the semantic approach of RustBelt with refinement of a type by a pure value that is parametrized over assignments on prophecy variables, which allows us to use the *prophecy*-based logic model in the style of RustHorn. Finally, in order to formalize and verify type-based translation of Rust programs into logic models, we introduce and semantically model a type system for Rust that manipulates a logic model, which we dub the *refined type system* (Chapter 6). Thanks to the semantic approach, we can verify each deduction rule of the type system *separately*, which brings about *extensibility*, allowing us to flexibly add new features and libraries of Rust.

The remaining part of this thesis is structured as follows. In Chapter 2, we introduce the higher-order separation logic Iris as preliminaries for the following chapters. In Chapter 3, we present a new platform for prophecy. In Chapter 4, we present a low-level foundation for verification of imperative programming. In Chapter 5, we define our notion of the semantic type and model various Rust types. In Chapter 6, we introduce and semantically model our refined type system for Rust, proving some deduction rules in detail and giving some examples of verification. In Chapter 7, we present the conclusion of the paper and discuss future work.

This work emerged during my research internship at the RustBelt team of Max Planck Institute for Software Systems from September to December 2020. Nevertheless, virtually all the core ideas of this thesis were conceived by the author unless otherwise noted, although I received valuable advice about the direction of this research. It is worth noting here that Jacques-Henri Jourdan conceived the idea of using cumulative time receipts for swelling persistent time receipts (§4.2; `CUMU``TIME`-`SWELL`-`PERS``TIME`), which was an essential technique for completing the formalization of this thesis.

## Chapter 2

# Preliminaries on the Higher-Order Separation Logic Iris

A *separation logic* is a logic for scalable and modular reasoning about the resources, which was introduced by O’Hearn et al. (2001) and Reynolds (2002). Various separation logics have been designed so far. The separation logic *Iris* (Jung et al., 2015, 2018c) is a general-purpose separation logic that attains outstanding extensibility by *impredicativity* (or higher-order ghost states), which is made possible by *step indexing* (Appel and McAllester, 2001; Birkedal et al., 2011). Also, *Iris* is *mechanized* in the Coq Proof Assistant with sophisticated tactics and proof modes (Krebbers et al., 2017, 2018).

*Iris* is applied to verification of various contexts; to name a few, C11’s weak memory model (Kaiser et al., 2017), refinement on fine-grained concurrency (Frumin et al., 2018), the ST monad (Timany et al., 2018), persistency of the Intel-x86 Architecture (Raad et al., 2020), low-level sandboxing (Sammler et al., 2020a), session types (Hinrichsen et al., 2020), distributed databases (Gondelman et al., 2021), effect handlers (de Vilhena and Pottier, 2021), and gradual typing (Jacobs et al., 2021). And notably, *Iris* is a vital foundation for RustBelt (Jung et al., 2018a). Our work RustHornBelt, which is based on the technique of RustBelt, is also built on top of the separation logic *Iris*.

In this chapter, we briefly describe how we can use the higher-order separation logic *Iris* at a fairly high level. In § 2.1, we introduce basic connectives and deduction rules of *Iris*. In § 2.2, we introduce the notion of the *resource algebra* and explain how we can customize *Iris* using resource algebras. In § 2.3, we describe how we can use the *lifetime logic*, which was built on top of *Iris* by Jung et al. (2018a) to discuss borrowing of resources under lifetimes in Rust. This chapter serves as technical preliminaries for the following chapters Chapter 3, Chapter 4, Chapter 5 and Chapter 6.

We still do not introduce any program logic in this chapter. The weakest precondition (and the Hoare triple) for imperative programs can be *defined* using the existing connectives of *Iris*. Because our definition of the weakest precondition is a part of our contribution, it is described later in § 4.2.

### 2.1 Basic Connectives and Deduction Rules of Iris

In this section, we introduce and describe basic logical connectives and deduction rules of *Iris*.

Before introducing the connectives, we explicitly write down here the precedence on the connectives: (i, the strongest)  $\Box, \triangleright, \diamond, \Im_{\mathcal{E}}$  (and similar connectives); (ii)  $*$ ,  $\wedge$ ,  $\vee$ ; (iii)  $\multimap$ ,  $\Rightarrow$ ,  $\multimap_{\mathcal{E}}$ ,  $\Rightarrow_{\mathcal{E}}$  (and similar connectives); (iv, the weakest)  $\forall x.$ ,  $\exists x.$ .

**Iris Proposition** An *Iris proposition*  $P, Q, R : \text{IProp}$  is roughly speaking a *predicate over resources*. More precisely, it is a *step-indexed, monotone* predicate over *valid* resources. Also, a resource for an *Iris proposition* is called a *ghost state*. The formal

model of Iris propositions is well described in (Jung et al., 2018c, §4). In this thesis, we often call an Iris proposition just a proposition.

We call a proposition of metalogic a *pure proposition* and represent it by variables  $\phi, \psi$ , to distinguish it from an Iris proposition. We write  $\text{Prop}$  for the type of pure propositions.

**Sequent** Iris has a *sequent*  $P \vdash Q$ , which is a (metalogic) binary relation on Iris propositions  $P, Q : \text{IProp}$ . It satisfies reflexivity, transitivity and anti-symmetry.

$$\begin{array}{ccc} \text{SEQ-REFL} & \text{SEQ-TRANS} & \text{SEQ-ANTISYM} \\ P \vdash P & \frac{P \vdash Q \quad Q \vdash R}{P \vdash R} & \frac{P \vdash Q \quad Q \vdash P}{P = Q} \end{array}$$

**Basic Connectives** Iris has the basic connectives of predicate logic: the truth  $\text{True}$ , falsehood  $\text{False}$ , conjunction  $P \wedge Q$ , universal quantification  $\forall x : T. P_x$ , disjunction  $P \vee Q$ , existential quantification  $\exists x : T. P_x$ , and implication  $P \Rightarrow Q$ . These connectives are defined based on Kripke semantics of intuitionistic logic and thus satisfy *all tautologies of intuitionistic logic*. In particular, the following properties hold.

$$\begin{array}{ccc} \text{TRUE-INTRO} & \text{FALSE-ELIM} & \text{AND-UNIV} \\ P \vdash \text{True} & \text{False} \vdash P & \frac{R \vdash P \quad R \vdash Q}{R \vdash P \wedge Q} \\ \text{OR-UNIV} & \text{EXIST-UNIV} & \text{ALL-UNIV} \\ \frac{P \vdash R \quad Q \vdash R}{P \vee Q \vdash R} & \frac{\forall x : T. (P_x \vdash Q)}{(\exists x : T. P_x) \vdash Q} & \frac{\forall x : T. (Q \vdash P_x)}{Q \vdash (\forall x : T. P_x)} \\ \text{IMP-UNIV} & & \\ \frac{P \wedge Q \vdash R}{P \vdash Q \Rightarrow R} & & \end{array}$$

Note that we can ‘throw away’ any proposition  $P$  by the rule **TRUE-INTRO**. A separation logic that satisfies this rule is called *affine*.

We can also promote a pure proposition  $\phi : \text{Prop}$  into an Iris proposition. This promotion preserves basic connectives – truth, falsehood, conjunction, universal quantification, disjunction, existential quantification and implication.

We also introduce the following shorthand.

$$\begin{aligned} P \Leftrightarrow Q & := (P \Rightarrow Q) \wedge (Q \Rightarrow P) \\ \exists x \text{ s.t. } \phi_x. P & := \exists x. \phi_x \wedge P & \forall x \text{ s.t. } \phi_x. P & := \forall x. \phi_x \Rightarrow P \end{aligned}$$

We say that an Iris proposition  $P$  is a tautology if it is equal to  $\text{True}$ .

**Separating Conjunction and Magic Wand** The *separating conjunction*  $P * Q$  of propositions  $P$  and  $Q$  represents a resource that can be written as a composition of some resource satisfying  $P$  and some resource satisfying  $Q$ .

The separating conjunction is associative, commutative and unital (on  $\text{True}$ ).

$$\begin{array}{ccc} \text{SEP-ASSOC} & \text{SEP-COMM} & \text{SEP-UNIT} \\ (P * Q) * R = P * (Q * R) & P * Q = Q * P & P * \text{True} = P \end{array}$$

Also, we can compose sequents in parallel according to separating conjunction, which is a very important property.

$$\text{SEP-PARALLEL} \quad \frac{P \vdash Q \quad P' \vdash Q'}{P * P' \vdash Q * Q'}$$

We can derive the following rules.

$$\begin{array}{c}
\text{SEP-LOSE} \quad \text{SEP-AND} \quad \text{SEP-ALL} \quad \text{SEP-FRAME} \\
P * Q \vdash P \quad P * Q \vdash P \wedge Q \quad P * \forall x. Q_x \vdash \forall x. P * Q_x \quad \frac{P \vdash Q}{P * R \vdash Q * R}
\end{array}$$

The separating conjunction has the right adjoint, which is written as  $P \multimap Q$  and called the *magic wand*.

$$\frac{\text{WAND-UNIV} \quad P * Q \vdash R}{P \vdash Q \multimap R}$$

A magic wand  $P \multimap Q$  from  $P$  to  $Q$  can be understood as a resource such that if we compose it with any resource satisfying  $P$  then we get a resource satisfying  $Q$ . Given that the implication is the right adjoint of the (plain) conjunction, roughly speaking, the magic wand is to the separating conjunction what the implication is to the conjunction.

We can derive the following rules.

$$\begin{array}{c}
\text{IMP-WAND} \quad \text{WAND-APPLY} \quad \text{WAND-CURRY} \\
P \Rightarrow Q \vdash P \multimap Q \quad P * (P \multimap Q) \vdash Q \quad P * Q \multimap R = P \multimap (Q * R) \\
\\
\text{WAND-ID} \quad \text{WAND-COMP} \\
\text{True} \vdash P \multimap P \quad (P \multimap Q) * (Q \multimap R) \vdash P \multimap R \\
\\
\text{WAND-PARALLEL} \\
(P \multimap Q) * (P' \multimap Q') \vdash P * P' \multimap Q * Q'
\end{array}$$

Since the separating conjunction is a left adjoint of the magic wand, by the ‘left adjoints preserve colimits’ law, we have the following.

$$\begin{array}{c}
\text{SEP-EXIST} \\
P * \exists x. Q_x = \exists x. P * Q_x
\end{array}$$

**Persistence Modality** Iris has the persistence modality  $\Box P$ . Intuitively,  $\Box P$  means that  $P$  holds in a ‘lightweight’ way.

The persistence modality satisfies the following property.

$$\begin{array}{c}
\text{PERS-AND-SEP} \\
Q \wedge \Box P = Q * \Box P
\end{array}$$

The proposition  $Q \wedge \Box P$  can be understood as a resource that satisfies both  $Q$  and  $\Box P$ . By this rule, we can *separate out* a resource that satisfies  $\Box P$  from that resource (note that  $Q \wedge \Box P$  is equivalent to  $(Q \wedge \Box P) \wedge \Box P$ ).

The persistence modality forms a comonad.

$$\begin{array}{c}
\text{PERS-MONO} \quad \text{PERS-ELIM} \quad \text{PERS-IDEM} \\
\frac{P \vdash Q}{\Box P \vdash \Box Q} \quad \Box P \vdash P \quad \Box P = \Box \Box P
\end{array}$$

The persistence modality commutes with the separating conjunction, conjunction, universal quantification, disjunction, existential quantification, and pure-proposition promotion.

$$\begin{array}{c}
\text{PERS-COMM-SEP} \quad \text{PERS-COMM-AND} \quad \text{PERS-COMM-ALL} \\
\Box(P * Q) = \Box P * \Box Q \quad \Box(P \wedge Q) = \Box P \wedge \Box Q \quad \Box(\forall x. P_x) = \forall x. \Box P_x \\
\\
\text{PERS-COMM-OR} \quad \text{PERS-COMM-EXIST} \quad \text{PERS-COMM-PURE} \\
\Box(P \vee Q) = \Box P \vee \Box Q \quad \Box(\exists x. P_x) = \exists x. \Box P_x \quad \Box \phi = \phi
\end{array}$$

Also, under the persistence modality, the separating conjunction and the conjunction coincide and so do the magic wand and the implication.

$$\begin{array}{cc} \text{PERS-EQ-SEP-AND} & \text{PERS-EQ-WAND-IMP} \\ \square(P * Q) = \square(P \wedge Q) & \square(P \multimap Q) = \square(P \Rightarrow Q) \end{array}$$

We say that  $P$  is *persistent* and write  $\text{persistent}(P)$  when it is equal to  $\square P$ .

$$\text{persistent}(P) := P = \square P$$

Persistent propositions satisfy the following properties.

$$\begin{array}{ccc} \text{PERSISTENT-EQ-SEP-AND} & \text{PERSISTENT-EQ-WAND-IMP} & \text{PERSISTENT-INTRO-PERS} \\ \frac{\text{persistent}(P)}{P * Q = P \wedge Q} & \frac{\text{persistent}(P)}{P \multimap Q = P \Rightarrow Q} & \frac{\text{persistent}(P) \quad P \vdash Q}{P \vdash \square Q} \\ \\ \text{PERSISTENT-DUP} & \text{PERSISTENT-RETAIN} & \\ \frac{\text{persistent}(P)}{P = P * P} & \frac{\text{persistent}(P) \quad Q \vdash P}{Q \vdash Q * P} & \end{array}$$

The rule [PERSISTENT-INTRO-PERS](#) says that if we can prove  $Q$  out of a persistent assumption, we can also prove  $\square Q$ .

We have the following lemmas for finding persistent propositions.

$$\begin{array}{ccc} \text{PERSISTENT-PERS} & \text{PERSISTENT-PURE} & \text{PERSISTENT-SEP} \\ \text{persistent}(\square P) & \text{persistent}(\phi) & \frac{\text{persistent}(P) \quad \text{persistent}(Q)}{\text{persistent}(P * Q)} \\ \\ \text{PERSISTENT-AND} & \text{PERSISTENT-OR} & \\ \frac{\text{persistent}(P) \quad \text{persistent}(Q)}{\text{persistent}(P \wedge Q)} & \frac{\text{persistent}(P) \quad \text{persistent}(Q)}{\text{persistent}(P \vee Q)} & \\ \\ \text{PERSISTENT-ALL} & \text{PERSISTENT-EXIST} & \\ \frac{\forall x. \text{persistent}(P_x)}{\text{persistent}(\forall x. P_x)} & \frac{\forall x. \text{persistent}(P_x)}{\text{persistent}(\exists x. P_x)} & \end{array}$$

**Step Indexing** Iris is a *step-indexed* logic. Step indexing allows us to introduce *general recursion* and *impredicativity*, which is an essential source of expressivity and extensibility of Iris. Later we introduce the *later modality*  $\triangleright P$ , which is an important connective for step indexing. Before introducing the later modality, we briefly introduce *step indexing* in Iris. See [Jung et al. \(2018c, §4\)](#) for details.

A *step index* is any natural number  $n$  (in the setting of Iris). Intuitively, the world is *stratified* over step indices and the world gets more *clearer* as the step index grows. Formally, a step-indexed type (also known as an ordered family of equivalences or OFE) is a type equipped with the *step-indexed equality*  $(\stackrel{n}{=})_n$ , which is a family of equivalence relations indexed by step indices which is anti-monotone over step indices and whose limit coincides with the equality. A function between step-indexed types is said to be *non-expansive* if it preserves the step-indexed equality. Also, a function  $f$  between step-indexed types is said to be *contractive* if it further satisfies  $f x \stackrel{0}{=} f y$  and  $x \stackrel{n}{=} y \Rightarrow f x \stackrel{n+1}{=} f y$ . We have the following fixed point theorem on contractive functions, which is very useful in practice. Taking the fixed point by contractiveness can be understood as *coinduction over the step indices*.

**Theorem 2.1** (Banach’s Fixed Point Theorem). *Any contractive function  $f: T \rightarrow T$  over the step-indexed type  $T$  has a unique fixed point, i.e., has a unique solution  $x$  for the equation  $x = f x$ .*

A step-indexed proposition  $X: \text{SProp}$  is a predicate over step indices that is anti-monotone over step indices. The type  $\text{SProp}$  is equipped with the step-indexed equality  $X \stackrel{n}{=} Y$ , which is defined as  $\forall m \leq n. X m = Y m$ , the equality at the step indices  $0, \dots, n$ . An Iris proposition  $P: \text{IProp}$  is a step-indexed predicate over resources, with some restrictions and quotients. The type  $\text{IProp}$  is also equipped with the step-indexed equality in a manner similar to  $\text{SProp}$ .

**Later Modality** Iris has the *later modality*  $\triangleright P$ . Intuitively,  $\triangleright P$  means that  $P$  holds *one step later*. More precisely,  $\triangleright P$  holds at the step index 0 regardless of  $P$  and holds at the step index  $n + 1$  if  $P$  holds at the step index  $n$ .

The later modality  $\triangleright$  is important for recursion and impredicativity because it is a *contractive* function over  $\text{IProp}$ . Other logical connectives are also designed to be *non-expansive*. So roughly speaking, a function over  $\text{IProp}$  composed by logical connectives of Iris is non-expansive *if all occurrences of the argument are under the later modality*. As we mentioned earlier, contractiveness allows us to take a *unique fixed point*, which is very useful for recursive definitions.

The later modality is monotone and can be introduced.

$$\frac{\text{LATER-MONO} \quad P \vdash Q}{\triangleright P \vdash \triangleright Q} \qquad \text{LATER-INTRO} \quad P \vdash \triangleright P$$

The later modality is not idempotent because it just says ‘one step later’ instead of ‘any steps later’.

Also, we have the following rule for *Löb induction*.

$$\frac{\text{LÖB}}{\triangleright P \Rightarrow P = P}$$

It means that whenever we prove  $P$  we can assume  $\triangleright P$ . This is a strong deduction rule when we deal with recursive definition with the later modality.

The later modality also commutes with the truth, separating conjunction, conjunction, universal quantification, disjunction, inhabited existential quantification, implication, and persistent modality.

$$\begin{array}{lll} \text{LATER-COMM-TRUE} & \text{LATER-COMM-SEP} & \text{LATER-COMM-AND} \\ \triangleright \text{True} = \text{True} & \triangleright (P * Q) = \triangleright P * \triangleright Q & \triangleright (P \wedge Q) = \triangleright P \wedge \triangleright Q \\ \\ \text{LATER-COMM-ALL} & \text{LATER-COMM-OR} & \text{LATER-COMM-INHEXIST} \\ \triangleright (\forall x. P_x) = \forall x. \triangleright P_x & \triangleright (P \vee Q) = \triangleright P \vee \triangleright Q & \frac{\exists \_ : T. \text{True}}{\triangleright (\exists x: T. P_x) = \exists x: T. \triangleright P_x} \\ \\ & \text{LATER-COMM-PERS} & \\ & \triangleright \Box P = \Box \triangleright P & \end{array}$$

The following also holds because of [LATER-COMM-PERS](#).

$$\frac{\text{PERSISTENT-LATER} \quad \text{persistent}(P)}{\text{persistent}(\triangleright P)}$$

In order to help discussion about the later modality, we introduce the following *except-0 modality*.

$$\diamond P := \triangleright \text{False} \vee P$$

It means that  $P$  holds at least except at the step index 0. It forms a monad.

$$\begin{array}{ccc} \text{EXCEPT0-MONO} & & \\ \frac{P \vdash Q}{\diamond P \vdash \diamond Q} & \text{EXCEPT0-INTRO} & \text{EXCEPT0-IDEM} \\ & P \vdash \diamond P & \diamond \diamond P = \diamond P \end{array}$$

It also satisfies the following properties.

$$\begin{array}{ccc} \text{EXCEPT0-LATER} & \text{EXCEPT0-COMM-SEP} & \text{EXCEPT0-COMM-AND} \\ \diamond \triangleright P = \triangleright P & \diamond(P * Q) = \diamond P * \diamond Q & \diamond(P \wedge Q) = \diamond P \wedge \diamond Q \\ \\ \text{EXCEPT0-COMM-ALL} & \text{EXCEPT0-COMM-OR} & \text{EXCEPT0-COMM-EXIST} \\ \diamond(\forall x. P_x) = \forall x. \diamond P_x & \diamond(P \vee Q) = \diamond P \vee \diamond Q & \diamond(\exists x. P_x) = \exists x. \diamond P_x \\ \\ & \text{EXCEPT0-COMM-PERS} & \\ & \diamond \square P = \square \diamond P & \end{array}$$

We say that  $P$  is *timeless* and write  $\text{timeless}(P)$  when  $\triangleright P$  is equal to  $\diamond P$ .

$$\text{timeless}(P) := \triangleright P = \diamond P$$

If we peek into the model,  $\text{timeless}(P)$  turns out to be equivalent to saying that  $P$  is constant over the step indices.

We have the following lemmas for finding timeless propositions.

$$\begin{array}{ccc} \text{TIMELESS-SEP} & & \text{TIMELESS-AND} \\ \frac{\text{timeless}(P) \quad \text{timeless}(Q)}{\text{timeless}(P * Q)} & & \frac{\text{timeless}(P) \quad \text{timeless}(Q)}{\text{timeless}(P \wedge Q)} \\ \\ \text{TIMELESS-PURE} & & \\ \text{timeless}(\phi) & & \\ \\ \text{TIMELESS-OR} & \text{TIMELESS-ALL} & \text{TIMELESS-EXIST} \\ \frac{\text{timeless}(P) \quad \text{timeless}(Q)}{\text{timeless}(P \vee Q)} & \frac{\forall x. \text{timeless}(P_x)}{\text{timeless}(\forall x. P_x)} & \frac{\forall x. \text{timeless}(P_x)}{\text{timeless}(\exists x. P_x)} \\ \\ \text{TIMELESS-WAND} & & \text{TIMELESS-IMP} \\ \frac{\text{timeless}(P)}{\text{timeless}(Q \multimap P)} & & \frac{\text{timeless}(P)}{\text{timeless}(Q \Rightarrow P)} \end{array}$$

**Fancy Update Modality** Iris also has the *fancy update modality*.

First, we introduce some notions. We use some infinite type  $\text{InvName}$  and call its elements an *invariant name* and represent them by  $\iota$ . Internally in the model of Iris, an invariant name is used for naming an *Iris proposition used for an invariant*. A *mask*  $\mathcal{E} : \mathbb{P} \text{InvName}$  is a set of invariant names. A *namespace*  $\mathcal{N}$  is an *infinite* set of invariant names.<sup>1</sup> We write  $\top$  for the universal set of the type  $\text{InvName}$ .

Iris has the *fancy update modality*  $\mathcal{E} \mapsto^{\mathcal{E}'} P$  for masks  $\mathcal{E}, \mathcal{E}'$ . Roughly speaking,  $\mathcal{E} \mapsto^{\mathcal{E}'} P$  represents a resource such that it can be updated into a resource satisfying in a frame-preserving way; before the update we *own* the invariant names in  $\mathcal{E}$ , which can be used for the update, and after the update we *own* the invariant names in  $\mathcal{E}'$ . Actually, the fancy update modality is defined in Iris combining simpler connectives, which is described in (Jung et al., 2018c, §7).

<sup>1</sup> This formulation is a bit different from the actual Coq implementation in Iris.

The fancy update modality forms an indexed monad.

$$\begin{array}{c}
\text{FUPD-MONO} \\
\frac{P \vdash Q}{\mathcal{E} \Rightarrow^{\mathcal{E}'} P \vdash \mathcal{E} \Rightarrow^{\mathcal{E}'} Q}
\end{array}
\quad
\begin{array}{c}
\text{FUPD-INTRO} \\
P \vdash \mathcal{E} \Rightarrow^{\mathcal{E}} P
\end{array}
\quad
\begin{array}{c}
\text{FUPD-JOIN} \\
\mathcal{E} \Rightarrow^{\mathcal{E}'} \mathcal{E}' \Rightarrow^{\mathcal{E}''} P \vdash \mathcal{E} \Rightarrow^{\mathcal{E}''} P
\end{array}$$

We can regard the fancy update modality as some *computation*, as is often the case for various monads in functional programming.

We also have the following stronger variant of [FUPD-INTRO](#).<sup>2</sup>

$$\begin{array}{c}
\text{FUPD-BOUNCE} \\
P \vdash \mathcal{E} + \mathcal{E}' \Rightarrow^{\mathcal{E}} \mathcal{E} \Rightarrow^{\mathcal{E} + \mathcal{E}'} P
\end{array}$$

The fancy update modality also satisfies the following *frame rule*.

$$\begin{array}{c}
\text{FUPD-FRAME} \\
Q * \mathcal{E} \Rightarrow^{\mathcal{E}'} P \vdash \mathcal{E} \Rightarrow^{\mathcal{E}'} (Q * P)
\end{array}$$

The fancy update modality can clear away the except-0 modality.

$$\begin{array}{c}
\text{FUPD-EXCEPT0} \\
\diamond \mathcal{E} \Rightarrow^{\mathcal{E}'} \diamond P = \mathcal{E} \Rightarrow^{\mathcal{E}'} P
\end{array}$$

We can also derive the following rules.

$$\begin{array}{c}
\text{FUPD-MERGE} \\
\mathcal{E} \Rightarrow^{\mathcal{E}'} P * \mathcal{E}' \Rightarrow^{\mathcal{E}''} Q \vdash \mathcal{E} \Rightarrow^{\mathcal{E}''} (P * Q)
\end{array}
\quad
\begin{array}{c}
\text{FUPD-WKN-MASK} \\
\mathcal{E} \Rightarrow^{\mathcal{E}'} P \vdash \mathcal{E} + \mathcal{E}'' \Rightarrow^{\mathcal{E}' + \mathcal{E}''} P
\end{array}$$

[FUPD-MERGE](#) follows from [FUPD-FRAME](#) and [FUPD-JOIN](#). [FUPD-WKN-MASK](#) follows from [FUPD-BOUNCE](#), [FUPD-FRAME](#) and [FUPD-JOIN](#).

The fancy update does *not* commute with any of the separating conjunction, magic wand, conjunction, universal quantification, disjunction, existential quantification, implication, persistence modality, and later modality.

We also use the following shorthand for propositions  $P, Q$  and masks  $\mathcal{E}, \mathcal{E}'$ .

$$P \overset{\mathcal{E}}{\Rightarrow^*} \overset{\mathcal{E}'}{Q} := P \multimap \mathcal{E} \Rightarrow^{\mathcal{E}'} Q \quad P \overset{\mathcal{E}}{\Rightarrow} \overset{\mathcal{E}'}{Q} := \Box (P \overset{\mathcal{E}}{\Rightarrow^*} \overset{\mathcal{E}'}{Q})$$

We call  $P \overset{\mathcal{E}}{\Rightarrow^*} \overset{\mathcal{E}'}{Q}$  a *view shift* from  $P$  to  $Q$ , under the input and output masks  $\mathcal{E}, \mathcal{E}'$ . Intuitively,  $P \overset{\mathcal{E}}{\Rightarrow^*} \overset{\mathcal{E}'}{Q}$  means that, by consuming a resource satisfying  $P$  and performing an update that inputs the mask  $\mathcal{E}$  and outputs the mask  $\mathcal{E}'$ , we get a resource satisfying  $Q$ . The proposition  $P \overset{\mathcal{E}}{\Rightarrow} \overset{\mathcal{E}'}{Q}$  means that the view shift  $P \overset{\mathcal{E}}{\Rightarrow^*} \overset{\mathcal{E}'}{Q}$  holds persistently; the symbol looks more like the implication than the magic wand, which is justified by the fact that under the persistent modality the implication and the magic wand coincide ([PERS-EQ-WAND-IMP](#)).

We can derive the following properties on the view shift.

$$\begin{array}{c}
\text{VSHIFT-APPLY} \\
P * (P \overset{\mathcal{E}}{\Rightarrow^*} \overset{\mathcal{E}'}{Q}) \vdash \mathcal{E} \Rightarrow^{\mathcal{E}'} Q
\end{array}
\quad
\begin{array}{c}
\text{VSHIFT-ID} \\
\text{True} \vdash P \overset{\mathcal{E}}{\Rightarrow^*} \overset{\mathcal{E}}{P}
\end{array}$$

$$\begin{array}{c}
\text{VSHIFT-COMP} \\
(P \overset{\mathcal{E}}{\Rightarrow^*} \overset{\mathcal{E}'}{Q}) * (Q \overset{\mathcal{E}'}{\Rightarrow^*} \overset{\mathcal{E}''}{R}) \vdash P \overset{\mathcal{E}}{\Rightarrow^*} \overset{\mathcal{E}''}{R}
\end{array}$$

$$\begin{array}{c}
\text{VSHIFT-MERGE} \\
(P \overset{\mathcal{E}}{\Rightarrow^*} \overset{\mathcal{E}'}{Q}) * (P' \overset{\mathcal{E}'}{\Rightarrow^*} \overset{\mathcal{E}''}{Q'}) \vdash P * P' \overset{\mathcal{E}}{\Rightarrow^*} \overset{\mathcal{E}''}{Q * Q'}
\end{array}$$

<sup>2</sup> For sets  $A, B$ , we write  $A+B$  for the disjoint union of them, i.e., the union  $A \cup B$  with the precondition that the argument sets are disjoint  $A \cap B = \emptyset$ .

$$\text{VSHIFT-WKN-MASK} \\ P \stackrel{\mathcal{E}}{\rightsquigarrow} Q \vdash P \stackrel{\mathcal{E}+\mathcal{E}''}{\rightsquigarrow} Q$$

$$\text{VSHIFT-EXCEPT0} \\ P \stackrel{\mathcal{E}}{\rightsquigarrow} Q \diamond P \stackrel{\mathcal{E}}{\rightsquigarrow} Q = P \stackrel{\mathcal{E}}{\rightsquigarrow} Q \diamond P \stackrel{\mathcal{E}}{\rightsquigarrow} Q$$

For common use cases, the input and output masks for the fancy update are the same. So we introduce the following shorthand.

$$\Rightarrow_{\mathcal{E}} P := \mathcal{E} \Rightarrow P \quad P \rightsquigarrow_{\mathcal{E}} Q := P \stackrel{\mathcal{E}}{\rightsquigarrow} Q \quad P \Rightarrow_{\mathcal{E}} Q := P \stackrel{\mathcal{E}}{\Rightarrow} Q$$

We also introduce the following *step-taking* fancy update modality.

$$\Rightarrow_{\mathcal{E}} P := \mathcal{E} \Rightarrow^{\emptyset} \triangleright \Rightarrow P$$

It means that we get  $P$  after performing an update consuming the mask  $\mathcal{E}$ , spending one step index, and performing an update that returns the mask  $\mathcal{E}$ . Informally, when we have  $\Rightarrow_{\mathcal{E}} P$ , we say that after one *logical step* under the mask  $\mathcal{E}$  we get  $P$ . We also introduce the following shorthand.

$$P \rightsquigarrow_{\mathcal{E}} Q := P \ast \Rightarrow_{\mathcal{E}} Q \quad P \Rightarrow_{\mathcal{E}} Q := \square (P \rightsquigarrow_{\mathcal{E}} Q)$$

Moreover, we also introduce the following shorthand for repeating logical steps.

$$\Rightarrow_{\mathcal{E}}^0 P := \Rightarrow_{\mathcal{E}} P \quad \Rightarrow_{\mathcal{E}}^{n+1} P := \Rightarrow_{\mathcal{E}} \Rightarrow_{\mathcal{E}}^n P \\ P \rightsquigarrow_{\mathcal{E}}^n Q := P \ast \Rightarrow_{\mathcal{E}}^n Q \quad P \Rightarrow_{\mathcal{E}}^n Q := \square (P \rightsquigarrow_{\mathcal{E}}^n Q)$$

Note that  $\Rightarrow_{\mathcal{E}}^0 P$  is defined as  $\Rightarrow_{\mathcal{E}} P$  instead of  $P$ . Note also that  $\Rightarrow_{\mathcal{E}}^1 P$ ,  $P \rightsquigarrow_{\mathcal{E}}^1 Q$  and  $P \Rightarrow_{\mathcal{E}}^1 Q$  are equal to  $\Rightarrow_{\mathcal{E}} P$ ,  $P \rightsquigarrow_{\mathcal{E}} Q$  and  $P \Rightarrow_{\mathcal{E}} Q$ , respectively.

We can *derive* the following properties on the step-taking fancy update modality.

$$\begin{array}{c} \text{STFUPD-MONO} \\ \frac{P \vdash Q}{\Rightarrow_{\mathcal{E}}^n P \vdash \Rightarrow_{\mathcal{E}}^n Q} \end{array} \quad \begin{array}{c} \text{LATER-STFUPD} \\ \triangleright^n P \vdash \Rightarrow_{\mathcal{E}}^n P \end{array} \quad \begin{array}{c} \text{STFUPD-FUPD} \\ \Rightarrow_{\mathcal{E}} \Rightarrow_{\mathcal{E}}^n \Rightarrow_{\mathcal{E}} P = \Rightarrow_{\mathcal{E}}^n P \end{array} \\ \\ \begin{array}{c} \text{STFUPD-WKN-MASK} \\ \Rightarrow_{\mathcal{E}}^n P \vdash \Rightarrow_{\mathcal{E}+\mathcal{E}'}^n P \end{array} \quad \begin{array}{c} \text{STFUPD-FRAME} \\ Q \ast \Rightarrow_{\mathcal{E}}^n P \vdash \Rightarrow_{\mathcal{E}}^n (Q \ast P) \end{array} \\ \\ \begin{array}{c} \text{STFUPD-MERGE} \\ \Rightarrow_{\mathcal{E}}^n P \ast \Rightarrow_{\mathcal{E}}^n Q \vdash \Rightarrow_{\mathcal{E}}^n (P \ast Q) \end{array} \quad \begin{array}{c} \text{STFUPD-ADD-NSTEP} \\ \Rightarrow_{\mathcal{E}}^n \Rightarrow_{\mathcal{E}}^m P = \Rightarrow_{\mathcal{E}}^{n+m} P \end{array}$$

Note that step-taking updates can be merged *in parallel* by the rule [STFUPD-MERGE](#), which follows from [FUPD-MERGE](#) and [LATER-COMM-SEP](#). The rule [STFUPD-WKN-MASK](#) follows from [FUPD-BOUNCE](#), [FUPD-FRAME](#), [LATER-INTRO](#) and [LATER-COMM-SEP](#).

The step-taking fancy update modality satisfies the following *soundness* theorem.

**Theorem 2.2** (Soundness of the Step-Taking Fancy Update Modality). *For any pure proposition  $\phi$ : Prop, if  $\Rightarrow_{\mathcal{E}}^n \phi$  is a tautology for some mask  $\mathcal{E}$  and natural number  $n$ , then  $\phi$  holds.*

**Invariant** In Iris, we can store an Iris proposition  $\triangleright P$  to the world as an *invariant* and share it by the *invariant token*  $\boxed{P}^{\mathcal{E}}$  which is *persistent*. The mask  $\mathcal{E}$  is what we need to consume to take out the Iris proposition  $\triangleright P$ .

Formally, the *invariant token*  $\boxed{P}^{\mathcal{E}}$  is defined as the following *persistent view shift*.

$$\boxed{P}^{\mathcal{E}} := \text{True} \stackrel{\mathcal{E}}{\Rightarrow} \triangleright P \ast (\triangleright P \stackrel{\emptyset}{\rightsquigarrow} \text{True})$$

By consuming the mask  $\mathcal{E}$ , we can take out  $\triangleright P$ , and after we put back  $\triangleright P$  we can retrieve the mask  $\mathcal{E}$ .

We can get an invariant token by the following rule.

$$\text{INV-INTRO} \\ \triangleright P \vdash \Leftrightarrow_{\emptyset} \boxed{P}^{\mathcal{N}}$$

We can use any namespace  $\mathcal{N}$ , which is an infinite set of invariant names. Internally, we choose some unused invariant name  $\iota$  from  $\mathcal{N}$  (always only a finite number of invariant names are used), associate  $\iota$  with  $P$ , and store  $\triangleright P$  to some globally shared place; when we temporarily take out  $\triangleright P$ , we temporarily deposit the ownership on the invariant name  $\iota$  to that shared place. Actually, in order to achieve this operation, the fancy update modality  $\mathcal{E} \Leftrightarrow^{\mathcal{E}'}$  is equipped with the masks  $\mathcal{E}, \mathcal{E}'$ .

Later in § 3.3, we use the mechanism of the invariant to formalize the prophecy resolution ([PROPH-RESOLVE](#)).

## 2.2 Customizing Iris by Resource Algebras

We can customize Iris by registering new *resource algebras* to the global resource. A resource algebra (RA) is a variant of a partial commutative monoid (PCM). In this section, we define the RA, introduce various RAs, and describe how we can customize Iris through RAs.

**Resource Algebra** A *resource algebra* (RA)  $Ra$  is a quintuple of (i) the underlying type  $|Ra|$ : Type, (ii) the *composition* operation  $\cdot : |Ra| \rightarrow |Ra| \rightarrow |Ra|$ , (iii) the *unit*  $\varepsilon : |Ra|$ , (iv) the *validity* predicate  $\checkmark : |Ra| \rightarrow \text{Prop}$ , and (v) the *core* of each item  $|-| : |Ra| \rightarrow |Ra|$ , satisfying the following properties.<sup>3</sup>

RA-ASSOC $(a \cdot b) \cdot c = a \cdot (b \cdot c)$	RA-COMM $a \cdot b = b \cdot a$	RA-UNIT $a \cdot \varepsilon = a$	RA-VALID-UNIT $a \cdot \varepsilon = a$
RA-VALID-MONO $\frac{a \leq b \quad \checkmark b}{\checkmark a}$	RA-CORE-UNIT $a \cdot  a  = a$	RA-CORE-IDEM $  a   =  a $	RA-CORE-MONO $\frac{a \leq b}{ a  \leq  b }$

$$\text{where } a \leq b := \exists c. a \cdot b = c$$

The composition is associative, commutative and unital ([RA-ASSOC](#), [RA-COMM](#), [RA-UNIT](#)). The validity predicate holds for the unit and anti-monotone over composition ([RA-VALID-UNIT](#), [RA-VALID-MONO](#)). The core  $|a|$  of an item  $a$  can be understood as what works like a unit in the presence of  $a$  ([RA-CORE-UNIT](#)). The core is idempotent and monotone over composition ([RA-CORE-IDEM](#), [RA-CORE-MONO](#)).

RA is similar to a partial commutative monoid (PCM) but different from that in that (i) RA uses a total composition operation and separately discusses validity, while PCM has a partial composition operation, and (ii) RA has the core operation, which PCM does not have.

We introduce the following relation  $a \rightsquigarrow B$  for an item  $a : |Ra|$  and a set of items  $B : \mathbb{P}|Ra|$ .

$$a \rightsquigarrow B := \forall c \text{ s.t. } \checkmark(a \cdot c). \exists b \in B. \checkmark(b \cdot c)$$

It means that we can safely update  $a$  into some element of the set  $B$  without losing validity for any *frame*  $c$ , i.e., the remaining part of the global resource. This kind of update is called *frame-preserving update*.

<sup>3</sup> This algebra is actually called a *unital* RA in Iris. Iris allows an RA to lack the unit and the totality of the core operation. In this thesis, we only use unital RAs for simplicity.

**Various RAs** Now we introduce various RAs, which are used later in the paper.

We call a positive rational number no greater than 1 a *fraction*. The type of fractions  $\mathbb{Q}_{(0,1]}$  is defined as the subtype  $\{a: \mathbb{Q} / 0 < a \leq 1\}$ .<sup>4</sup> We use the letter  $q$  to represent a fraction.

The *fraction RA*  $\text{FRAC}$  is defined as follows.

$$a: |\text{FRAC}| ::= 0 \mid q \ (q: \mathbb{Q}_{(0,1]}) \mid \zeta \quad a \cdot b := \begin{cases} a + b & a, b \neq \zeta, a + b \leq 1 \\ \zeta & \text{otherwise} \end{cases}$$

$$\varepsilon := 0 \quad \surd a := a \neq \zeta \quad |a| := 0$$

An item of  $\text{FRAC}$  is (i) 0, (ii) a fraction  $q: \mathbb{Q}_{(0,1]}$ , or (iii)  $\zeta$ , the invalid element.

The *exclusive RA*  $\text{EX}(T)$  on the type  $T$  is defined as follows.

$$a: |\text{EX}(T)| ::= \varepsilon \mid \text{ex } x \ (x: T) \mid \zeta \quad a \cdot b := \begin{cases} a & b = \varepsilon \\ b & a = \varepsilon \\ \zeta & \text{otherwise} \end{cases}$$

$$\varepsilon := \varepsilon \quad \surd a := a \neq \zeta \quad |a| := \varepsilon$$

An item of  $\text{EX}(T)$  is (i)  $\varepsilon$ , the unit, (ii)  $\text{ex } x$ , witnessing  $x: T$  exclusively, or (iii)  $\zeta$ , invalid. The exclusive RA admits the following frame-preserving update for any  $x, y: T$ .

$$\text{ex } x \rightsquigarrow \{\text{ex } y\}$$

The *fractional ownership RA*  $\text{FRACOWN}(T)$  on the type  $T$  is defined as follows.

$$a: |\text{FRACOWN}(T)| ::= \varepsilon \mid \text{fown}_q x \ (q: \mathbb{Q}_{(0,1]}, x: T) \mid \zeta$$

$$a \cdot b := \begin{cases} a & b = \varepsilon \\ b & a = \varepsilon \\ \text{fown}_{q+q'} x & a = \text{fown}_q x, b = \text{fown}_{q'} x, q + q' \leq 1 \\ \zeta & \text{otherwise} \end{cases}$$

$$\varepsilon := \varepsilon \quad \surd a := a \neq \zeta \quad |a| := \varepsilon$$

An item of  $\text{FRACOWN}(T)$  is (i)  $\varepsilon$ , the unit, (ii)  $\text{fown}_q x$ , witnessing  $x: T$  with the fraction  $q$ , or (iii)  $\zeta$ , invalid. When we have  $\text{fown}_q x$  and  $\text{fown}_{q'} y$  in a valid way, the sum of fractions  $q + q'$  should be no more than 1 and  $x$  and  $y$  should coincide. This fractional ownership RA admits the following frame-preserving update for any  $x, y: T$ , just like the exclusive RA.

$$\text{fown}_1 x \rightsquigarrow \{\text{fown}_1 y\}$$

Also note that we have the following rules on the fraction RA.

$$\frac{\surd(\text{fown}_q x \cdot \text{fown}_{q'} y)}{x = y \wedge q + q' \leq 1} \quad \frac{q + q' \leq 1}{\text{fown}_{q+q'} x = \text{fown}_q x \cdot \text{fown}_{q'} x}$$

The *product RA*  $\prod_i \text{Cmra}_i$  of an indexed family of RAs  $(\text{Cmra}_i)_i$  is defined point-wise as follows.

$$|\prod_i \text{Cmra}_i| := \prod_i |\text{Cmra}_i|$$

$$a \cdot b := \lambda i. a i \cdot b i \quad \varepsilon := \lambda i. \varepsilon \quad \surd a := \forall i. \surd(a i) \quad |a| := \lambda i. |a i|$$

<sup>4</sup> In this thesis, we use the notation  $\{x: T / \phi_x\}$  for the subtype of  $T$ : Type that consists of any value  $x: T$  satisfying the condition  $\phi_x$ . We implicitly promote a value of the type  $\{x: T / \phi_x\}$  into the type  $T$ .

We can combine the frame-preserving updates of the component RAs in the product RA as follows.

$$\frac{\forall i. a_i \rightsquigarrow B_i}{a \rightsquigarrow \prod_i B_i}$$

The *finite-support map* RA  $T \xrightarrow{\text{fin}} Ra$  on the type  $T$  and the RA  $Ra$  is defined as follows.

$$\begin{aligned} |T \xrightarrow{\text{fin}} Ra| &:= \{ f: T \rightarrow |Ra| \mid \{ x \mid f x \neq \varepsilon \} \text{ is finite} \} & a \cdot b &:= \lambda x. a x \cdot b x \\ \varepsilon &:= \lambda x. \varepsilon & \checkmark a &:= \forall x. \checkmark (a x) & |a| &:= \lambda x. |a x| \end{aligned}$$

An item of  $T \xrightarrow{\text{fin}} Ra$  is a function  $f: T \rightarrow |Ra|$  such that  $f x$  is non-unit for only a finite number of  $x$ . The operations on this RA are defined pointwise. For  $x: T$  and  $a: Ra$ , we write  $[x \leftarrow a]$  for the element of  $T \xrightarrow{\text{fin}} Ra$  that returns  $a$  if the input is  $x$  and returns  $\varepsilon$  otherwise.<sup>5</sup> The finite-support map RA admits the following frame-preserving update, which ‘allocates’ a valid element  $b$  of the target RA  $Ra$  to some fresh place  $x: T$  ( $a\{x \leftarrow b\}$  denotes the map  $a$  with the output for  $x$  replaced with  $b$ ).

$$\frac{\checkmark b}{a \rightsquigarrow \{ a\{x \leftarrow b\} \mid x: T, a x = \varepsilon \}}$$

The RA also admits the following frame-preserving update for modifying the item at one place.

$$\frac{a x \rightsquigarrow B}{a \rightsquigarrow \{ a\{x \leftarrow b\} \mid b \in B \}}$$

**Global Camera** We can also define a step-indexed version of a resource algebra, which is called a *camera*. The definition of the camera can be found in (Jung et al., 2018c, §4.4); in this thesis we omit it. Any resource algebra can be promoted into a *discrete* camera. We use  $Cmra$  to represent a camera.

An Iris proposition is parametrized over some camera, which we call the *global camera* and write as  $\text{GlobCmra}$  in this thesis. The global camera  $\text{GlobCmra}$  is constructed as follows for a finite set of *component cameras*  $\text{CompCmras}$ .<sup>6</sup>

$$\text{GlobCmra} := \prod_{Cmra \in \text{CompCmras}} (\text{GhostName} \xrightarrow{\text{fin}} Cmra)$$

Here,  $\text{GhostName}$  is some infinite type, whose elements are called *ghost names* and are represented by  $\gamma$ . The type of Iris propositions  $\text{IProp}$  is defined based on this global camera  $\text{GlobCmra}$ . Actually, Iris has a mechanism for *impredicativity*, which allows  $\text{GlobCmra}$  to depend on  $\text{IProp}$ , using the fixed point theorem on locally contractive profunctors. For the purpose of this thesis, we don’t need to get deeply into this mechanism.

A proof in Iris can be *extensible* by keeping it parametrized over the set of component cameras  $\text{CompCmras}$ , assuming only that the set contains the cameras we need. We benefit a lot from going *open-world*.

In  $\text{RustHornBelt}$ , we register our RAs to the global camera, i.e., add them to the set  $\text{CompCmras}$  —  $\text{PROPH}$  in §3.3,  $\text{HEAP}$  and  $\text{TIME}$  in §4.2, and  $\text{UNQ}$  in §5.3.

<sup>5</sup> For some  $T$ , we need some classical axioms to construct  $[x \leftarrow a]$ .

<sup>6</sup>The actual Coq implementation of Iris uses a finite sequence of component cameras possibly with duplication in order to avoid equality on cameras.

**Customizing Iris by Cameras/RAs** Iris has the connective  $\text{Own}(a)$  for *owning* the resource  $a$ :  $|\text{GlobCmra}|$  in the global camera. It satisfies the following properties.<sup>7</sup>

$$\begin{array}{c}
\text{OWN-COMP} \\
\text{Own}(a) * \text{Own}(b) = \text{Own}(a \cdot b)
\end{array}
\qquad
\begin{array}{c}
\text{OWN-UNIT} \\
\text{True} = \text{Own}(\varepsilon)
\end{array}
\qquad
\begin{array}{c}
\text{OWN-CORE-PERS} \\
\text{persistent}(\text{Own}(|a|))
\end{array}$$

$$\begin{array}{c}
\text{OWN-VALID} \\
\text{Own}(a) \vdash \checkmark a
\end{array}
\qquad
\begin{array}{c}
\text{OWN-UPD} \\
\frac{a \rightsquigarrow B}{\text{Own}(a) \vdash \exists b \in B. \text{Own}(b)}
\end{array}$$

We can merge and split  $\text{Own}$  according to the composition of the camera (**OWN-COMP**). We always own the unit (**OWN-UNIT**). Owning the core of any item is persistent (**OWN-CORE-PERS**). Owning an item ensures that the item is valid (**OWN-VALID**). When the frame-preserving update  $a \rightsquigarrow B$  holds, consuming  $\text{Own}(a)$  we get  $\text{Own}(b)$  for some  $b$  in  $B$  (**OWN-UPD**).

Importantly, the frame-preserving update  $a \rightsquigarrow B$  we can use for **OWN-UPD** can be extended by registering new cameras to the global camera.

Usually, instead of directly using  $\text{Own}$ , we use the connective  $\llbracket a \rrbracket_{\text{Cmra}}^\gamma$ , which owns the resource  $a$ :  $|\text{Cmra}|$  at the specific ghost name  $\gamma$  and the camera  $\text{Cmra}$ , where  $\text{Cmra}$  is in the set of component cameras  $\text{CompCmras}$ .

$$\llbracket a \rrbracket_{\text{Cmra}}^\gamma := [\text{Cmra} \leftarrow [\gamma \leftarrow a]]$$

It satisfies the following properties, which follows from the properties on  $\text{Own}$ .

$$\begin{array}{c}
\text{GOWN-COMP} \\
\llbracket a \rrbracket_{\text{Cmra}}^\gamma * \llbracket b \rrbracket_{\text{Cmra}}^\gamma = \llbracket a \cdot b \rrbracket_{\text{Cmra}}^\gamma
\end{array}
\qquad
\begin{array}{c}
\text{GOWN-UNIT} \\
\text{True} = \llbracket \varepsilon \rrbracket_{\text{Cmra}}^\gamma
\end{array}
\qquad
\begin{array}{c}
\text{GOWN-CORE-PERS} \\
\text{persistent}(\llbracket |a| \rrbracket_{\text{Cmra}}^\gamma)
\end{array}$$

$$\begin{array}{c}
\text{GOWN-VALID} \\
\llbracket a \rrbracket_{\text{Cmra}}^\gamma \vdash \checkmark a
\end{array}
\qquad
\begin{array}{c}
\text{GOWN-TIMELESS} \\
\text{Cmra is a discrete camera} \\
\hline
\text{timeless}(\llbracket a \rrbracket_{\text{Cmra}}^\gamma)
\end{array}
\qquad
\begin{array}{c}
\text{GOWN-ALLOC} \\
\checkmark a \\
\hline
\text{True} \vdash \exists \gamma. \llbracket a \rrbracket_{\text{Cmra}}^\gamma
\end{array}$$

$$\begin{array}{c}
\text{GOWN-UPD} \\
\frac{a \rightsquigarrow B}{\llbracket a \rrbracket_{\text{Cmra}}^\gamma \vdash \exists b \in B. \llbracket b \rrbracket_{\text{Cmra}}^\gamma}
\end{array}$$

$$\begin{array}{c}
\text{GOWN-UPD-TWO} \\
\frac{(a, a') \rightsquigarrow B \quad \text{Cmra} \neq \text{Cmra}'}{\llbracket a \rrbracket_{\text{Cmra}}^\gamma * \llbracket a' \rrbracket_{\text{Cmra}'}^{\gamma'} \vdash \exists (b, b') \in B. (\llbracket b \rrbracket_{\text{Cmra}}^\gamma * \llbracket b' \rrbracket_{\text{Cmra}'}^{\gamma'})}
\end{array}$$

**GOWN-ALLOC** says that a valid element  $a$  can be allocated to some ghost name  $\gamma$ . **GOWN-UPD-TWO** can discuss interaction of two different cameras  $\text{Cmra}, \text{Cmra}'$  registered to the global camera; the frame preserving update  $(a, a') \rightsquigarrow B$  is discussed in the product camera  $\text{Cmra} \times \text{Cmra}'$ .

## 2.3 Lifetime Logic – Borrowing Propositions Under Lifetimes

The *lifetime logic* was introduced by Jung et al. (2018a,b) to discuss borrowing of resources under *lifetimes* in Rust. The lifetime logic is formalized as a group of definitions and lemmas in the separation logic Iris with Coq mechanization. Using the *impredicativity* of the Iris, the lifetime logic allows us to *borrow any Iris proposition*  $P$ , which is very important for *extensibility* of the proofs with regard to borrowing. Our work, RustHornBelt, depends a lot on the lifetime logic (Chapter 5 and Chapter 6).

<sup>7</sup> The validity predicate  $\checkmark a$  used in **OWN-VALID** is actually step-indexed.

The Iris model of the lifetime logic is highly complicated and subtle, which is well described and explained by Jung (2020, Chapter 11). However, if we just *use* the lifetime logic, we just need to understand the deduction rules of the lifetime logic.

In this chapter, we describe how we can use the lifetime logic. To aid understanding, the formulas and names of the deduction rules are a bit modified from those of (Jung et al., 2018b; Jung, 2020).

**Presenting Deduction Rules in Iris** A deduction rule like the following for an *Iris* proposition  $P$  means that  $P$  is a tautology.

$$\begin{array}{c} \text{SOME-SIMPLE-RULE} \\ P \end{array}$$

A deduction rule like the following for *Iris* propositions  $\vec{Q}, P$  means that  $Q_0 \wedge Q_1 \wedge \dots \wedge Q_{n-1} \Rightarrow P$  is a tautology.

$$\begin{array}{c} \text{SOME-RULE-WITH-ASSUMPTIONS} \\ \frac{Q_0 \quad Q_1 \quad \dots \quad Q_{n-1}}{P} \end{array}$$

To aid understanding, we only use *persistent propositions* for the assumptions  $\vec{Q}$ .

Sometimes we may hide some common assumptions that are persistent. For the lifetime logic, we hide the *lifetime context*  $\text{Ctx}_{\text{lft}}$ , which is an invariant token on some complicated proposition.

**Lifetimes** A *lifetime*  $\alpha, \beta$  is an object of some type  $\text{Lft}$ .

We have the *intersection* operation over lifetimes  $\alpha \sqcap \beta$ , which is associative, commutative, and unital under the *eternity* lifetime  $\infty$ .

$$(\alpha \sqcap \alpha') \sqcap \alpha'' = \alpha \sqcap (\alpha' \sqcap \alpha'') \quad \alpha \sqcap \beta = \beta \sqcap \alpha \quad \infty \sqcap \alpha = \alpha$$

**Lifetime Tokens and Dead-Lifetime Tokens** In order to manage the information about the liveness of lifetimes, we use the lifetime token  $[\alpha]_q$  and the dead lifetime token  $[\dagger\alpha]$ .

The lifetime token  $[\alpha]_q$  witnesses with the fraction  $q$  that the lifetime  $\alpha$  is still alive. The dead-lifetime token  $[\dagger\alpha]$  persistently witnesses that the lifetime  $\alpha$  has been dead.

Both types of tokens are timeless. A dead-lifetime token is persistent.

$$\text{timeless}([\alpha]_q) \quad \text{persistent}([\dagger\alpha]) \quad \text{timeless}([\dagger\alpha])$$

We can introduce a new lifetime  $\alpha$  by the following rule.

$$\begin{array}{c} \text{LFT-INTRO} \\ \text{True} \Rightarrow_{\mathcal{N}_{\text{lft}}} \exists \alpha. [\alpha]_1 * ([\alpha]_1 \Rightarrow_{\mathcal{N}_{\text{lft}}}^* [\dagger\alpha]) \end{array}$$

We get the full lifetime token  $[\alpha]_1$  and the step-taking view shift  $[\alpha]_1 \Rightarrow_{\mathcal{N}_{\text{lft}}}^* [\dagger\alpha]$ . The view shift allows us to end the lifetime one step after we consume the full lifetime token  $[\alpha]_1$ .

A lifetime token can be split and merged according to the fraction.

$$\begin{array}{c} \text{LFTOKEN-FRAC} \\ [\alpha]_{q+q'} \Leftrightarrow [\alpha]_q * [\alpha]_{q'} \end{array}$$

A lifetime cannot be alive and dead at the same time.

$$\frac{\text{LFTOKEN-DEADLFT} \quad [\dagger\alpha]}{[\alpha]_q \Rightarrow \text{False}}$$

The intersection operation and the null lifetime interact with the lifetime and dead-lifetime tokens as follows.

$$\begin{array}{cc} \text{LFTOKEN-INTRSCT} & \text{DEADLFT-INTRSCT} \\ [\alpha \sqcap \beta]_q \Leftrightarrow [\alpha]_q * [\beta]_q & [\dagger(\alpha \sqcap \beta)]_q \Leftrightarrow [\dagger\alpha] \vee [\dagger\beta] \\ \\ \text{LFTOKEN-ETER} & \text{DEADLFT-ETER} \\ [\infty]_q & [\dagger\infty] \Rightarrow \text{False} \end{array}$$

**Lifetime Inclusion** We define *inclusion between lifetimes*  $\alpha \sqsubseteq \beta$  as follows.

$$\alpha \sqsubseteq \beta := (\forall q. [\alpha]_q \Rightarrow_{\mathcal{N}_{\text{fit}}} \exists q'. [\beta]_{q'} * ([\beta]_{q'} \Rightarrow_{\mathcal{N}_{\text{fit}}}^* [\alpha]_q)) \wedge ([\dagger\beta] \Rightarrow_{\mathcal{N}_{\text{fit}}} [\dagger\alpha])$$

It means that (i) out of a fractional lifetime token on  $\alpha$  of any fraction  $q$ , we can temporarily take out a fractional lifetime token on  $\beta$  of some fraction  $q'$ , and (ii) if we have the dead lifetime token on  $\beta$ , we can make the dead lifetime token on  $\alpha$ . Intuitively, it means that the lifetime  $\beta$  lives longer than  $\alpha$ . We say the lifetime  $\beta$  *outlives* the lifetime  $\alpha$  when  $\alpha \sqsubseteq \beta$  holds.

We can use lifetime inclusion by the following rules.

$$\frac{\text{LFTINCL-LFTOKEN} \quad \alpha \sqsubseteq \beta}{[\alpha]_q \Rightarrow_{\mathcal{N}_{\text{fit}}} \exists q'. [\beta]_{q'} * ([\beta]_{q'} \Rightarrow_{\mathcal{N}_{\text{fit}}}^* [\alpha]_q)} \quad \frac{\text{LFTINCL-DEADLFT} \quad \beta \sqsubseteq \alpha \quad [\dagger\alpha]}{\text{True} \Rightarrow_{\mathcal{N}_{\text{fit}}} [\dagger\beta]}$$

We can *derive* the following rules for deducing lifetime inclusion.

$$\begin{array}{ccc} \text{LFTINCL-REFL} & \text{LFTINCL-TRANS} & \\ \alpha \sqsubseteq \alpha & \frac{\alpha \sqsubseteq \alpha' \quad \alpha' \sqsubseteq \alpha''}{\alpha \sqsubseteq \alpha''} & \\ \\ \text{LFTINCL-INTRSCT-L} & \text{LFTINCL-INTRSCT-R} & \text{LFTINCL-ETER} \\ \alpha \sqcap \beta \sqsubseteq \alpha & \frac{\alpha \sqsubseteq \beta \quad \alpha \sqsubseteq \beta'}{\alpha \sqsubseteq \beta \sqcap \beta'} & \alpha \sqsubseteq \infty \end{array}$$

**Full Borrows** A *full borrow*  $\&_{\text{full}}^\alpha P$  fully borrows the *Iris proposition*  $P$ :  $\text{IProp}$  under the lifetime  $\alpha$ . It corresponds to a unique reference in Rust.

For soundness of impredicativity, the full borrow owns  $\triangleright P$  instead of just  $P$ , just like invariants.

We can introduce a full borrow by the following rule.

$$\frac{\text{FULLBOR-INTRO}}{\triangleright P \Rightarrow_{\mathcal{N}_{\text{fit}}} \&_{\text{full}}^\alpha P * ([\dagger\alpha] \Rightarrow_{\mathcal{N}_{\text{fit}}}^* \triangleright P)}$$

By consuming  $\triangleright P$ , we get a full borrow  $\&_{\text{full}}^\alpha P$  and the view shift  $[\dagger\alpha] \Rightarrow_{\mathcal{N}_{\text{fit}}}^* \triangleright P$ . Using the view shift, we can retrieve  $\triangleright P$  after the lifetime  $\alpha$  ends.

We can temporarily access  $\triangleright P$  out of a full borrow  $\&_{\text{full}}^\alpha P$  with help of a fractional lifetime token  $[\alpha]_q$ .

$$\frac{\text{FULLBOR-ACCESS}}{\&_{\text{full}}^\alpha P * [\alpha]_q \Rightarrow_{\mathcal{N}_{\text{fit}}} \triangleright P * (\triangleright P \Rightarrow_{\mathcal{N}_{\text{fit}}}^* \&_{\text{full}}^\alpha P * [\alpha]_q)}$$

The fractional lifetime token  $[\alpha]_q$  is returned after we put back  $\triangleright P$ .

We can also *reborrow* a full borrow  $\&_{\text{full}}^{\beta} P$  from a full borrow  $\&_{\text{full}}^{\alpha} P$  instead of  $\triangleright P$ .

$$\text{FULLBOR-REBOR} \quad \frac{\beta \sqsubseteq \alpha}{\&_{\text{full}}^{\alpha} P \Rightarrow_{\mathcal{N}_{\text{ft}}} \&_{\text{full}}^{\beta} P * ([\dagger\beta] \Rightarrow_{\mathcal{N}_{\text{ft}}} \&_{\text{full}}^{\alpha} P)}$$

It corresponds to reborrowing a unique reference out of a unique reference in Rust.

We can modify the lifetime and the proposition of a full borrow.

$$\begin{array}{c} \text{FULLBOR-MONO-LFT} \\ \beta \sqsubseteq \alpha \\ \hline \&_{\text{full}}^{\alpha} P \Rightarrow \&_{\text{full}}^{\beta} P \end{array} \quad \begin{array}{c} \text{FULLBOR-IFF} \\ \triangleright \square (P \Leftrightarrow Q) \\ \hline \&_{\text{full}}^{\alpha} P \Leftrightarrow \&_{\text{full}}^{\alpha} Q \end{array}$$

Note that a version of **FULLBOR-IFF** with  $\Leftrightarrow$  replaced with one-side implication  $\Rightarrow$  does not hold, because the target proposition  $P$  of a full borrow  $\&_{\text{full}}^{\alpha} P$  represents both what we get from and what we have to put into the full borrow.

We also have various operations on full borrows.

We can split and merge full borrows according to separating conjunction.

$$\begin{array}{c} \text{FULLBOR-SPLIT} \\ \&_{\text{full}}^{\alpha} (P * Q) \Rightarrow_{\mathcal{N}_{\text{ft}}} \&_{\text{full}}^{\alpha} P * \&_{\text{full}}^{\alpha} Q \end{array} \quad \begin{array}{c} \text{FULLBOR-MERGE} \\ \&_{\text{full}}^{\alpha} P * \&_{\text{full}}^{\alpha} Q \Rightarrow_{\mathcal{N}_{\text{ft}}} \&_{\text{full}}^{\alpha} (P * Q) \end{array}$$

We can freeze an internal parameter inside the target proposition of a full borrow, as long as the type  $T$  of the variable is inhabited.

$$\text{FULLBOR-FREEZE} \quad \frac{\exists \_ : T. \text{True}}{\&_{\text{full}}^{\alpha} (\exists a : T. P_a) \Rightarrow_{\mathcal{N}_{\text{ft}}} \exists a : T. \&_{\text{full}}^{\alpha} P_a}$$

A later modality inside a full borrow can be removed in one logical step.

$$\text{FULLBOR-UNLATER} \quad \&_{\text{full}}^{\alpha} (\triangleright P) \Rightarrow_{\mathcal{N}_{\text{ft}}} \&_{\text{full}}^{\alpha} P$$

A full borrow of a full borrow can be unnested into a full borrow with the intersection lifetime in one logical step.

$$\text{FULLBOR-UNNEST} \quad \&_{\text{full}}^{\alpha} \&_{\text{full}}^{\beta} P \Rightarrow_{\mathcal{N}_{\text{ft}}} \&_{\text{full}}^{\alpha \sqcap \beta} P$$

We have the following strong rule for *subdividing* a full borrow  $\&_{\text{full}}^{\alpha} P$  into a new full borrow  $\&_{\text{full}}^{\alpha} Q$ .

$$\text{FULLBOR-SUBDIV} \quad \&_{\text{full}}^{\alpha} P * [\alpha]_q \Rightarrow_{\mathcal{N}_{\text{ft}}} \triangleright P * \forall Q. (\triangleright Q * (\triangleright Q \Rightarrow_{\emptyset} \triangleright P) \Rightarrow_{\mathcal{N}_{\text{ft}}} \&_{\text{full}}^{\alpha} Q * [\alpha]_q)$$

We first consume the full borrow  $\&_{\text{full}}^{\alpha} P$  with help of a fractional lifetime token  $[\alpha]_q$ . We access the content  $\triangleright P$ . Then we can decide the Iris proposition  $Q$ , possibly depending on internal parameters of  $P$ . We have to construct  $\triangleright Q$  and also a view shift  $\triangleright Q \Rightarrow_{\emptyset} \triangleright P$  that turns  $\triangleright Q$  into  $\triangleright P$  under the empty mask  $\emptyset$ . By consuming them, we finally get the new subdivided full borrow  $\&_{\text{full}}^{\alpha} Q$  and also get back  $[\alpha]_q$ .

**Fractured Borrows** A *fractured borrow*  $\&_{\text{frac}}^\alpha \Phi$  is a *persistent* proposition that borrows an Iris proposition parametrized over a fraction  $\Phi: \mathbb{Q}_{(0,1]} \rightarrow \text{IProp}$ .

$$\text{persistent}(\&_{\text{frac}}^\alpha \Phi)$$

We can make a fractured borrow of  $\Phi$  out of a full borrow of  $\Phi(1)$ .

$$\begin{array}{c} \text{FULLBOR-FRACBOR} \\ \&_{\text{full}}^\alpha \Phi(1) \Rightarrow_{\mathcal{N}_{\text{ft}}} \&_{\text{frac}}^\alpha \Phi \end{array}$$

With help of a lifetime token, we can temporarily access the content of a fractured borrow  $\&_{\text{frac}}^\alpha \Phi$ , getting a proposition  $\triangleright \Phi(q)$  for *some* fraction  $q$ .

$$\begin{array}{c} \text{FRACBOR-ACCESS} \\ \&_{\text{frac}}^\alpha \Phi \quad \square (\forall q_0, q_1. \Phi(q_0 + q_1) \Leftrightarrow \Phi(q_0) * \Phi(q_1)) \\ \hline [\alpha]_{q'} \Rightarrow_{\mathcal{N}_{\text{ft}}} \exists q. \triangleright \Phi(q) * (\triangleright \Phi(q) \Rightarrow_{\mathcal{N}_{\text{ft}}}^* [\alpha]_{q'}) \end{array}$$

The second assumption persistently witnesses that  $\Phi$  can be split and merged according to the fraction.

We can modify the lifetime and the target predicate of fractured borrows.

$$\begin{array}{c} \text{FRACBOR-MONO-LFT} \\ \&_{\text{frac}}^\alpha \Phi \quad \beta \sqsubseteq \alpha \\ \hline \&_{\text{frac}}^\beta \Phi \end{array} \qquad \begin{array}{c} \text{FRACBOR-IFF} \\ \triangleright \square (\forall q. \Phi(q) \Leftrightarrow \Phi'(q)) \\ \hline \&_{\text{frac}}^\alpha \Phi \Leftrightarrow \&_{\text{frac}}^\alpha \Phi' \end{array}$$

Also, we can *dynamically create lifetime inclusion*  $\alpha \sqsubseteq \beta$  by introducing a fractured borrow of a fractional lifetime token of  $\beta$  under the lifetime  $\alpha$ .

$$\begin{array}{c} \text{FRACBOR-LFTINCL} \\ \&_{\text{frac}}^\alpha (\lambda q'. [\beta]_{q \cdot q'}) \\ \hline \alpha \sqsubseteq \beta \end{array}$$

Note that we can create the fractured borrow  $\&_{\text{frac}}^\alpha (\lambda q'. [\beta]_{q \cdot q'})$  by the following view shift, which can be derived from [FULLBOR-INTRO](#) and [FULLBOR-FRACBOR](#).

$$[\beta]_q \Rightarrow_{\mathcal{N}_{\text{ft}}} \&_{\text{frac}}^\alpha (\lambda q'. [\beta]_{q \cdot q'}) * ([\dagger \alpha] \Rightarrow_{\mathcal{N}_{\text{ft}}}^* [\beta]_q)$$

## Chapter 3

### A New Formulation of Prophecy

In program verification, our reasoning usually goes forward, along the execution of the program. For some advanced reasoning, however, we need to know and use some information about *what happens in the future* (i.e., events occurring later in the execution). For example, we sometimes need to know in advance what non-deterministic choices will be made in the future. Also, even if the verified program is deterministic, we can often avoid elaborate analysis of the program by peeking out information about the future. To address that need, *prophecy variables* are employed. A prophecy variable is a special variable that conveys information about the future for verification.<sup>1</sup> Roughly speaking, we use a prophecy variable in the following way. We create a prophecy variable  $x$  at some point. From then on, we can use the value of  $x$  for verification to refer to some information about the future. Later at a certain point, after some program execution, we construct some value  $v$  based on information we have at the moment, and we *resolve*  $x$  to  $v$ , i.e., we finally assign to  $x$  the value  $v$ .

The idea of prophecy variables has been long used in various contexts. [Abadi and Lamport \(1988, 1991\)](#) introduced the idea of prophecy variables for a new technique of verifying refinement between state machines. Later, prophecy variables were used by [Sezgin et al. \(2010\)](#) to verify non-interference on concurrent programs and by [Cook and Koskinen \(2011\)](#) and [Lamport and Merz \(2017\)](#) to automatically verify temporal properties of transition systems. [Jung et al. \(2020b\)](#) presented a separation logic with prophecy variables by extending Iris, which was used by them to verify logical atomicity of the RDCSS operation and by [de Vilhena et al. \(2020\)](#) to verify that the local generic solver computes the least fixed point. And as we discussed in §1.3, *RustHorn* by [Matsushita et al. \(2020a\)](#) is a notable application of prophecy variables for verification of imperative programs in Rust.

For our work RustHornBelt, aiming at unifying RustHorn and RustBelt, we newly built a new flexible framework of prophecy in the *separation logic Iris*, which is different in many aspects from existing work. Prophecy variables live only in the *ghost state* and allows resolution of a prophecy variable to a value that *depends on other prophecy variables*. In this chapter, we introduce this new framework of prophecy. In §3.1, we explain the motivation of our framework in light of what we need for RustHornBelt. In §3.2, we present a high-level overview of our framework. In §3.3, we present the formulation of our framework. In §3.4, we discuss related work.

---

<sup>1</sup> Some literature uses the term prophecy variable only for a form of *auxiliary variable*, i.e., a variable employed in *ghost code* that was added to the original program for verification. In this thesis, we use the term prophecy variable in a broader sense to refer to any kind of variable that describes something about the future for verification.

### 3.1 Motivation

What kind of framework of prophecy do we need for our work RustHornBelt? To model Rust in the style of RustHorn on top of the separation logic Iris, we should be able to resolve the prophecy variable of a unique reference into a *value described in the ghost state* modeling the target object of the unique references. Moreover, when we release a unique reference to some object that itself contains other unique references, which we call a *nested unique reference*, we have to resolve its prophecy variable into a value depending on *other prophecy variables* (later we show a detailed example of this), which we call *dependent resolution*. Also, in order to tackle Rust’s elaborate lifetime-based ownership principle in an *extensible* way, we need to build our framework of prophecy in an extensible and powerful separation logic like *Iris*. In particular, it will be great if we can reuse the *lifetime logic* (introduced in §2.3), which is built on top of Iris.

However, none of the existing formulations of prophecy variables were sufficiently flexible and powerful for our project RustHornBelt. The framework by Jung et al. (2020b) targets a OCaml-like *value-based* language, and requires us to add to the original program some ghost code manipulating prophecy variables, in order to observe *physical values* in the execution. This is far from sufficient for modeling Rust’s unique references with prophecies. The proof of soundness and completeness of RustHorn by Matsushita et al. (2020b) tackles this difficulty by focusing on a formalization of a *safe* subset of Rust (not allowing unsafe code) and building a bisimulation between execution of a Rust program and some deduction algorithm (named SLDC resolution) on the CHCs, or *logic formulas*, generated by RustHorn as a model of the program. In the proof, prophecy variables of unique references are encoded as universally quantified *syntactic variables* in a logic formula, which can be specialized later in the deduction algorithm. The use of syntactic variables is justified by soundness and completeness of the deduction algorithm (with respect to the least-fixed-point model of the CHCs). In this way, we can support dependent resolution in a natural way. Although this proof works well in that setting, assuming that verified Rust programs do not have unsafe code, it is rather *fragile*. For the proof, we have to carefully design the transition systems of the formalization of Rust and the deduction algorithm, as well as the CHCs generated for a Rust program, and we also have to depend heavily on syntactic structures of logic formulas. So it was unclear how to extend the proof to support Rust programs with unsafe code for in general.

This is the reason why we decided to try to design a new formulation of prophecy that is sufficient for our work RustHornBelt. Let us discuss more in detail the features we want for our new formulation of prophecy through some examples of Rust code.

To begin with, consider the following very simple code.

```
1 let mut a : i32 = 1; let ua : &mut i32 = &mut a; /* (i) */
2 *ua += 2; /* (ii) */ print!("{}", a); // 3
```

By the unique borrow `&mut a`, we introduce a new prophecy variable  $x$  that represents the value of `a` at the end of the borrow.<sup>2</sup> At (i), `ua` is modeled as  $(1, x)$ , a pair of the current value 1 and the prophecy variable  $x$ , so that  $x$  can be resolved to an appropriate value when `ua` is released. Also, at (i), `a` is modeled as  $x$  so that we can access the value of `a` when we use `a` again. By the operation `*ua += 2`, the model of `ua` is updated into  $(3, x)$ . Then at (ii), we release `ua` and we resolve the prophecy variable  $x$  to the value 3. Now we can end the borrow and retrieve `a`. As we know that  $x$  has been resolved to 3, at last we successfully know that `a` has the value 3.

---

<sup>2</sup> In this thesis, hereinafter, we usually use  $x, y, z$  for prophecy variables and  $a, b, c$  for program variables.

When we handle *nested unique references* (i.e., unique references to an object containing unique references), the situation can be very tricky and we need *dependent resolution*.

*Example 3.1* (Nested Unique Reference). Let us consider the following Rust code.

```

1 let mut a = 0; let mut b = 1;
2 let mut uc = &mut a; *uc += 2; let ub = &mut b; /* (i) */
3 let uuc = &mut uc; /* (ii) */ *uuc = ub; /* (iii) */
4 *uc += 3; /* (iv) */ print!("{}", a, b); // 2, 4

```

Here, `a` and `b` are typed `i32`, `uc` and `ub` are typed `&mut i32`, and `uuc` is typed `&mut &mut i32` — `uuc` is a nested unique reference. We omit lifetime information since it is not relevant to the discussion. Let  $x$ ,  $y$  and  $z$  be the prophecy variables for the unique borrows `uc = &mut a`, `ub = &mut b` and `uuc = &mut uc`. At (i), `uc` is modeled as  $(2, x)$ , `ub` as  $(1, y)$ , `a` as  $x$ , and `b` as  $y$ . At (ii), after the borrow for `uuc`, `uuc` is modeled  $((2, x), z)$  and the model of `uc` is updated into  $z$ .

At (iii), after the update `*uuc = ub`, the model of `uuc` is updated into  $((1, y), z)$ . Since  $(2, x)$  is lost here, we resolve  $x$  to 2. Now, we release the nested unique reference `uuc`, modeled as  $((1, y), z)$ , and thus resolve  $z$  to  $(1, y)$  — *resolution of a prophecy variable  $z$  to a value depending on another prophecy variable  $y$* . This is an instance of *dependent resolution*. Through this resolution, `uc` is now modeled as  $(1, y)$ .

At (iv), the model of `uc` is updated into  $(4, y)$ . Here we release `uc` and resolve  $y$  to 4. Now, since  $x$  and  $y$  have been resolved to 2 and 4, we know that `a` and `b` have the values 2 and 4 at last.

If we allow dependent resolution without restriction, we can accidentally introduce circular dependencies between prophecies, causing unsoundness. For example, if we can resolve  $x$  to  $y + 1$  and  $y$  to  $x + 1$ , we fall into a contradiction. Naturally, we can use the following simple and intuitive restriction: a prophecy variable can only be resolved into a value depending on *other unresolved prophecy variables*. This is a good idea in modeling Rust, because resolution dependencies can be determined in a dynamic way in Rust.

*Example 3.2* (Dynamically Determined Resolution Dependency). For example, suppose the following code, which uses a tricky recursive type `MNat` featuring the unique-reference type:

```

1 enum MNat<'a> { MZero, MSucc( &'a mut MNat<'a> ) } use MNat::*;
2 let mut a = MZero; let mut b = MZero;
3 let ua = &mut a; let ub = &mut b;
4 if rand() { *ua = MSucc(ub); } else { *ub = MSucc(ua); }
5 print!("{}", a, b);

```

Here, we assume some non-deterministic function `fn rand() -> bool`. Let  $x$  and  $y$  be the prophecy variables for `ua` and `ub`. Now `ua` is modeled as  $(MZero, x)$  and `ub` is modeled as  $(MZero, y)$ . If `rand()` returns true, `*ua` is updated into `MSucc(ub)` and `ua` is released; thereby the prophecy variable  $x$  is resolved into `MSucc((MZero, y))`, making  $x$  depend on  $y$ . Otherwise, we do exactly the other way around;  $y$  is resolved into `MSucc((MZero, x))`, making  $y$  depend on  $x$ . Importantly, whether  $x$  depends on  $y$  or  $y$  depends on  $x$  is determined by a *dynamic condition* `rand()`. In particular, the lifetime information is insufficient to capture this kind of dependency, as both `ua` and `ub` have the same lifetime.

### 3.2 Overview of Our Formulation of Prophecy

The difficulty with prophecy is that we *don't know* what happens in the future in a usual sense, especially when the program is non-deterministic. Then how can we model prophecy variables *within the ghost state*?

Intuitively, it seems natural to model prophecies in the following way. Every time we take a new prophecy variable  $x$ , the current world *ramifies* into many *possible worlds* with different values assigned to  $x$ . When we resolve  $x$  to some value  $v$ , we *prune* all the possible worlds that assign to  $x$  any value different from  $v$ . The soundness of this model is guaranteed by the fact that we never lose all the possible worlds, i.e., the constraint over possible worlds is always satisfiable. We can even perform dependent resolution as long as we can ensure satisfiability of the constraint. Although this approach works, ramification and pruning of possible worlds seems to be a *global* operation which conflicts with the *modular, local reasoning of separation logic*.

Instead, we can take a slightly different approach. First, we fix a *closed set of prophecy variables*, which is infinite. Each time we want to introduce a new prophecy variable, we can choose a fresh prophecy variable from this set. Next, we introduce the notion of a *prophecy assignment*  $\pi$ , a map that assigns to each prophecy variable a value of some expected type, representing a possible world. The key idea is that we *retain a set of possible worlds* that remain valid (i.e., have not yet been pruned) in each Hoare triple, instead of choosing one specific possible world for each Hoare triple like Jung et al. (2020b) did. We can share information about the conditions that the valid possible worlds satisfy. Note that, since we have a closed set of possible worlds, we never ramify possible worlds more and only prune them. For prophecy-based reasoning, we pervasively use a *value parametrized over a possible world / prophecy assignment*, which we call a  $\pi$ -*parametrized value* and represent by a letter with a hat like  $\hat{v}$ ,  $\hat{w}$ . A  $\pi$ -parametrized value can be regarded as a value depending on values of some prophecy variables. This framework supports *dependent resolution*: we can resolve a prophecy variable  $x$  to a  $\pi$ -parametrized value  $\hat{v}$  depending only on other unresolved prophecy variables; after the resolution, we know that  $\pi x = \hat{v} \pi$  holds for any valid prophecy assignment  $\pi$ .

We can naturally adapt this idea to the modular reasoning of separation logic, especially Iris. In order to manage resolution of prophecy variables, we use a *prophecy token*  $[x]_q$ , which observes with a fraction  $q$  that the prophecy variable  $x$  has not yet been resolved. When we introduce a fresh prophecy variable  $x$ , we get a full prophecy token  $[x]_1$ . We also use a persistent token  $\langle \pi. \phi_\pi \rangle$  called a *prophecy observation*, which means that a pure proposition  $\phi_\pi$  holds for any valid prophecy assignment  $\pi$ . In the notation  $\langle \pi. \phi_\pi \rangle$ , the first part  $\pi$  binds the variable  $\pi$  and the second part  $\phi_\pi$  describes a condition on  $\pi$ . Now, dependent resolution is supported in the following way: we can resolve  $x$  into a  $\pi$ -parametrized value  $\hat{v}$  to get a prophecy observation  $\langle \pi. \pi x = \hat{v} \pi \rangle$ , by consuming the full prophecy token  $[x]_1$  and temporarily using a partial prophecy token  $[y]_q$  (which ensures that  $y$  is unresolved and different from  $x$ ) for every dependency  $y$  of  $\hat{v}$ . We can modify the prophecy observations acquired through resolution by merging (i.e., getting  $\langle \pi. \phi_\pi \wedge \psi_\pi \rangle$  out of  $\langle \pi. \phi_\pi \rangle$  and  $\langle \pi. \psi_\pi \rangle$ ) and weakening (i.e., getting  $\langle \pi. \phi'_\pi \rangle$  out of  $\langle \pi. \phi_\pi \rangle$  and  $\forall \pi. \phi_\pi \Rightarrow \phi'_\pi$ ). We use the following lemma to link a prophecy observation to the ‘reality’: when we have  $\langle \pi. \hat{\phi} \pi \rangle$ , there exists some prophecy assignment  $\pi_*$  satisfying  $\hat{\phi} \pi_*$ . This lemma intuitively means that there always remains some possible world that has not been pruned. Notably, this formulation of prophecies is completely independent of how we define the weakest precondition predicate or the Hoare triple.

This very simple formulation of prophecies is flexible enough to model unique references in Rust. Each variable of Rust is now modeled as a  $\pi$ -*parametrized value*. Each

unique reference is equipped with the full prophecy token  $[x]_1$  of its prophecy variable  $x$ .

*Example 3.3* (Nested Unique Reference, Revisited). To see how our framework works, let us reconsider the Rust program of [Example 3.1](#), an example of using a nested unique reference.

```

1 let mut a = 0; let mut b = 1;
2 let mut uc = &mut a; *uc += 2; let ub = &mut b; /* (i) */
3 let uuc = &mut uc; /* (ii) */ *uuc = ub; /* (iii) */
4 *uc += 3; /* (iv) */ println!("{}", a, b); // 2, 4

```

On the unique borrows `&mut a` and `&mut b`, we create full prophecy tokens  $[x]_1$  and  $[y]_1$  for fresh prophecy variables  $x$  and  $y$ , which are given to the unique references `uc` and `ub`. At (i), `uc` is modeled as a  $\pi$ -parametrized value  $\lambda\pi.(2, \pi x)$ , `ub` as  $\lambda\pi.(1, \pi y)$ , `a` as  $\lambda\pi.\pi x$  and `b` as  $\lambda\pi.\pi y$ .

At (ii), letting  $z$  be the prophecy variable for the borrow `&mut uc`, `uc` is modeled as  $\lambda\pi.\pi z$  and `uuc` as  $\lambda\pi.((2, \pi x), \pi z)$ . Through the update `*uuc = ub`, consuming the full prophecy token  $[x]_1$  previously owned by `*uuc`, we *resolve*  $x$  into 2 to get a prophecy observation  $\langle\pi.\pi x = 2\rangle$ .

At (iii), `uuc` is modeled as  $\lambda\pi.((1, \pi y), \pi z)$ , and we release `uuc` here. We *resolve*  $z$  into  $(1, \pi y)$  by consuming  $[z]_1$  owned by `uuc` and temporarily accessing  $[y]_1$  owned by `*uuc`, and get  $\langle\pi.\pi z = (1, \pi y)\rangle$ . This is *dependent resolution*. We temporarily use  $[y]_1$  to ensure that  $y$  has not been resolved.

When we access again `uc` in the line 4, we know that `uc` is modeled as *some*  $\pi$ -parametrized value  $\hat{\nu}$  satisfying  $\langle\pi.\hat{\nu}\pi = \pi z\rangle$ . The type of `uc` tells us that  $\hat{\nu}$  is of the form  $\lambda\pi.(n, \pi y')$  for some integer  $n$  and prophecy variable  $y'$ . We get  $\langle\pi.(n, \pi y') = (1, \pi y)\rangle$  from the two previously introduced prophecy observations, which entails  $n = 1$  and  $\langle\pi.\pi y' = \pi y\rangle$ . Here, we got  $n = 1$  from  $\langle\pi.n = 1\rangle$ . Although we don't obtain  $y = y'$  here, it suffices to just have  $\langle\pi.\pi y' = \pi y\rangle$  for verification.

At (iv), after the update `*uc += 3`, `uc` is modeled as  $\lambda\pi.(4, \pi y')$ , and by releasing `uc` we get  $\langle\pi.\pi y' = 3\rangle$ . Now, from the prophecy observations we have acquired, we get  $\langle\pi.\pi x = 2 \wedge \pi y = 4\rangle$ . Therefore, we successfully know that the values of `a` and `b` are set to 2 and 4 at last.

### 3.3 Technical Details

Now we discuss technical details of our formulation of prophecy variables described in the previous section. Although the idea was originally conceived for our project `RustHornBelt`, this formulation is designed so that it can be employed for a general purpose. We assume that we work on a dependent type system like `Coq`, but our notation does not strictly follow the style of `Coq`.

**Basic Concepts** A *prophecy variable*  $x, y, z$  is modeled as an object of the record type

$$\text{ProphVar} := (\text{discr}: \mathbb{N}, \text{type}: \text{InhType}),$$

consisting of (`discr`) the discriminator number and (`type`) the type of a value assigned to the variable. Here, `InhType` denotes the type of *inhabited* types and is defined as the subtype  $\{T: \text{Type} / \exists \_ : T. \text{True}\}$ .<sup>3</sup> Notably, for the value type of a prophecy variable,

<sup>3</sup> Note that `Coq` prevents unsoundness about a ‘type of types’ by having a *hierarchy of universes*. The type `Type` actually has an implicit parameter  $i$  representing some universe, being of the form `Typei`. For any  $i$ , the type `Typei` belongs to the type `Typei+1`, but not to `Typei`. In the definition of `InhType`, `Type` is in fact `Typei` for some universe level  $i$ .

we can use any inhabited type; we can use especially a predicate type  $T \rightarrow \text{Prop}$ , which is essential for modeling Rust's function type in RustHornBelt (see §6.1.4).

A *prophecy assignment*  $\pi$  is an object of the dependent function type

$$\text{ProphAsn} := (x: \text{ProphVar}) \rightarrow x.\text{type},$$

which assigns to each prophecy variable some value of the expected type. A prophecy assignment can be regarded as a representation of a possible world. Since for any prophecy variable  $x$  the type  $x.\text{type}$  is inhabited, it is guaranteed that there exists some prophecy assignment (i.e.,  $\exists \pi: \text{ProphAsn}.$  True holds),<sup>4</sup> which is a vital property for this framework (and is later used for the proof of [Theorem 3.1](#)).

A  $\pi$ -*parametrized value*  $\hat{v}, \hat{w}$  over a type  $T$  is a map of the type  $\text{ProphAsn} \rightarrow T$ , which is an important machinery in this framework. In particular, later in the formulation of RustHornBelt, each object in Rust is associated with a  $\pi$ -parametrized value (see [Chapter 5](#)).

A key notion about  $\pi$ -parametrized values is the relation  $\text{Dep}(\hat{v}, X)$ , meaning that a  $\pi$ -parametrized value  $\hat{v}: \text{ProphAsn} \rightarrow T$  depends only on values assigned to the prophecy variables in the set  $X: \mathbb{P} \text{ProphVar}$ . For example, the following holds.

$$\text{Dep}(\lambda \pi. \text{inj}_0((1, \pi x), \pi y), \{x, y\})$$

Let us define this relation. First, we introduce an equivalence relation on prophecy assignments  $\pi \equiv_X \pi'$ , meaning that  $\pi$  and  $\pi'$  assign exactly the same value to every prophecy variable in the set  $X$ . It is simply defined as follows.

$$\pi \equiv_X \pi' := \forall x \in X. \pi x = \pi' x$$

Each equivalence class on  $(\equiv_X)$  is characterized exactly by what values are assigned to the prophecy variables in  $X$ . Therefore,  $\text{Dep}(\hat{v}, X)$  can be defined as the property that  $\hat{v}$  stays constant within each equivalence class on  $(\equiv_X)$ , which is formally described as follows.

$$\text{Dep}(\hat{v}, X) := \forall \pi, \pi' \text{ s.t. } \pi \equiv_X \pi'. \hat{v} \pi = \hat{v} \pi'$$

We have the following lemmas on the relation  $\text{Dep}(\hat{v}, X)$ .

$$\begin{array}{c} \text{DEP-ONE} \\ \text{Dep}(\lambda \pi. \pi x, \{x\}) \end{array} \quad \frac{\text{DEP-CONSTRUCT} \quad \forall i. \text{Dep}(\hat{v}_i, X_i) \quad f \text{ does not depend on } \pi}{\text{Dep}(\lambda \pi. f(\vec{\hat{v}} \pi), \bigcup_i X_i)}$$

$$\frac{\text{DEP-DESTRUCT} \quad \text{Dep}(\lambda \pi. f(\vec{\hat{v}} \pi), X) \quad f \text{ is injective}}{\text{Dep}(\hat{v}_i, X)}$$

The basic  $\pi$ -parametrized value  $\lambda \pi. \pi x$  depends only on  $x$  ([DEP-ONE](#)). [DEP-CONSTRUCT](#) says that, when we have  $\pi$ -parametrized values  $\vec{\hat{v}}$  depending only on  $\vec{X}$  respectively, the  $\pi$ -parametrized value constructed by applying a function  $f$  (which does not depend on  $\pi$ ) to them depends only on the union of the sets  $\vec{X}$ . Conversely, [DEP-DESTRUCT](#) says that, if the  $\pi$ -parametrized value constructed by applying an injective function  $f$  (which does not depend on  $\pi$ ) to  $\vec{\hat{v}}$  depends only on  $X$ , then each  $\hat{v}_i$  depends only on  $X$ .

<sup>4</sup> In Coq, we use the functional choice axiom to prove the existence.

**Prophecy Log** To manage resolution of prophecy variables, we introduce the notion of a *prophecy log*  $\mathcal{P}$ . It is a list  $[(x_0, \hat{v}_0), \dots, (x_{n-1}, \hat{v}_{n-1})]$  which records the resolutions that have been performed in reverse chronological order; each item  $(x_i, \hat{v}_i)$  means that a prophecy variable  $x_i$  has been resolved to a  $\pi$ -parametrized value  $\hat{v}_i$ . When we resolve a prophecy variable  $x$  to a  $\pi$ -parametrized value  $\hat{v}$ , we update the prophecy log  $\mathcal{P}$  into  $(x, \hat{v}) :: \mathcal{P}$ . Formally, each item of a prophecy log is of the dependent pair type

$$\text{ProphLogItem} := (x: \text{ProphVar}) \times (\text{ProphAsn} \rightarrow x.\text{type})$$

and a prophecy log has the type

$$\text{ProphLog} := \text{List ProphLogItem}.$$

The set of resolutions in a prophecy log  $\mathcal{P}$  forms a constraint over prophecy assignments. We use the following relation  $\pi \triangleleft \mathcal{P}$  to describe the constraint.

$$\pi \triangleleft \mathcal{P} := \forall (x, \hat{v}) \text{ in } \mathcal{P}. \pi x = \hat{v} \pi$$

It is essential to always make sure that the prophecy log  $\mathcal{P}$  is satisfiable, i.e., there exists some prophecy assignment  $\pi_*$  that satisfies  $\pi_* \triangleleft \mathcal{P}$ . We can achieve that by ensuring the following property: each time we resolve a prophecy variable  $x$  to a  $\pi$ -parametrized value  $\hat{v}$ ,  $x$  has never been resolved before that and  $\hat{v}$  depends only on other unresolved prophecy variables. Formally, we model this idea as the following *validity* condition  $\checkmark \mathcal{P}$  over a prophecy log  $\mathcal{P}$ .

$$\checkmark((x, \hat{v}) :: \mathcal{P}) := x \in \text{unres}(\mathcal{P}) \wedge \text{Dep}(\hat{v}, \text{unres}(\mathcal{P}) - \{x\}) \wedge \checkmark \mathcal{P} \quad \checkmark \text{nil} := \text{True}$$

Here,  $\text{unres}(\mathcal{P})$  denotes the set of prophecy variables that have not been resolved by  $\mathcal{P}$ , i.e., the complement of  $\{x \mid \exists \hat{v}. (x, \hat{v}) \text{ in } \mathcal{P}\}$ . As we intended, we achieve the following theorem.

**Theorem 3.1** (Satisfiability of a Valid Prophecy Log). *For any prophecy log  $\mathcal{P}_*$  that satisfies  $\checkmark \mathcal{P}_*$ , there exists a prophecy assignment  $\pi_*$  that satisfies  $\pi_* \triangleleft \mathcal{P}_*$ .*

*Proof.* We introduce the following function  $\text{modify}_{\mathcal{P}} \pi$  that modifies a prophecy assignment  $\pi$  based on the information of  $\mathcal{P}$ .

$$\text{modify}_{(x, \hat{v}) :: \mathcal{P}} \pi := \text{modify}_{\mathcal{P}}(\pi \{x \leftarrow \hat{v} \pi\}) \quad \text{modify}_{\text{nil}} \pi := \pi$$

Here,  $\pi \{x \leftarrow \hat{v} \pi\}$  is a prophecy assignment that assigns  $\hat{v} \pi$  to  $x$  and  $\pi y$  to any  $y \neq x$ . We inductively show  $\text{modify}_{\mathcal{P}} \pi \equiv_{\text{unres}(\mathcal{P})} \pi$  for any  $\mathcal{P}$ . Using that, we inductively show  $\text{modify}_{\mathcal{P}} \pi \triangleleft \mathcal{P}$  for any  $\mathcal{P}$  satisfying  $\checkmark \mathcal{P}$ . By taking some prophecy assignment  $\pi_0$ , we set  $\pi_* := \text{modify}_{\mathcal{P}_*}(\pi_0)$ .<sup>5</sup>  $\square$

**Machinery for Prophecy on Iris** Now, using the notions formulated so far, we build the machinery for prophecy upon the separation logic Iris.

First, we introduce the following RA  $\text{PROPH}$ , named the *prophecy RA*.<sup>6</sup>

$$|\text{PROPH}| := |\text{EX}(\text{ProphLog})| \times (\text{ProphAsn} \rightarrow \text{Prop}) \times |\text{ProphVar} \xrightarrow{\text{fin}} \text{FRAC}|$$

$$(a, \hat{\phi}, f) \cdot (b, \hat{\psi}, g) := (a \cdot b, \lambda \pi. \hat{\phi} \pi \wedge \hat{\psi} \pi, f \cdot g)$$

<sup>5</sup> We need classical axioms to discuss equality over Type.

<sup>6</sup> We can use  $\mathbb{N} \xrightarrow{\text{fin}} \text{InhType} \rightarrow \text{FRAC}$  instead of  $\text{ProphVar} \xrightarrow{\text{fin}} \text{FRAC}$  for this RA to simplify formalization about finiteness.

$$\varepsilon := (\varepsilon, \lambda\pi.\text{True}, \varepsilon) \quad |(a, \hat{\phi}, f)| := (\varepsilon, \hat{\phi}, \varepsilon)$$

$$\begin{aligned} \checkmark(a, \hat{\phi}, f) &:= \checkmark a \wedge (\exists \pi. \hat{\phi} \pi) \wedge \checkmark f \wedge \\ &\quad \forall \mathcal{P} \text{ s.t. } \text{ex } \mathcal{P} = a. \checkmark \mathcal{P} \wedge (\forall \pi \triangleleft \mathcal{P}. \hat{\phi} \pi) \wedge (\forall x \notin \text{unres}(\mathcal{P}). f x = 0) \end{aligned}$$

An item  $(a, \hat{\phi}, f)$  of this RA consists of (i) an item  $a: |\text{EX}(\text{ProphLog})|$  (later used for the *prophesy context*), (ii) a known condition on valid prophecy assignments  $\hat{\phi}: \text{ProphAsn} \rightarrow \text{Prop}$  (later used for *prophesy observations*), and (iii) a finite-support map from prophecy variables to extended fractions  $f: |\text{ProphVar}| \xrightarrow{\text{fin}} \text{FRAC}$  that manages unresolved prophecy variables (later used for *prophesy tokens*).

For definition, we piggyback the unit elements  $\varepsilon$  and the compositions  $(\cdot)$  of the RAs  $\text{EX}(\text{ProphLog})$  and  $\text{ProphVar} \xrightarrow{\text{fin}} \text{FRAC}$ . The composition of the prophecy RA is defined component-wise; for the  $\text{ProphAsn} \rightarrow \text{Prop}$  part, we take the (pointwise) logical conjunction. Since logical conjunction is idempotent, the core  $|(a, \hat{\phi}, f)|$  can be set to  $(\varepsilon, \hat{\phi}, \varepsilon)$  (later this makes prophecy observations persistent).

The validity predicate  $\checkmark$  of this RA is important. An item  $(\text{ex } \mathcal{P}, \hat{\phi}, f)$  is valid when (i)  $\mathcal{P}$  is valid, (ii)  $\hat{\phi} \pi$  holds for any  $\pi$  satisfying  $\mathcal{P}$  (which entails satisfiability of  $\hat{\phi}$  by [Theorem 3.1](#)), and (iii)  $f$  is valid and  $f x = 0$  holds for any  $x$  resolved by  $\mathcal{P}$ . An item  $(\varepsilon, \hat{\phi}, \varepsilon)$  is valid when  $\hat{\phi}$  is satisfiable and  $f$  is valid.

Now we register the prophecy RA  $\text{PROPH}$  as a component of the global camera and take some fresh namespace  $\mathcal{N}_{\text{proph}}$ . We define the following *prophesy context*  $\text{Ctx}_{\text{proph}}$ .

$$\text{Ctx}_{\text{proph}} := \boxed{\exists \mathcal{P}. \boxed{(\text{ex } \mathcal{P}, \lambda\pi.\text{True}, \varepsilon)}_{\text{PROPH}}^{\gamma_{\text{proph}}}}^{\mathcal{N}_{\text{proph}}}$$

It is a persistent token that refers to an invariant in  $\mathcal{N}_{\text{proph}}$  that governs the prophecy log  $\mathcal{P}$  at a ghost name  $\gamma_{\text{proph}}$ . We need this context to resolve prophecy variables (see the lemma [PROPH-RESOLVE](#)). The ghost name  $\gamma_{\text{proph}}$  is chosen when we introduce the context  $\text{Ctx}_{\text{proph}}$  by the following lemma.

$$\begin{array}{c} \text{PROPH-CTX-INTRO} \\ \text{True} \Rightarrow_{\emptyset} \exists \gamma_{\text{proph}}. \text{Ctx}_{\text{proph}} \end{array}$$

Since  $\text{Ctx}_{\text{proph}}$  is persistent, later on we implicitly assume  $\text{Ctx}_{\text{proph}}$ .

After  $\gamma_{\text{proph}}$  is fixed, we can define the following *prophesy observation*  $\langle \pi. \phi_{\pi} \rangle$  and *prophesy token*  $[x]_q$ .

$$\langle \pi. \phi_{\pi} \rangle := \boxed{(\varepsilon, \lambda\pi.\phi_{\pi}, \varepsilon)}_{\text{PROPH}}^{\gamma_{\text{proph}}} \quad [x]_q := \boxed{(\varepsilon, \lambda\pi.\text{True}, [x \leftarrow q])}_{\text{PROPH}}^{\gamma_{\text{proph}}}$$

A prophecy observation  $\langle \pi. \phi_{\pi} \rangle$  persistently witnesses that the pure proposition  $\phi_{\pi}$  holds for any valid prophecy assignment  $\pi$ . A prophecy token  $[x]_q$  witnesses with the fraction  $q: \mathbb{Q}_{(0,1]}$  that the prophecy variable  $x$  has not been resolved. We also use the shorthand  $[X]_q := \ast_{x \in X} [x]_q$  for  $X: \mathbb{P}_{\text{fin}} \text{ProphVar}$ .

The following properties hold.

$$\text{persistent}(\langle \pi. \phi_{\pi} \rangle) \quad \text{timeless}(\langle \pi. \phi_{\pi} \rangle) \quad \text{timeless}([x]_q)$$

$$\begin{array}{c} \text{PROPH-TOKEN-INTRO} \\ \exists \_ : T. \text{True} \\ \hline \text{True} \Rightarrow_{\emptyset} \exists x \text{ s.t. } x.\text{type} = T. [x]_1 \end{array}$$

$$\begin{array}{c} \text{PROPH-TOKEN-FRAC} \\ [x]_{q+q'} \Leftrightarrow [x]_q \ast [x]_{q'} \end{array}$$

$$\begin{array}{c} \text{PROPH-RESOLVE} \\ \text{Dep}(\hat{v}, Y) \\ \hline [x]_1 \ast [Y]_q \Rightarrow_{\mathcal{N}_{\text{proph}}} \langle \pi. \pi x = \hat{v} \pi \rangle \ast [Y]_q \end{array}$$

$$\begin{array}{c} \text{PROPH-OBS-MERGE} \\ \langle \pi. \phi_{\pi} \rangle \quad \langle \pi. \psi_{\pi} \rangle \\ \hline \langle \pi. \phi_{\pi} \wedge \psi_{\pi} \rangle \end{array}$$

$$\begin{array}{ccc}
\text{PROPHOBS-WKN} & \text{PROPHOBS-SAT} & \text{PROPHOBS-FACT} \\
\frac{\forall \pi. \phi_\pi \Rightarrow \psi_\pi \quad \langle \pi. \phi_\pi \rangle}{\langle \pi. \psi_\pi \rangle} & \frac{\langle \pi. \hat{\phi} \pi \rangle}{\exists \pi_*. \hat{\phi} \pi_*} & \frac{\langle \pi. \phi \rangle}{\phi}
\end{array}$$

**PROPHOKEN-INTRO** says that we can always take a fresh prophecy variable  $x$  of any expected (inhabited) type  $T$  and get a full prophecy token  $[x]_1$ . A prophecy token can be split and merged according to the fraction (**PROPHOKEN-FRAC**). **PROPH-RESOLVE** says that, by consuming a full prophecy token of  $x$ , as long as we have a partial prophecy token on the prophecy variables that a  $\pi$ -parametrized value  $\hat{v}$  may depend on, we can resolve  $x$  to  $\hat{v}$  and obtain a prophecy observation  $\pi x = \hat{v} \pi$  (note that we assume here that  $\hat{v}$  has the type  $\text{ProphAsn} \rightarrow x.\text{type}$ ). Prophecy observations can be merged (**PROPHOBS-MERGE**) and weakened (**PROPHOBS-WKN**). **PROPHOBS-SAT** says that a prophecy observation  $\langle \pi. \hat{\phi} \pi \rangle$  implies that there exists a prophecy assignment  $\pi_*$  satisfying the predicate  $\hat{\phi}$ . **PROPHOBS-FACT** says that having a prophecy observation  $\langle \pi. \phi \rangle$ , where  $\phi$  does not depend on  $\pi$ , is the same thing as knowing that  $\phi$  holds.

*Proof of **PROPHOKEN-INTRO**.* Because only a finite number of prophecy variables have been *used* in the global resource, i.e., resolved by the prophecy log or included in the support of the map  $\text{ProphVar} \xrightarrow{\text{fin}} \text{FRAC}$ , we can always find a prophecy variable that has not been used yet.  $\square$

*Proof of **PROPH-RESOLVE**.* We update the prophecy log in the invariant from  $\mathcal{P}$  to  $(x, \hat{v}) :: \mathcal{P}$ . The validity condition on the prophecy log is retained, because the prophecy tokens  $[x]_1$  and  $[Y]_q$  ensure that  $x$  has not been resolved, that each  $y$  in  $Y$  has not been resolved, and that  $x$  is not included in  $Y$ . Also, by consuming the full token  $[x]_1$ , we are allowed to switch the state of  $x$  from unresolved to resolved. Because the constraint on a prophecy assignment  $\pi$  denoted by the prophecy log just changes from  $\pi \triangleleft \mathcal{P}$  to  $\pi x = \hat{v} \pi \wedge \pi \triangleleft \mathcal{P}$ , we can obtain the prophecy observation  $\langle \pi. \pi x = \hat{v} \pi \rangle$ .  $\square$

*Proof of **PROPHOBS-SAT**.* By the definition of the validity predicate  $\checkmark$  of **PROPH**.  $\square$

*Proof of **PROPHOBS-FACT**.* The forward implication follows by **PROPHOBS-SAT**. The backward implication follows by the definition of  $\varepsilon$  in **PROPH**.  $\square$

### 3.4 Related Work

The idea of prophecy variables was first introduced in the theoretical work by **Abadi and Lamport (1988, 1991)** as the ‘mirror image’ of history variables (auxiliary variables describing the past events), for a new technique of verifying refinement between state machines. In general, we can verify that a state machine  $S_1$  refines  $S_2$  (i.e., the observable behaviors of  $S_1$  are a subset of those of  $S_2$ ) by finding a refinement mapping from  $S_1$  to  $S_2$  (i.e., a behavior-preserving map from  $S_1$ ’s states to  $S_2$ ’s states). We can get a more powerful verification method by using history and prophecy variables. A state machine  $S_1^*$  obtained by adding history and prophecy variables to  $S_1$  has the same traces with  $S_1$ , and existence of a refinement mapping from  $S_1^*$  to  $S_2$  entails refinement of  $S_2$  by  $S_1$ . Actually, this new method is complete: under some conditions, if  $S_1$  refines  $S_2$ , there exist a state machine  $S_1^*$ , obtained by adding some history and prophecy variables to  $S_1$ , and a refinement mapping from  $S_1^*$  to  $S_2$ . Here, prophecies are employed to know in advance what non-deterministic choices  $S_2$  makes in the future. This completeness theorem is the main theoretical result of their work. Note that in their verification method they make a modified state machine  $S_1^*$  with prophecy variables, which can be understood as addition of some prophecy-related ghost code to programs.

The studies of Vafeiadis (2008, §5.3.3), Zhang et al. (2012) and Jung et al. (2020b) employed prophecy variables for verification in separation logic. At a high level, all of them basically took the same approach. They added to the program some *ghost code* that introduces prophecy variables as *physical* variables and resolves them by information observed *physically*. They performed Hoare-style verification over the modified program with the prophecy-handling ghost code, letting each Hoare triple choose *one specific possible world* on prophecies (in contrast to our approach of retaining a *set of possible worlds*). Also, they separately proved the *erasure theorem*, saying that wished properties on the original program can be ensured by verifying the modified program with ghost code. As a key use case, all of them targeted verification of linearizability of the RDCSS (restricted double-compare single-swap) operation proposed by Harris et al. (2002). We say a concurrent operation is *linearizable* if it has a single physical step called the *linearization point* where the relevant part of the *ghost state* is *atomically* updated, even though the operation may not be physically atomic, i.e., may take multiple physical steps. Verifying linearizability of the RDCSS operation naturally requires prophecies because the linearization point depends on information about the future (specifically, whether the current thread wins the race to perform some completion operation). Because atomicity over the ghost state, or *logical atomicity*, is technically subtle, their approach of adding prophecy-handling ghost code is fairly reasonable. In fact, Jung et al. (2020b) successfully verified a very strong version of linearizability for the RDCSS operation. Still, we did not take their approach for our project RustHornBelt, because resolution of the prophecy variables of unique references actually depends on information about the *ghost state*, rather than the physical state, in the presence of unsafe code. If we were to take their approach for our purpose, we would have to prove our version of the *erasure theorem* with regard to the ghost state, which seems highly non-trivial.

The idea of retaining a set of possibilities about the future has been employed in the context of separation logic by Turon et al. (2013), in the name of *speculation*. They used speculation to verify linearizability (via contextual refinement) of a simplified version of the RDCSS operation, called conditional increment. Later, Liang and Feng (2013) extended the idea to verify the original version of RDCSS. Both studies took basically the same approach, which is quite different from ours. Each separation-logic proposition  $P$  is parametrized over a *set of speculative states*  $\Sigma$ , which is non-empty and *finite*. They use the *speculative choice*  $P \oplus Q$  over propositions  $P$  and  $Q$ , which says that the set of speculative states  $\Sigma$  is the union of some  $\Sigma_0$  satisfying  $P$  and  $\Sigma_1$  satisfying  $Q$ . When they perform some speculation about the future, where each case corresponds to  $P_0, \dots, P_n$ , they obtain  $P_0 \oplus \dots \oplus P_n$ . When an expression  $e$  can turn  $P$  into  $P'$  and  $Q$  into  $Q'$ ,  $e$  can turn  $P \oplus Q$  into  $P' \oplus Q'$ . When they have  $P \oplus Q$  and (physically) observe that the speculation of  $P$  is correct, they can get  $P$  out of  $P \oplus Q$ . They successfully used the mechanism to do case analysis for verifying linearizability of conditional increment or RDCSS. However, in their frameworks they could not do speculation on an unbounded integer value, for example, which introduces an *infinite* number of possibilities, especially because of the nature of the speculative choice ( $\oplus$ ) operator. Therefore, their frameworks were unsuitable for our project RustHornBelt.

Matsushita et al. (2020b) proved soundness and completeness of RustHorn’s reduction from Rust programs to CHCs by establishing bisimulation between the execution of a Rust program and the course of some deduction algorithm, named SLDC resolution, on the CHCs generated by RustHorn. SLDC resolution was specially designed for this work, extending the idea of SLD resolution proposed by Kowalski (1974). In general, a resolution<sup>7</sup> algorithm over CHCs can be understood as *top-down* construction (i.e.,

---

<sup>7</sup> Resolution over CHCs is not quite related to *resolution* over prophecy variables.

construction from the goal/root) of a proof tree on CHCs, where each CHC is regarded as a deduction rule. In the SLDC resolution algorithm, intermediate states can contain some *syntactic logic variables*, which can model the *prophecy variables* of unique references in Rust. Because these logic variables are universally quantified semantically, each intermediate state of SLDC resolution virtually handles a *set of possible worlds* that is parametrized over values assigned to the variables. Their proof is rather fragile, because the formalized execution of Rust programs and the SLDC resolution algorithm are carefully *designed* to form bisimulation and SLDC resolution depends heavily on *syntactic* structures of logic formulas. Also, their proof omits detailed discussion about locality of the effects of various Rust operations, without formal machinery like separation logic. Our approach to prophecies for RustHornBelt can be understood as a *semantic* reformulation of their syntactic approach built upon the *mechanized* separation logic Iris.

# Chapter 4

## A Low-level Foundation for Verification

Before building the verification platform specifically for Rust, we introduce a simple lambda calculus for imperative programming and build a low-level program logic on the calculus on top of the separation logic Iris. The calculus supports heap manipulation, concurrency and non-determinism. Although our program logic serves as a low-level foundation for verifying Rust programs later in [Chapter 5](#) and [Chapter 6](#), the logic is designed for a general purpose.

Although the construction of the program logic in Iris largely follows the standard technique, we present a new technique for spending an *unbounded number of logical steps* for each physical step in the *finite*-step-indexed logic Iris, which is technically important for our work.

In [§4.1](#), we present our lambda calculus for imperative programming. In [§4.2](#), we present our program logic for the calculus, with the new technique for spending an unbounded number of logical steps at a time. In [§4.3](#) we discuss some related work.

### 4.1 A Lambda Calculus for Imperative Programming

We introduce a simple *lambda calculus* for imperative programming, which models the *unsafe* part of Rust. To ease verification, we use this *substitution-based* calculus, just like RustBelt ([Jung et al., 2018a](#)) did (see also ([Jung, 2020, §7.3](#))).<sup>1</sup> Advanced operations can be represented as *syntax sugar* in this calculus. The so-called *safe part* of Rust can be understood as some subset of this calculus.

**Syntax** An *address*  $l^2$  is an object of the record type

$$\text{Addr} := (\text{block}: \mathbb{Z}, \text{cell}: \mathbb{Z}),$$

consisting of (block) the id of the *memory block* and (cell) the index of a *memory cell*. (These notions get clearer when we formalize the heap memory later in the paragraph [Operational Semantics](#).) We write  $\text{addr}(n, i)$  for the address whose block field equals  $n$  and whose cell field equals  $i$ . For  $l: \text{Addr}$  and  $n: \mathbb{Z}$ , we define  $l + n$  as follows.

$$l + n := \text{addr}(l.\text{block}, l.\text{cell} + n)$$

A *cell value*  $a, b, c: \text{CellVal}$  is a value stored in a single memory cell. An *expression*  $e: \text{Expr}$  is a structural program that may perform some imperative operations. Cell values and expressions are mutually inductively defined as follows.

---

<sup>1</sup> To be precise, this deviates from the behavior the real-world systems (e.g., a function definition is stored in the text segment of the program memory instead of carried around as a value). Still, we believe that this is a reasonable simplification.

<sup>2</sup> The letter ‘l’ stands for ‘location’.

(cell value)  $\mathbf{a}, \mathbf{b}, \mathbf{c} : \text{CellVal} ::= n$  (integer;  $n: \mathbb{Z}$ ) |  $\mathbf{I}$  (address) |  $\mathbf{f}$  (function) |  $\frac{1}{\downarrow}$  (invalid)

(function value)  $\mathbf{f} : \text{FnVal} ::= \text{fn } f(\vec{a})\{e\}$

(expression)  $e : \text{Expr} ::= a$  (variable) |  $\mathbf{a}$  (value)

|  $e(\vec{e}')$  (function call) |  $\text{alloc } e$  (allocate) |  $\text{free } e$  (free)

|  $*e$  (load) |  $e \leftarrow e'$  (store) |  $e.e'$  (address shift)

|  $\text{case } e \text{ of } \{0 \rightarrow e', 1 \rightarrow e''\}$  (conditional branching)

|  $e \parallel e'$  (concurrent execution)

|  $e \text{ iop } e'$  (integer operation) |  $e \text{ irel } e'$  (integer relation)

|  $\text{ndint}$  (non-deterministic integer)

(program variable)  $a, b, c, f : \text{Var}$

$\text{iop} ::= + \mid - \mid \times \mid \dots$        $\text{irel} ::= \leq \mid < \mid = \mid \dots$

A cell value  $\mathbf{a}, \mathbf{b}, \mathbf{c}$  can be an integer  $n$ , an address  $\mathbf{I}$ , a function value  $\mathbf{f}$ , or the invalid value  $\frac{1}{\downarrow}$ . We use variables  $\vec{a}, \vec{b}, \vec{c}$  for a list of cell values. The invalid value cannot be used in a meaningful way. A function value  $\mathbf{f} = \text{fn } f(\vec{a})\{e\}$  consists of (i) the variable  $f$  binding the function itself, which is used for recursion in  $e$ , (ii) the parameter variables  $\vec{a}$ , and (iii) the function body  $e$ .

We describe here the meaning of the expression primitives. By  $e(\vec{e}')$ , we call a function  $e$  with arguments  $\vec{e}'$ . By  $\text{alloc } e$ , we allocate a new memory block of the size  $e$ . By  $\text{free } e$ , we free the memory block starting at the address  $e$ . By  $*e$ , we load the value of the memory cell at the address  $e$ . By  $e \leftarrow e'$ , we update the value of the memory cell at the address  $e$  into the value  $e'$ . By  $e.e'$ , we get the address that is ahead of the address  $e$  by the integer  $e'$ . By  $\text{case } e \text{ of } \{0 \rightarrow e', 1 \rightarrow e''\}$ , we conditionally branch by the value of  $e$ , executing  $e'$  if the value is 0 and  $e''$  when if 1. By  $e \parallel e'$ , we execute the expressions  $e$  and  $e'$  concurrently; the expression finally returns the cell value of  $e$  after both  $e$  and  $e'$  terminates. We calculate a binary integer operation (that inputs and outputs an integer) by  $e \text{ iop } e'$  and a binary integer relation by  $e \text{ irel } e'$ . By  $\text{ndint}$ , we get a non-deterministic integer value.

We write  $\text{irel}_{\text{bool}}$  for the function of the type  $\mathbb{Z} \rightarrow \mathbb{Z} \rightarrow \mathbb{B}$  that corresponds to the binary integer relation  $\text{irel}$ .

We use the following shorthand for cell values.

$\text{tt} := 1$        $\text{ff} := 0$        $\text{fn}(\vec{a})\{e\} := \text{fn } \_(\vec{a})\{e\}$

We use the underscore  $\_$  for a special program variable from which we never take out the value (i.e.,  $\_$  is never used as a variable expression).

We use the following shorthand for expressions.

$\text{let } a = e \text{ in } e' := (\text{fn}(a)\{e'\})(e)$        $e; e' := \text{let } \_ = e \text{ in } e'$

$\text{if } e \{e'\} \text{ else } \{e''\} := \text{case } e \text{ of } \{0 \rightarrow e'', 1 \rightarrow e'\}$        $\text{if } e \{e'\} := \text{if } e \{e'\} \text{ else } \{\frac{1}{\downarrow}\}$

$\text{for } a \leftarrow e \dots e' \{e''\} :=$

$(\text{fn } \text{for}(b, c, f) \{ \text{if } b \geq c \{ \frac{1}{\downarrow} \} \text{ else } \{ f(b); \text{for}(b+1, c, f) \} \})(e, e', \text{fn}(a)\{e''\})$

$e \leftarrow_{e''}^* e' := (\text{fn}(a, b, c) \{ \text{for } i \leftarrow 0 \dots c \{ a.i \leftarrow *(b.i) \} \})(e, e', e'')$

For simplicity, let binding  $\text{let } a = e \text{ in } e'$  is defined using a function call. The for loop  $\text{for } a \leftarrow e \dots e' \{e''\}$  iteratively performs  $e''$  setting  $a$  to each integer value from the integer

$e$  to the integer  $e'$  exclusive. The sequential store  $e \leftarrow_{e''}^*$ ,  $e'$  copies  $e''$  consecutive cell values from the source address  $e'$  to the target address  $e$ .

**Operational Semantics** A *heap* (or heap memory)  $H$  is an object of the type

$$\text{Heap} := \mathbb{Z} \stackrel{\text{fin}}{\mapsto} \text{List CellVal},$$

a finite map from *memory block* ids to a list of cell values, representing the sequence of *memory cells* for the memory block.

For a heap  $H$ , we define the address domain  $\text{Dom } H: \mathbb{P}_{\text{fin}} \text{Addr}$  as follows.

$$\text{Dom } H := \{ I: \text{Addr} \mid I.\text{block} \in \text{dom } H \wedge 0 \leq I.\text{cell} < \text{len } H[I.\text{block}] \}$$

Here,  $\text{dom } H: \mathbb{P}_{\text{fin}} \mathbb{N}$  denotes the domain of  $H$  as a finite map over natural numbers. For an address  $I$  in  $\text{Dom } H$ , we introduce the following shorthand.

$$H[I] := H[I.\text{block}][I.\text{cell}] \quad H\{I \leftarrow \mathbf{a}\} := H\{I.\text{block} \leftarrow H[I.\text{block}]\{I.\text{cell} \leftarrow \mathbf{a}\}\}$$

Here,  $H[I]$  denotes the cell value at the address  $I$  and  $H\{I \leftarrow \mathbf{a}\}$  denotes the heap with the value of the memory cell at  $I$  updated into  $\mathbf{a}$ .

Also, for a heap  $H$  and a memory block id  $n$ ,  $H - n$  denotes the heap obtained from  $H$  by erasing the memory block at  $n$ .

The *evaluation context*  $K: \text{EvCtx}$  is defined as follows.

$$\begin{aligned} K: \text{EvCtx} ::= & \bullet \mid e(\vec{\mathbf{a}}, K, \vec{e}) \mid K(\vec{\mathbf{a}}) \mid \text{alloc } K \mid \text{free } K \\ & \mid *K \mid e \leftarrow K \mid K \leftarrow \mathbf{a} \mid K.e \mid \mathbf{a}.K \mid \text{case } K \text{ of } \{0 \rightarrow e, 1 \rightarrow e'\} \\ & \mid K \parallel e \mid e \parallel K \mid K \text{ iop } e \mid \mathbf{a} \text{ iop } K \mid K \text{ irel } e \mid \mathbf{a} \text{ irel } K \end{aligned}$$

For the concurrent execution  $e \parallel e'$ , we can choose whether to reduce  $e$  or  $e'$  for each reduction step. In the store expression  $e \leftarrow e'$ , the source expression  $e'$  is evaluated earlier than the target expression  $e$ .

Now the reduction relation  $(e, H) \rightarrow (e', H')$  is defined inductively by the following rules.<sup>3</sup>

$$\frac{(e, H) \rightarrow (e', H)}{(K[e], H) \rightarrow (K[e'], H)} \quad \frac{\vec{\mathbf{f}} = \text{fn } f(a_0, \dots, a_{n-1})\{e\}}{(\vec{\mathbf{f}}(a_0, \dots, a_{n-1}), H) \rightarrow (e[\vec{\mathbf{f}}/f, \vec{\mathbf{a}}/\vec{a}], H)}$$

$$\frac{m \notin \text{dom } H \quad H' = H\{m \leftarrow \vec{\mathbf{a}}\}}{(\text{alloc } n, H) \rightarrow (\text{addr}(m, 0), H')} \quad \frac{l.\text{cell} = 0 \quad H' = H - l.\text{block}}{(\text{free } l, H) \rightarrow (\zeta, H')}$$

$$\frac{I \in \text{Dom } H}{(*I, H) \rightarrow (H[I], H)} \quad \frac{I \in \text{Dom } H}{(I \leftarrow \mathbf{a}, H) \rightarrow (\zeta, H\{I \leftarrow \mathbf{a}\})} \quad \frac{I' = I + n}{(I.n, H) \rightarrow (I', H)}$$

$$\frac{i = 0 \vee i = 1}{(\text{case } i \text{ of } \{0 \rightarrow e_0, 1 \rightarrow e_1\}, H) \rightarrow (e_i, H)} \quad (\mathbf{a} \parallel \mathbf{b}, H) \rightarrow (\mathbf{a}, H)$$

$$\frac{l = m \text{ iop } n}{(m \text{ iop } n, H) \rightarrow (l, H)} \quad \frac{bl = m \text{ irel}_{\text{bool}} n}{(m \text{ irel } n, H) \rightarrow (bl, H)} \quad (\text{ndint}, H) \rightarrow (n, H)$$

Note that the reduction rules expect some *preconditions* on the form of the expression and the heap. When the preconditions are not satisfied, the reduction gets *stuck*. Also note that we cannot perform any reduction when the expression  $e$  is simply a cell value.

<sup>3</sup> In the rule on *irel*, the Boolean value  $bl$  is interpreted as 1 or 0 as a cell value.

For convenience, we introduce the following set  $\text{Next}(e, H)$ , which is the set of configurations to which the configuration  $(e, H)$  can reduce.

$$\text{Next}(e, H) := \{ (e', H') \mid (e, H) \rightarrow (e', H') \}$$

Also, we introduce the following pure predicate  $\text{red}(e, H)$ , which means that the configuration  $(e, H)$  is reducible.

$$\text{red}(e, H) := \text{Next}(e, H) \neq \emptyset$$

## 4.2 Verification in Iris

We model the weakest precondition and the Hoare triple for the imperative lambda calculus presented in the previous section.

We introduce a simple but new mechanism that enables us to *spend unboundedly many logical steps for each physical step* in Iris, in a sense. It is an important technique in Iris, because Iris employs *finite* step indexing; in Iris, the proposition  $\exists n. \text{E}^n_{\mathcal{E}} P$  is equivalent to  $\text{True}$  for any  $P$ , because at each step index  $i$  we can set  $n$  to  $i$ , which makes all the ‘visible part’ at the step index  $\text{True}$ . More specifically, our machinery enables spending  $\text{step}(n) + 1$  logical steps for the  $n$ -th physical step, for some fixed monotone function  $\text{step}$ , e.g.,  $\lambda n. n$ . This comes from the observation that (i) typically the number of logical steps we need to spend at once is proportional to the *depth* or complexity of the manipulated object and that (ii) the depth of the object lower-bounds the number of program steps we have spent so far. In order to manage how many physical steps have passed, we introduce a new *time receipt*  $RA$ , extending the idea of Mével et al. (2019). This machinery is used later in RustHornBelt to temporarily take out prophecy tokens deep inside an object for dependent resolution (see `SEM-TY-OWN-PROPH-TOKEN` and `SWKN-UNQREF-RELEASE`).

**Heap RA** In order to manage information about the heap, we introduce the following *heap RA*  $\text{HEAP}$ .

$$|\text{HEAP}| := |\text{EX}(\text{Heap})| \times |\mathbb{Z} \rightarrow \text{EX}(\mathbb{N})| \times |\text{Addr} \rightarrow \text{FRACOWN}(\text{CellVal})|$$

$$(a, f, h) \cdot (b, g, h') := (a \cdot b, f \cdot g, h \cdot h') \quad \varepsilon := (\varepsilon, \varepsilon, \varepsilon) \quad |(a, f, h)| := (\varepsilon, \varepsilon, \varepsilon)$$

$$\begin{aligned} \checkmark(a, f, h) &:= \checkmark a \wedge \checkmark f \wedge \checkmark h \wedge \forall H \text{ s.t. } \text{ex } H = a. \\ &(\forall m, n \text{ s.t. } \text{ex } n = f m. m \in \text{dom } H \wedge \text{len } H[m] = n) \wedge \\ &(\forall I, \mathbf{a}, q \text{ s.t. } \text{fown}_q \mathbf{a} = h I. I \in \text{Dom } H \wedge H[I] = \mathbf{a}) \end{aligned}$$

We register the heap RA  $\text{HEAP}$  to the global camera. Using some ghost name  $\gamma_{\text{heap}}$ , we introduce the following *exclusive heap token*  $\text{Ex}(H)$ , *free-right token*  $\text{Free}_n(I)$  and *maps-to token*  $I \xrightarrow{q} \mathbf{a}$ .

$$\text{Ex}(H) := \left[ \left( \text{ex } H, \varepsilon, \varepsilon \right) \right]_{\text{HEAP}}^{\gamma_{\text{heap}}} \quad \text{Free}_n(I) := I.\text{cell} = 0 * \left[ \left( \varepsilon, [I.\text{block} \leftarrow n], \varepsilon \right) \right]_{\text{HEAP}}^{\gamma_{\text{heap}}}$$

$$I \xrightarrow{q} \mathbf{a} := \left[ \left( \varepsilon, \varepsilon, [I \leftarrow \text{fown}_q \mathbf{a}] \right) \right]_{\text{HEAP}}^{\gamma_{\text{heap}}}$$

An exclusive heap token  $\text{Ex}(H)$  registers the heap  $H$ . A free-right token  $\text{Free}_n(I)$  owns the right to free the memory block starting at  $I$  of the size  $n$ . A points-to token  $I \xrightarrow{q} \mathbf{a}$  witnesses with the fraction  $q$  that the value of the memory cell at  $I$  is  $\mathbf{a}$ . We also introduce the following sequential points-to token  $I \xrightarrow{q} \bar{\mathbf{a}}$ , which witnesses with the fraction  $q$  that the cell values of the  $\text{len } \bar{\mathbf{a}}$  consecutive memory cells starting at  $I$  are  $\bar{\mathbf{a}}$ .

$$I \xrightarrow{q} \bar{\mathbf{a}} := *_{i < \text{len } \bar{\mathbf{a}}} I + i \xrightarrow{q} \bar{\mathbf{a}}[i]$$

We also use the following shorthand.

$$I \xrightarrow{q} \_ := \exists \mathbf{a}. I \xrightarrow{q} \mathbf{a} \quad I \xrightarrow{q} \_{}^n := \exists \bar{\mathbf{a}} \text{ s.t. } \text{len } \bar{\mathbf{a}} = n. I \xrightarrow{q} \bar{\mathbf{a}}$$

The ghost name  $\gamma_{\text{heap}}$  is fixed when we introduce the exclusive heap token  $\text{Ex}(H)$  using the following rule.

$$\begin{array}{c} \text{EXHEAP-INTRO} \\ \text{True} \Rightarrow_{\emptyset} \exists \gamma_{\text{heap}}. \text{Ex}(H) \end{array}$$

The following properties hold.

$$\text{timeless}(\text{Ex}(H)) \quad \text{timeless}(\text{Free}_n(I)) \quad \text{timeless}(I \xrightarrow{q} \mathbf{a})$$

$$\begin{array}{c} \text{MAPSTO-FRAC} \\ I \xrightarrow{q+q'} \mathbf{a} \Leftrightarrow I \xrightarrow{q} \mathbf{a} * I \xrightarrow{q'} \mathbf{a} \end{array} \quad \begin{array}{c} \text{MAPSTO-AGREE} \\ I \xrightarrow{q} \mathbf{a} * I \xrightarrow{q'} \mathbf{a}' \Rightarrow \mathbf{a} = \mathbf{a}' \end{array}$$

$$\begin{array}{c} \text{HEAP-ALLOC} \\ \frac{m \notin \text{dom } H \quad \text{len } \bar{\mathbf{a}} = n \quad I = \text{addr}(m, 0)}{\text{Ex}(H) \Rightarrow_{\emptyset} \text{Free}_n(I) * I \xrightarrow{1} \bar{\mathbf{a}} * \text{Ex}(H\{m \leftarrow \bar{\mathbf{a}}\})} \end{array}$$

$$\begin{array}{c} \text{HEAP-FREE} \\ \text{Free}_n(I) * I \xrightarrow{1} \_{}^n * \text{Ex}(H) \Rightarrow_{\emptyset} \text{Ex}(H - I.\text{block}) \end{array}$$

$$\begin{array}{c} \text{HEAP-LOAD} \\ I \xrightarrow{q} \mathbf{a} * \text{Ex}(H) \Rightarrow I \in \text{Dom } H \wedge H[I] = \mathbf{a} \end{array}$$

$$\begin{array}{c} \text{HEAP-STORE} \\ I \xrightarrow{1} \mathbf{a} * \text{Ex}(H) \Rightarrow_{\emptyset} I \xrightarrow{1} \mathbf{b} * \text{Ex}(H\{I \leftarrow \mathbf{b}\}) \end{array}$$

Points-to tokens can be split and merged according to the fraction ([MAPSTO-FRAC](#)). Two points-to tokens on the same address should agree on the cell value ([MAPSTO-AGREE](#)). [HEAP-ALLOC](#) says that, when we allocate in the heap a new memory block of the length  $n$  and the cell values  $\bar{\mathbf{a}}$ , letting  $I$  be the head address of the memory block, we get a free-right token  $\text{Free}_n(I)$  and a full sequential points-to token  $I \xrightarrow{1} \bar{\mathbf{a}}$ . [HEAP-FREE](#) says that, by consuming a free-right token on an address  $I$  and the size  $n$  and a full points-to token on the  $n$  consecutive memory cells starting at  $I$ , we can free the memory block of  $I$ . [HEAP-LOAD](#) says that, when we have a fractional points-to token  $I \xrightarrow{q} \mathbf{a}$ , we know that there exists a memory cell at  $I$  storing  $\mathbf{a}$ . [HEAP-STORE](#) says that, by consuming a full points-to token  $I \xrightarrow{1} \mathbf{a}$ , we can update the value of the memory cell at  $I$  into any value  $\mathbf{b}$  and obtain a new full points-to token  $I \xrightarrow{1} \mathbf{b}$ .

**Time Receipt RA** In order to manage information about *how many physical steps have passed*, we introduce the following *time receipt RA*  $\text{TIME}$ .

$$|\text{TIME}| := |\text{EX}(\mathbb{N})| \times \mathbb{N} \times \mathbb{N}$$

$$(a, m, n) \cdot (b, m', n') := (a \cdot b, m + m', \max\{n, n'\}) \quad \varepsilon := (\varepsilon, 0, 0)$$

$$|(a, m, n)| := (\varepsilon, 0, n) \quad \checkmark(a, m, n) := \checkmark a \wedge \forall l \text{ s.t. } \text{ex } l = a. l \geq m + n$$

We register the time receipt RA  $\text{TIME}$  to the global camera. Using some ghost name  $\gamma_{\text{time}}$ , we introduce the following *exclusive time receipt*  $\mathbf{\times} n$ , *cumulative time receipt*  $\mathbf{\Sigma} n$ , and *persistent time receipt*  $\mathbf{\bar{\Sigma}} n$ .<sup>4</sup>

$$\mathbf{\times} n := \left[ \left[ \text{ex } n, 0, 0 \right] \right]_{\text{TIME}}^{\gamma_{\text{time}}} \quad \mathbf{\Sigma} n := \left[ \left[ \varepsilon, n, 0 \right] \right]_{\text{TIME}}^{\gamma_{\text{time}}} \quad \mathbf{\bar{\Sigma}} n := \left[ \left[ \varepsilon, 0, n \right] \right]_{\text{TIME}}^{\gamma_{\text{time}}}$$

<sup>4</sup> The symbols represent sandglasses.

An exclusive time receipt  $\mathbf{\times} n$  registers the number  $n$  of physical steps that have passed. When we have a cumulative time receipt  $\mathbf{\bar{\times}} m$  and a persistent time receipt  $\mathbf{\bar{\times}} n$ , we witness that at least  $m + n$  physical steps have passed. Cumulative time receipts follow the addition law  $\mathbf{\bar{\times}} m * \mathbf{\bar{\times}} n \Leftrightarrow \mathbf{\bar{\times}}(m + n)$  but are not persistent. Persistent time receipts are persistent but do not follow such an addition law. We fix the ghost name  $\gamma_{\text{time}}$  when we introduce the exclusive time receipt by the following rule (typically  $n$  is set to 0).

$$\begin{array}{c} \text{EXTIME-INTRO} \\ \text{True} \Rightarrow_{\emptyset} \exists \gamma_{\text{time}}. \mathbf{\times} n \end{array}$$

The following properties hold.

$\text{timeless}(\mathbf{\times} n)$	$\text{timeless}(\mathbf{\bar{\times}} n)$	$\text{persistent}(\mathbf{\bar{\times}} n)$	$\text{timeless}(\mathbf{\bar{\times}} n)$
$\text{EXTIME-INCREMENT-CUMUETIME}$	$\text{CUMUETIME-ADD}$	$\text{PERS TIME-0}$	
$\mathbf{\times} n \Rightarrow_{\emptyset} \mathbf{\times}(n + 1) * \mathbf{\bar{\times}} 1$	$\mathbf{\bar{\times}} m * \mathbf{\bar{\times}} n \Leftrightarrow \mathbf{\bar{\times}}(m + n)$	$\mathbf{\bar{\times}} 0$	
$\text{CUMUETIME-SWELL-PERS TIME}$	$\text{PERS TIME-BOUND-EX TIME}$		
$\frac{\mathbf{\bar{\times}} n}{\mathbf{\bar{\times}} m \Rightarrow_{\emptyset} \mathbf{\bar{\times}}(m + n)}$	$\frac{\mathbf{\bar{\times}} m}{\mathbf{\times} n \Rightarrow n \geq m}$		

**EXTIME-INCREMENT-CUMUETIME** says that, when we increment the number of the exclusive time receipt, we get a cumulative time receipt of one step  $\mathbf{\bar{\times}} 1$ . Cumulative time receipts follow the addition law **CUMUETIME-ADD**. A persistent receipt of zero step  $\mathbf{\bar{\times}} 0$  can be freely obtained (**PERS TIME-0**). The rule **CUMUETIME-SWELL-PERS TIME** is interesting and important; by consuming a cumulative time receipt  $\mathbf{\bar{\times}} m$ , we can *swell* a persistent time receipt  $\mathbf{\bar{\times}} n$  into  $\mathbf{\bar{\times}}(m + n)$ .<sup>5</sup> **PERS TIME-BOUND-EX TIME** says that the number  $m$  of a persistent time receipt  $\mathbf{\bar{\times}} m$  lower-bounds the number  $n$  of the exclusive time receipt  $\mathbf{\times} n$ .

**Weakest Precondition and Hoare Triple** We fix some *monotone* function over natural numbers  $\text{step}: \mathbb{N} \rightarrow \mathbb{N}$ . For RustHornBelt, we particularly set  $\text{step}(n) = 2n$ .

Now we define the *weakest precondition*  $\text{wp}_{\mathcal{E}} e \{ \Phi \}$ , where  $\Phi$  is the postcondition over the returned cell value (typed  $\text{CellVal} \rightarrow \text{IProp}$ ).

$$\begin{aligned} \text{wp}_{\mathcal{E}} e \{ \Phi \} &:= (\exists \mathbf{a} \text{ s.t. } \mathbf{a} = e. \Rightarrow_{\mathcal{E}} \Phi \mathbf{a}) \vee ((\forall \mathbf{a}. \mathbf{a} \neq e) \wedge \\ &\forall H, n. \text{Ex}(H) * \mathbf{\times} n \not\Rightarrow_{\mathcal{E}}^{\text{step}(n)+1} \text{red}(e, H) * \\ &\forall (e', H') \in \text{Next}(e, H). \Rightarrow_{\mathcal{E}} (\text{Ex}(H') * \mathbf{\times}(n+1) * \text{wp}_{\mathcal{E}} e' \{ \Phi \})) \end{aligned}$$

This recursive equation over  $\text{wp}_{\mathcal{E}} \cdots \{ \Phi \}$  has the unique solution due to the *contractiveness* introduced by the step-taking view shift  $\not\Rightarrow_{\mathcal{E}}^{\text{step}(n)+1}$  (by Banach's fixed point theorem [Theorem 2.1](#)).

This definition of the weakest precondition  $\text{wp}_{\mathcal{E}} e \{ \Phi \}$  can be understood as follows. When the given expression  $e$  is the form of some value  $\mathbf{a}$ , the postcondition  $\Phi \mathbf{a}$  should be satisfied after a fancy update ( $\Rightarrow_{\mathcal{E}}$ ). When  $e$  is not a value, and when the heap is  $H$  and it is the  $n$ -th physical step, having the tokens  $\text{Ex}(H)$  and  $\mathbf{\times} n$ , after fancy update *with*  $\text{step}(n) + 1$  *logical steps* ( $\Rightarrow_{\mathcal{E}}^{\text{step}(n)+1}$ ), we know that the configuration  $(e, H)$  is not stuck, and for any next possible configuration  $(e', H')$ , along with the tokens  $\text{Ex}(H')$  and  $\mathbf{\times}(n+1)$  we get the weakest precondition  $\text{wp}_{\mathcal{E}} e' \{ \Phi \}$ , which gives us the postcondition  $\Phi$  after executing  $e'$ .

<sup>5</sup> This use of a cumulative time receipt was conceived by Jacques-Henri Jourdan.

We introduce the following *super fancy update modality*  $\mathbb{H}_{\mathcal{E}}^{\#} P$ .

$$\mathbb{H}_{\mathcal{E}}^{\#} P := \mathbb{H}_{\mathcal{E}}(\exists n. \bar{\times} n * \mathbb{H}_{\mathcal{E}}^{\text{step}(n)} P)$$

It means that, after one logical step, for some  $n$  such that we know that  $n$  physical steps have passed by  $\bar{\times} n$ , we can get  $P$  after  $\text{step}(n)$  logical steps, under the mask  $\mathcal{E}$ . Note that  $\exists n. \bar{\times} n * \mathbb{H}_{\mathcal{E}}^{\text{step}(n)} P$  is *not* equal to True, unlike  $\exists n. \mathbb{H}_{\mathcal{E}}^{\text{step}(n)} P$ , because the persistent time receipt  $\bar{\times} n$  prevents us from taking unboundedly large  $n$  as the step index grows.

We can construct a proposition under a super fancy update modality using the following properties.

$$\begin{array}{c} \text{SPFUPD-DEF} \\ \mathbb{H}_{\mathcal{E}}(\exists n. \bar{\times} n * \mathbb{H}_{\mathcal{E}}^{\text{step}(n)} P) \Leftrightarrow \mathbb{H}_{\mathcal{E}}^{\#} P \end{array} \qquad \begin{array}{c} \text{SPFUPD-ZERO} \\ \mathbb{H}_{\mathcal{E}}^{\text{step}(0)+1} P \Rightarrow \mathbb{H}_{\mathcal{E}}^{\#} P \end{array}$$

$$\begin{array}{c} \text{SPFUPD-MERGE} \\ \mathbb{H}_{\mathcal{E}}^{\#} P * \mathbb{H}_{\mathcal{E}}^{\#} Q \Rightarrow \mathbb{H}_{\mathcal{E}}^{\#} (P * Q) \end{array}$$

We also introduce the following shorthand.

$$P \mathbb{H}_{\mathcal{E}}^{\#} Q := P * \mathbb{H}_{\mathcal{E}}^{\#} Q \qquad P \mathbb{H}_{\mathcal{E}}^{\#} Q := \square(P \mathbb{H}_{\mathcal{E}}^{\#} Q)$$

We can strip off the super fancy update modality using the following property.

$$\begin{array}{c} \text{EXTIME-SPFUPD-CUMU TIME} \\ \mathbb{H}_{\mathcal{E}}^{\#} P * \bar{\times} n \mathbb{H}_{\mathcal{E}}^{\text{step}(n)+1} P * \bar{\times} 1 * \bar{\times}(n+1) \end{array}$$

It means that, spending  $\text{step}(n) + 1$  logical steps under the mask  $\mathcal{E}$ , if we update  $\bar{\times} n$  into  $\bar{\times}(n+1)$ , we can strip off one super fancy update modality  $\mathbb{H}_{\mathcal{E}}^{\#}$  and get a cumulative time receipt  $\bar{\times} 1$ .

*Proof.* Assume that we have  $\mathbb{H}_{\mathcal{E}}^{\#} P$  and  $\bar{\times} n$ . We can decompose  $\mathbb{H}_{\mathcal{E}}^{\#} P$  into  $\bar{\times} m$  and  $\mathbb{H}_{\mathcal{E}}^{\text{step}(m)+1} P$  for some  $m$ . By [PERTIME-BOUND-EX TIME](#) we know  $n \geq m$ . Since  $\text{step}(n) + 1 \geq \text{step}(m) + 1$  holds by monotonicity of  $\text{step}$ , we can update  $\mathbb{H}_{\mathcal{E}}^{\text{step}(m)+1} P$  into  $P$  in  $\text{step}(n) + 1$  logical steps. Also, by [EX TIME-INCREMENT-CUMU TIME](#), we can update  $\bar{\times} n$  into  $\bar{\times} 1$  and  $\bar{\times}(n+1)$ .  $\square$

By [EX TIME-SPFUPD-CUMU TIME](#), we can prove the following basic lemma for discussing the weakest precondition over an operation that takes one reduction step.

$$\frac{\text{WP-STEP} \quad \text{Ex}(H) * P \Rightarrow_{\mathcal{E}} \text{red}(e, H) * \forall (e', H') \in \text{Next}(e, H). \mathbb{H}_{\mathcal{E}}(\text{Ex}(H') * \text{wp}_{\mathcal{E}} e' \{\Phi\})}{\mathbb{H}_{\mathcal{E}}^{\#} P \Rightarrow \text{wp}_{\mathcal{E}} e \{\lambda a. \bar{\times} 1 * \Phi a\}}$$

In one physical step, we can strip off the super fancy later modality and gain a cumulative time receipt  $\bar{\times} 1$ .

We also have the following lemmas for the weakest precondition on an evaluation context and a value expression.

$$\begin{array}{c} \text{WP-EVCTX} \\ \text{wp}_{\mathcal{E}} e \{\lambda a. \text{wp}_{\mathcal{E}} K[a] \{\Phi\}\} \Rightarrow \text{wp}_{\mathcal{E}} K[e] \{\Phi\} \end{array} \qquad \begin{array}{c} \text{WP-VAL} \\ \Phi a \Leftrightarrow \text{wp}_{\mathcal{E}} a \{\Phi\} \end{array}$$

We can modify the postcondition part of the weakest precondition using a view shift.

$$\begin{array}{c} \text{WP-VSHIFT} \\ (\forall a. \Phi a \mathbb{H}_{\mathcal{E}}^{\#} \Psi a) * \text{wp}_{\mathcal{E}} e \{\Phi\} \Rightarrow \text{wp}_{\mathcal{E}} e \{\Psi\} \end{array}$$

We can prove this rule using Löb induction ([LÖB](#)).

The weakest precondition satisfies the following *adequacy theorem*.

**Theorem 4.1** (Adequacy of the Weakest Precondition). *Let  $\phi$  be a pure predicate on a cell value. If*

$$\models_{\mathcal{E}} \exists \gamma_{\text{heap}}, \gamma_{\text{time}}. \text{Ex}(H_0) * \mathbf{\Sigma}k * \text{wp}_{\mathcal{E}} e_0 \{ \phi \}$$

*is a tautology, for any reduction sequence  $(e_0, H_0) \rightarrow (e_1, H_1) \rightarrow \dots \rightarrow (e_n, H_n)$  such that  $\neg \text{red}(e_n, H_n)$  holds,  $e_n$  is a cell value satisfying  $\phi$ .*

Here, the tokens  $\text{Ex}(H_0)$  and  $\mathbf{\Sigma}k$  depend on the ghost names  $\gamma_{\text{heap}}$  and  $\gamma_{\text{time}}$ . We can use the fancy update  $\models_{\mathcal{E}}$  and the existential quantification  $\exists \gamma_{\text{heap}}, \gamma_{\text{time}}$  for introducing the tokens  $\text{Ex}(H_0)$  and  $\mathbf{\Sigma}k$  by the rules [EXHEAP-INTRO](#) and [EXTIME-INTRO](#).

*Proof.* By the definition of the weakest precondition, the following holds for each reduction step  $(e_i, H_i) \rightarrow (e_{i+1}, H_{i+1})$ .

$$\text{Ex}(H_i) * \mathbf{\Sigma}(k+i) * \text{wp}_{\mathcal{E}} e_i \{ \phi \} \xRightarrow{\mathcal{E}}^{\text{step}(k+i)+1} \text{Ex}(H_{i+1}) * \mathbf{\Sigma}(k+i+1) * \text{wp}_{\mathcal{E}} e_{i+1} \{ \phi \}$$

Therefore, using the assumption tautology  $\text{Ex}(H_0) * \mathbf{\Sigma}k * \text{wp}_{\mathcal{E}} e_0 \{ \phi \}$ , we get the following tautology.

$$\models_{\mathcal{E}}^{\sum_{i=0}^{n-1} (\text{step}(k+i)+1)} \exists \gamma_{\text{heap}}, \gamma_{\text{time}}. (\text{Ex}(H_n) * \mathbf{\Sigma}(k+n) * \text{wp}_{\mathcal{E}} e_n \{ \phi \})$$

Also, since  $\neg \text{red}(e_n, H_n)$  holds, when we have the tokens  $\text{Ex}(H_n)$  and  $\mathbf{\Sigma}(k+n)$ , it turns out after  $\text{step}(k+n) + 1$  logical steps that the weakest precondition cannot take the right-hand disjunct. So we have the following tautology.

$$(\exists \gamma_{\text{heap}}, \gamma_{\text{time}}. \text{Ex}(H_n) * \mathbf{\Sigma}(k+n) * \text{wp}_{\mathcal{E}} e_n \{ \phi \}) \xRightarrow{\mathcal{E}}^{\text{step}(k+n)+1} \exists \mathbf{a} \text{ s.t. } \mathbf{a} = e_n. \phi \mathbf{a}$$

Combining the two tautologies, we have the following tautology.

$$\models_{\mathcal{E}}^{\sum_{i=0}^n (\text{step}(k+i)+1)} \exists \mathbf{a} \text{ s.t. } \mathbf{a} = e_n. \phi \mathbf{a}$$

By the soundness theorem on the step-taking fancy update modality [Theorem 2.2](#), we finally get  $\exists \mathbf{a} \text{ s.t. } \mathbf{a} = e_n. \phi \mathbf{a}$ .  $\square$

For verification, it is more convenient to use the *Hoare triple*  $\{ P \} e \{ \mathbf{a}. Q_{\mathbf{a}} \}_{\mathcal{E}}$  than directly use the weakest precondition. It is simply defined as follows.

$$\{ P \} e \{ \mathbf{a}. Q_{\mathbf{a}} \}_{\mathcal{E}} := P \xRightarrow{\mathcal{E}} \text{wp}_{\mathcal{E}} e \{ \lambda \mathbf{a}. Q \}$$

Note that the Hoare triple is persistent. Also, we introduce the following shorthand for the case where the returned value is ignored by the postcondition.

$$\{ P \} e \{ Q \}_{\mathcal{E}} := \{ P \} e \{ \_ . Q \}_{\mathcal{E}}$$

We have the following structural rules on the Hoare triple.

$\frac{\text{HOARE-FRAME} \quad \{ P \} e \{ \mathbf{a}. Q_{\mathbf{a}} \}_{\mathcal{E}}}{\{ P * R \} e \{ \mathbf{a}. Q_{\mathbf{a}} * R \}_{\mathcal{E}}}$	$\frac{\text{HOARE-VSHIFT} \quad P' \xRightarrow{\mathcal{E}} P \quad \{ P \} e \{ \mathbf{a}. Q_{\mathbf{a}} \}_{\mathcal{E}} \quad \forall \mathbf{a}. Q_{\mathbf{a}} \xRightarrow{\mathcal{E}} Q'_{\mathbf{a}}}{\{ P' \} e \{ \mathbf{a}. Q'_{\mathbf{a}} \}_{\mathcal{E}}}$
$\frac{\text{HOARE-CTX-IN} \quad \text{persistent}(P) \quad P \quad \{ P * Q \} e \{ \mathbf{a}. R_{\mathbf{a}} \}_{\mathcal{E}}}{\{ Q \} e \{ \mathbf{a}. R_{\mathbf{a}} \}_{\mathcal{E}}}$	$\frac{\text{HOARE-CTX-OUT} \quad P \xRightarrow{\mathcal{E}} \{ Q \} e \{ \mathbf{a}. R_{\mathbf{a}} \}_{\mathcal{E}}}{\{ P * Q \} e \{ \mathbf{a}. R_{\mathbf{a}} \}_{\mathcal{E}}}$

$$\begin{array}{c}
\text{HOARE-FALSE} \\
\frac{}{\{ \text{False} \} e \{ \mathbf{a}. P \}_{\mathcal{E}}} \\
\\
\text{HOARE-EXIST} \\
\frac{\forall x. \{ P_x \} e \{ \mathbf{a}. Q_x \}_{\mathcal{E}}}{\{ \exists x. P_x \} e \{ \mathbf{a}. Q_x \}_{\mathcal{E}}} \\
\\
\text{HOARE-DISJ} \\
\frac{\{ P \} e \{ \mathbf{a}. R_a \}_{\mathcal{E}} \quad \{ Q \} e \{ \mathbf{a}. R_a \}_{\mathcal{E}}}{\{ P \vee Q \} e \{ \mathbf{a}. R_a \}_{\mathcal{E}}} \\
\\
\text{HOARE-MONO-MASK} \\
\frac{\{ P \} e \{ \mathbf{a}. Q_a \}_{\mathcal{E}} \quad \mathcal{E} \subseteq \mathcal{E}'}{\{ P \} e \{ \mathbf{a}. Q_a \}_{\mathcal{E}'}}
\end{array}$$

**HOARE-FRAME** is an important rule that removes a proposition  $R$  that is shared under *separating conjunction* by the precondition and postcondition, which is called a *frame*. **HOARE-VSHIFT** modifies the precondition and postcondition of the Hoare triple using view shifts. Both **HOARE-FRAME** and **HOARE-VSHIFT** follow from **WP-VSHIFT**. **HOARE-CTX-IN** adds persistent knowledge  $P$  to the precondition. **HOARE-CTX-OUT** takes out a part  $P$  of the precondition into an assumption of the Hoare triple. **HOARE-FALSE**, **HOARE-DISJ** and **HOARE-EXIST** eliminate false, disjunction and existential quantification in the precondition. **HOARE-MONO-MASK** weakens the mask.

We also have the following rules for an evaluation context and a value expression.

$$\begin{array}{c}
\text{HOARE-EVCTX} \\
\frac{\{ P \} e \{ \mathbf{a}. Q_a \}_{\mathcal{E}} \quad \forall \mathbf{a}. \{ Q_a \} K[\mathbf{a}] \{ \mathbf{b}. R_b \}_{\mathcal{E}}}{\{ P \} K[e] \{ \mathbf{b}. R_b \}_{\mathcal{E}}} \\
\\
\text{HOARE-VAL} \\
\frac{P \Rightarrow_{\mathcal{E}} Q_a}{\{ P \} \mathbf{a} \{ \mathbf{a}'. Q_{a'} \}_{\mathcal{E}}}
\end{array}$$

**HOARE-EVCTX** follows from **WP-EVCTX** and **HOARE-VAL** follows from **WP-VAL**.

We have the following rule for concurrent execution.

$$\text{HOARE-CONCUR} \\
\frac{\{ P \} e \{ \mathbf{a}. Q_a \}_{\mathcal{E}} \quad \{ P' \} e' \{ Q' \}_{\mathcal{E}}}{\{ P * P' \} e \parallel e' \{ \mathbf{a}. Q_a * Q' \}_{\mathcal{E}}}$$

This rule separately conjoins the preconditions and postconditions of the assumption Hoare triples on  $e$  and  $e'$ . It is a standard rule of *concurrent separation logic* (O'Hearn, 2007).

Using the lemma **WP-STEP**, we can prove the following Hoare-triple rules on basic operations.

$$\text{HOARE-FNCALL} \\
\frac{\mathfrak{f} = \text{fn } f(a_0, \dots, a_{n-1})\{e\} \quad \{ P * \nabla 1 \} e[\mathfrak{f}/f, \overline{\mathbf{a}}/\overline{a}] \{ \mathbf{c}. Q_c \}_{\mathcal{E}}}{\{ \boxRightarrow_{\mathcal{E}}^{\#} P \} \mathfrak{f}(a_0, \dots, a_{n-1}) \{ \mathbf{c}. Q_c \}_{\mathcal{E}}}$$

$$\text{HOARE-ALLOC} \\
\{ \boxRightarrow_{\mathcal{E}}^{\#} P \} \text{alloc } n \{ \mathbf{a}. \exists \mathbf{l} \text{ s.t. } \mathbf{l} = \mathbf{a}. \text{Free}_n(\mathbf{l}) * \mathbf{l} \mapsto \_{}^n * \nabla 1 * P \}_{\mathcal{E}}$$

$$\text{HOARE-FREE} \\
\{ \boxRightarrow_{\mathcal{E}}^{\#} (\text{Free}_n(\mathbf{l}) * \mathbf{l} \mapsto \_{}^n * P) \} \text{free } \mathbf{l} \{ \nabla 1 * P \}_{\mathcal{E}}$$

$$\text{HOARE-LOAD} \\
\{ \boxRightarrow_{\mathcal{E}}^{\#} (\mathbf{l} \xrightarrow{g} \mathbf{a} * P) \} * \mathbf{l} \{ \mathbf{a}'. \mathbf{a}' = \mathbf{a} * \mathbf{l} \xrightarrow{g} \mathbf{a} * \nabla 1 * P \}_{\mathcal{E}}$$

$$\text{HOARE-STORE} \\
\{ \boxRightarrow_{\mathcal{E}}^{\#} (\mathbf{l} \mapsto \_{} * P) \} \mathbf{l} \leftarrow \mathbf{a} \{ \mathbf{l} \mapsto \mathbf{a} * \nabla 1 * P \}_{\mathcal{E}}$$

$$\text{HOARE-ADDRSHIFT} \\
\{ \boxRightarrow_{\mathcal{E}}^{\#} P \} \mathbf{l}.n \{ \mathbf{a}. \mathbf{a} = \mathbf{l} + n * \nabla 1 * P \}_{\mathcal{E}}$$

$$\frac{\text{HOARE-CASE} \quad i = 0 \vee i = 1 \quad \{P * \nabla 1\} e_i \{a. Q_a\}_{\mathcal{E}}}{\{\mathbb{E}_{\mathcal{E}}^{\#} P\} \text{ case } i \text{ of } \{0 \rightarrow e_0, 1 \rightarrow e_1\} \{a. Q_a\}_{\mathcal{E}}}$$

$$\frac{\text{HOARE-INTOP}}{\{\mathbb{E}_{\mathcal{E}}^{\#} P\} m \text{ iop } n \{a. a = m \text{ iop } n * \nabla 1 * P\}_{\mathcal{E}}}$$

$$\frac{\text{HOARE-INTREL}}{\{\mathbb{E}_{\mathcal{E}}^{\#} P\} m \text{ irel } n \{a. a = m \text{ irel } n * \nabla 1 * P\}_{\mathcal{E}}}$$

$$\frac{\text{HOARE-NDINT}}{\{\mathbb{E}_{\mathcal{E}}^{\#} P\} \text{ ndint } \{a. \exists n. a = n * \nabla 1 * P\}_{\mathcal{E}}}$$

For each physical step, we can strip off the super fancy update modality  $\mathbb{E}_{\mathcal{E}}^{\#}$  and obtain a cumulative time receipt  $\nabla 1$ . The rules [HOARE-ALLOC](#), [HOARE-FREE](#), [HOARE-LOAD](#) and [HOARE-STORE](#) follow particularly from the rules [HEAP-ALLOC](#), [HEAP-FREE](#), [HEAP-LOAD](#) and [HEAP-STORE](#), respectively.

Using the rules above, we can derive the following rules.

$$\frac{\text{HOARE-LET} \quad \{P\} e \{a. \mathbb{E}_{\mathcal{E}}^{\#} Q_a\}_{\mathcal{E}} \quad \forall a. \{Q_a * \nabla 1\} e' [a/a] \{b. R_b\}_{\mathcal{E}}}{\{P\} \text{ let } a = e \text{ in } e' \{b. R_b\}_{\mathcal{E}}}$$

$$\frac{\text{HOARE-SEQ} \quad \{P\} e \{\mathbb{E}_{\mathcal{E}}^{\#} Q\}_{\mathcal{E}} \quad \{Q * \nabla 1\} e' \{a. R_a\}_{\mathcal{E}}}{\{P\} e; e' \{a. R_a\}_{\mathcal{E}}}$$

$$\frac{\text{HOARE-IF} \quad \{P * \nabla 1\} e_{bl} \{a. Q_a\}_{\mathcal{E}}}{\{\mathbb{E}_{\mathcal{E}}^{\#} P\} \text{ if } bl \{e_{tt}\} \text{ else } \{e_{ff}\} \{a. Q_a\}_{\mathcal{E}}}$$

$$\frac{\text{HOARE-SEQCOPY} \quad \text{len } \bar{a} = \text{len } \bar{b} = n}{\{\mathbb{E}_{\mathcal{E}}^{\# n+1} (I \mapsto \bar{a} * I' \mapsto \bar{b} * P)\} I \leftarrow_n^* I' \{I \mapsto \bar{b} * I' \mapsto \bar{b} * \nabla(n+1) * P\}_{\mathcal{E}}}$$

### 4.3 Related Work

**Time Receipts** The idea of *time receipts* came from [Mével et al. \(2019\)](#). A time receipt *lower*-bounds the number of physical steps that have elapsed. It was conceived as a dual of a *time credit* ([Atkey, 2011](#)), which *upper*-bounds the number of physical steps that have elapsed and is typically used for verifying an upper bound on the execution time of a program. [Mével et al. \(2019\)](#) used time receipts to ensure that undesirable events like *integer overflows* cannot occur within  $N$  physical steps, where  $N$  is a *global constant* that is typically set to a very large number like  $2^{63}$ . Unlike their use of time receipts, we use time receipts to ensure that we do not need to consume an undesirably large number of logical steps *with respect to* the number of physical steps that have elapsed. Also, our time receipt RA combines the *cumulative* and *persistent* time receipts through the *swelling* rule [CUMU<sub>TIME</sub>-SWELL-PERS<sub>TIME</sub>](#), which is a new idea.

**Transfinite Iris** A recent emerging study by [Spies et al. \(2020\)](#) introduces a variant of Iris called *Transfinite Iris*, which incorporates *transfinite* step indexing (i.e., machinery where a step index can be a transfinite ordinal) into Iris, which currently uses *finite* step

indexing (i.e., machinery where a step index is a natural number or finite ordinal). In Transfinite Iris, we can easily spend any finite number of logical steps for each physical step. This is due to the *existential property* of Transfinite Iris, which says that if  $\exists x: T. P_x$  is a tautology then  $P_x$  is a tautology for some  $x: T$ , as long as the cardinality of  $T$  is not too large. In particular,  $\exists n: \mathbb{N}. \mathbb{P}_{\mathcal{E}}^n P$  is not equivalent to True in Transfinite Iris, unlike Iris. They also provide Coq mechanization of Transfinite Iris. However, in Transfinite Iris we lose many useful deduction rules of the original Iris, including commutativity of the later modality over the inhabited existential quantifier [LATER-COMM-INHEXIST](#) and the separating conjunction [LATER-COMM-SEP](#). Because of this, it remains unclear how we can model the *lifetime logic* in Transfinite Iris, which prevents us from using Transfinite Iris for RustHornBelt.

## Chapter 5

# A New Design of Semantic Types for Rust

*Semantically modeling the guarantees of program types* is one approach to proving correctness (e.g., reducibility, memory safety) of a type system in an *extensible* way. We call such a semantic model of a program type a *semantic type*. For a simple example, if we target a simple functional programming language instead of Rust, we can semantically model each program type as the *set of values* that the type accepts. In general, this semantic approach is more extensible to new features than common syntactic approaches, such as type soundness proof by progress and preservation.

When we target Rust, one challenge is that an object of a program type can have some *ownership* or *permission* on resources. For example, an object of `Box<i32>` is a pointer that owns a memory block of an integer value with permission to update and read the integer data and release the memory block. More subtly, a *unique reference* of the type `&'a mut i32` can update and read its target integer data *only while the lifetime 'a is alive*. This is exactly the challenge that *RustBelt* (Jung et al., 2018a) solved. They model the guarantees of program types as propositions in the separation logic Iris (Chapter 2). In order to handle borrows based on lifetimes, they built the lifetime logic (§2.3) on Iris.

For our work RustHornBelt, we have another challenge. We want to verify *functional correctness* of Rust programs, not only safety, unlike RustBelt. Moreover, we want to model each object of Rust as a *pure value* without irrelevant low-level information, even using *prophecy* to model unique references. For example, an object of the box pointer type `Box<i32>` should be modeled simply as its target integer value, without the information about the address. Furthermore, a unique reference `&'a mut i32` should be modeled as a pair of two integers, the current target value and the final target value that is prophesied.

For this purpose, we extend RustBelt’s approach using our new formulation of prophecy in Iris introduced in Chapter 3. A key idea is to assign to each Rust object a  $\pi$ -*parametrized pure value*, a pure value that is parametrized over assignments on prophecy variables, as we roughly discussed in §3.2.

In this chapter, we introduce our new design of semantic types for Rust. In §5.1, we give an overview and a formal definition of our notion of the semantic type. In §5.2, we present the semantic types for basic program types in Rust. In §5.3, we present the semantic type for the *unique reference type* in Rust, introducing some additional machinery on prophecy. In §5.4, we present the semantic type for the general recursive type.

Defining semantic types itself does not conclude the functional correctness proof for Rust programs. Later in Chapter 6, we introduce a type system that handles logic models, which we dub a *refined type system*, for Rust programs. It uses a typing judgment emitting a logic model, dubbed a refined typing judgment, which is modeled as persistent Iris propositions using the Iris predicates of the semantics types we define.

Related work of the combination of this chapter and [Chapter 6](#) is discussed in [§6.5](#). The function type is defined later in [§6.2.4](#), because it is modeled using the refined typing judgment.

## 5.1 Our Notion of the Semantic Type

Now we introduce our new notion of the semantic type for RustHornBelt, extending the approach of RustBelt ([Jung et al., 2018a](#)) with our formulation of prophecy introduced in [Chapter 3](#).

First, a semantic type for a Rust program type has an *Iris* predicate *Own* describing the *ownership* of the object of that type, which we call the *ownership predicate*. The predicate *Own* has the type  $LowVal \rightarrow (ProphAsn \rightarrow PureVal) \rightarrow IProp$ . Here, *LowVal* and *PureVal* are types of low-level values and pure values designed for the program type. The *Own* predicate takes as an argument a  $\pi$ -parametrized pure value  $\hat{v}$  instead of just a pure value, in order to handle prophecy in the style described in [§3.2](#).

For example, let us discuss the semantic type for the box pointer type `Box<i32>`. We can set the type *LowVal* to the address type and set the type *PureVal* to the integer type  $\mathbb{Z}$ . The ownership predicate *Own*  $\mathbb{I} \hat{v}$  for the box pointer type is set to  $\exists n \text{ s.t. } \hat{v} = \text{const } n. \mathbb{I} \xrightarrow{1} n$ , saying that  $\hat{v}$  is a constant function over some integer  $n$  and the object fully owns the address  $\mathbb{I}$  with the value  $n$ .

For another example, let us roughly discuss the semantic type for the the unique reference type `&'a mut i32`. We can set *LowVal* again to the address type and set *PureVal* to the *pair type of two integers*, in the style of RustHorn. Defining the ownership predicate *Own*  $\mathbb{I} \hat{v}$  for `&'a mut i32` is fairly challenging. Roughly speaking, it contains existential quantification like  $\exists x \text{ s.t. } (.1) \circ \hat{v} = \lambda \pi. \pi x. \dots$ , checking that there is some prophecy variable  $x$  that represents the final target value on  $\hat{v}$ , and the full prophecy token  $[x]_1$  is in some sense owned by the unique reference. Later this part is elaborated in [§5.3](#).

The story does not end here. In order to deal with shared references `&'a T` in general, we equip the semantic type with another predicate *Share*:  $Lft \rightarrow LowVal \rightarrow PureVal \rightarrow IProp$ , representing the condition to *share* an object of the program type with some low-level and pure values. We need this because in Rust a program type works quite differently under the ownership (unique permission) and the sharing permission. For example, a shared reference to a unique reference `&'b &'a mut i32` is copyable but only has the sharing permission on the inner integer data (as explained by [Example 1.8](#)), unlike a bare unique reference `&'a mut i32`. The idea of using two predicates *Own* and *Share* for unique and sharing permissions comes from RustBelt.

Our actual formulation is a bit different from the explanation above. As a technical workaround to mitigate *later modalities* in *Iris*, we add to *Own* and *Share* an extra natural-number parameter dubbed the *depth*. Also, we use symbolic notation  $\bar{a} \triangleleft \tau \{\hat{v}\}_d$ , which corresponds to something like  $Own_{\tau} \bar{a} \hat{v} d$  in the explanation above, and write  $\bar{a} \triangleleft^{\alpha} \tau \{\hat{v}\}_d$  for the *Share* counterpart.

For simplicity and flexibility, in this thesis, we formalize program types only in terms of *semantic* types, not introducing any fixed *syntax* of program types. We can do so particularly because later (in [Chapter 6](#)) we formulate our deductive type system for Rust only in semantic terms, without introducing any syntactic structures about deduction.

**Formal Definition of a Semantic Type** We write  $List_n \text{ CellVal}$  for the subtype  $\{\bar{a}: List \text{ CellVal} / \text{len } \bar{a} = n\}$ , i.e., the type of a list of cell values that has the length  $n$ .

A semantic type  $\tau : \text{SemTy}$  is a quintuple

$$\begin{aligned} & ( \lfloor \tau \rfloor : \text{Type}, \quad |\tau| : \mathbb{N}, \quad - \sqsubseteq \tau : \text{Lft} \rightarrow \text{IProp}, \\ & - \blacktriangleleft \tau \{-\}_- : \text{List}_{|\tau|} \text{CellVal} \rightarrow (\text{ProphAsn} \rightarrow \lfloor \tau \rfloor) \rightarrow \mathbb{N} \rightarrow \text{IProp} \\ & - \blacktriangleleft^\tau \tau \{-\}_- : \text{List}_{|\tau|} \text{CellVal} \rightarrow \text{Lft} \rightarrow (\text{ProphAsn} \rightarrow \lfloor \tau \rfloor) \rightarrow \mathbb{N} \rightarrow \text{IProp} ) \end{aligned}$$

that satisfies the following properties.

$$\begin{array}{c} \text{SEM-TY-OUTLV-PERS} \quad \frac{\text{SEM-TY-OUTLV-MONO} \quad \alpha \sqsubseteq \tau \quad \beta \sqsubseteq \alpha}{\beta \sqsubseteq \tau} \quad \frac{\text{SEM-TY-OWN-MONO} \quad d \leq d'}{\bar{\mathbf{a}} \blacktriangleleft \tau \{\hat{v}\}_d \Rightarrow \bar{\mathbf{a}} \blacktriangleleft \tau \{\hat{v}\}_{d'}} \\ \\ \text{SEM-TY-SHR-PERS} \quad \frac{\text{SEM-TY-SHR-MONO} \quad \bar{\mathbf{a}} \blacktriangleleft^\alpha \tau \{\hat{v}\}_d \quad \beta \sqsubseteq \alpha \quad d \leq d'}{\bar{\mathbf{a}} \blacktriangleleft^\beta \tau \{\hat{v}\}_{d'}} \\ \\ \text{SEM-TY-OWN-SHR} \quad \frac{\alpha \sqsubseteq \tau \quad \alpha \sqsubseteq \beta}{\&_{\text{full}}^\beta (\bar{\mathbf{a}} \blacktriangleleft \tau \{\hat{v}\}_d) * [\alpha]_q \Rightarrow_{\mathcal{N}_{\text{fit}}}^d \bar{\mathbf{a}} \blacktriangleleft^\beta \tau \{\hat{v}\}_d * [\alpha]_q} \\ \\ \text{SEM-TY-OWN-PROPH-TOKEN} \quad \frac{\alpha \sqsubseteq \tau}{\bar{\mathbf{a}} \blacktriangleleft \tau \{\hat{v}\}_d * [\alpha]_q \Rightarrow_{\mathcal{N}_{\text{fit}}}^d \exists X \text{ s.t. } \text{Dep}(\hat{v}, X). \exists q'. [X]_{q'} * ([X]_{q'} \Rightarrow_{\mathcal{N}_{\text{fit}}} \bar{\mathbf{a}} \blacktriangleleft \tau \{\hat{v}\}_d * [\alpha]_q)} \\ \\ \text{SEM-TY-SHR-PROPH-TOKEN} \quad \frac{\alpha \sqsubseteq \tau \quad \alpha \sqsubseteq \beta}{\bar{\mathbf{a}} \blacktriangleleft^\beta \tau \{\hat{v}\}_d * [\alpha]_q \Rightarrow_{\mathcal{N}_{\text{fit}}}^d \exists X \text{ s.t. } \text{Dep}(\hat{v}, X). \exists q'. [X]_{q'} * ([X]_{q'} \Rightarrow_{\mathcal{N}_{\text{fit}}} \bar{\mathbf{a}} \blacktriangleleft^\beta \tau \{\hat{v}\}_d * [\alpha]_q)} \end{array}$$

A semantic type  $\tau$  consists of: (i) the *pure type*  $\lfloor \tau \rfloor$ , which is used for values modeling data of  $\tau$ ; (ii) the *size*  $|\tau| : \text{int}$ , which is the number of memory cells used for data typed  $\tau$  at the shallowest level; (iii) the *outliving predicate*  $\alpha \sqsubseteq \tau : \text{IProp}$  for  $\alpha : \text{Lft}$ , which is the persistent condition for the type  $\tau$  to outlive the lifetime  $\alpha$ ; (iv) the *ownership predicate*  $\bar{\mathbf{a}} \blacktriangleleft \tau \{\hat{v}\}_d$ , which is the condition to own data typed  $\tau$  that has the list of cell values  $\bar{\mathbf{a}}$  at the low level, is modeled as the  $\pi$ -parametrized value  $\hat{v}$ , and has the *depth*  $d$ ; and (v) the *sharing predicate*  $\bar{\mathbf{a}} \blacktriangleleft^\alpha \tau \{\hat{v}\}_d$ , which is the persistent condition to share data typed  $\tau$  associated with  $\bar{\mathbf{a}}$ ,  $\hat{v}$  and  $d$  until the lifetime  $\alpha$  ends.

We equip the semantic type with an outliving predicate  $\alpha \sqsubseteq \tau$  instead of directly giving the lifetime of the type. It enables us to use *later modalities*, which facilitates definition of the recursive type (§5.4).<sup>1</sup>

The outliving predicate  $\alpha \sqsubseteq \tau$  should be persistent (**SEM-TY-OUTLV-PERS**) and anti-monotone over  $\alpha$  (**SEM-TY-OUTLV-MONO**). The ownership predicate  $\bar{\mathbf{a}} \blacktriangleleft \tau \{\hat{v}\}_d$  should be monotone over  $d$  (**SEM-TY-OWN-MONO**). The sharing predicate  $\bar{\mathbf{a}} \blacktriangleleft^\alpha \tau \{\hat{v}\}_d$  should be persistent (**SEM-TY-SHR-PERS**) as well as anti-monotone over  $\alpha$  and monotone over  $d$  (**SEM-TY-SHR-MONO**). **SEM-TY-OWN-SHR** says that we should be able to transform a full borrow of the ownership predicate  $\&_{\text{full}}^\beta \bar{\mathbf{a}} \blacktriangleleft \tau \{\hat{v}\}_d$  into a sharing predicate  $\bar{\mathbf{a}} \blacktriangleleft^\beta \tau \{\hat{v}\}_d$  taking  $d$  logical steps, by temporarily depositing a partial token on  $\alpha \sqcap \beta$ . **SEM-TY-OWN-PROPH-TOKEN** says that, out of the ownership predicate  $\bar{\mathbf{a}} \blacktriangleleft \tau \{\hat{v}\}_d$ , we should be able to temporarily take a partial prophecy token  $[X]_{q'}$  on prophecy variables  $X$  that  $\hat{v}$  depends on, taking  $d$  logical steps, by temporarily depositing a partial token on  $\alpha$ . **SEM-TY-SHR-PROPH-TOKEN** is a similar rule for a sharing predicate. Note that the depth of an

<sup>1</sup> This idea was conceived by Jacques-Henri Jourdan.

object determines how many logical steps we can use for [SEM-TY-OWN-SHR](#), [SEM-TY-OWN-PROPH-TOKEN](#) and [SEM-TY-SHR-PROPH-TOKEN](#).

## 5.2 Basic Semantic Types

In this section, we introduce a number of basic semantic types.

We introduce notation for list slicing. For a list  $v$  and  $i$  such that  $0 \leq i < \text{len } v$ ,  $v[..i]$  and  $v[i..]$  denote  $[v[0], \dots, v[i-1]]$  and  $[v[i], \dots, v[\text{len } v - 1]]$  respectively. Also, for a list  $v$  and  $i, j$  such that  $0 \leq i \leq j < \text{len } v$ ,  $v[i..j]$  denotes  $[v[i], \dots, v[j-1]]$ .

**Integer Type** The (unbounded) *integer type* `int` is an idealization of Rust’s bounded integer types like `i32`. Note that in our low-level formulation one memory cell can contain any integer, idealizing the real-world hardware. We define `int` as the following semantic type.

$$\begin{aligned} \lfloor \text{int} \rfloor &:= \mathbb{Z} & |\text{int}| &:= 1 & \alpha \sqsubseteq \text{int} &:= \text{True} \\ [\mathbf{a}] \blacktriangleleft \text{int} \{\hat{v}\}_d &= [\mathbf{a}] \triangleleft^\alpha \text{int} \{\hat{v}\}_d &:= \exists n \text{ s.t. } n = \mathbf{a}. \hat{v} = \text{const } n \end{aligned}$$

We use an integer of the mathematical type  $\mathbb{Z}$  to model an object of the program type `int`. At the low level, `int` occupies one memory cell that has an integer value  $n$ . Any lifetime outlives `int`. The  $\pi$ -parametrized integer value for `int` should be a constant function that returns the expected integer  $n$ . Any depth is permitted for `int`.

**Boolean Type** The *boolean type* `bool`, which corresponds to Rust’s `bool`, is defined as follows, in a similar way to the integer type.

$$\begin{aligned} \lfloor \text{bool} \rfloor &:= \mathbb{B} & |\text{bool}| &:= 1 & \alpha \sqsubseteq \text{bool} &:= \text{True} \\ [\mathbf{a}] \blacktriangleleft \text{bool} \{\hat{v}\}_d &= [\mathbf{a}] \triangleleft^\alpha \text{bool} \{\hat{v}\}_d &:= \exists bl \text{ s.t. } bl = \mathbf{a}. \hat{v} = \text{const } bl \end{aligned}$$

To model boolean data, we use the mathematical boolean type  $\mathbb{B}$ , which has two values `tt` (true) and `ff` (false).

**Invalid-Data Type** We introduce the *invalid-data type*  $\zeta_n$  of the size  $n$ .

$$\begin{aligned} \lfloor \zeta_n \rfloor &:= \text{Unit} & |\zeta_n| &:= n & \alpha \sqsubseteq \zeta_n &:= \text{True} \\ \bar{\mathbf{a}} \blacktriangleleft \zeta_n \{\hat{v}\}_d &= \bar{\mathbf{a}} \triangleleft^\alpha \zeta_n \{\hat{v}\}_d &:= \text{True} \end{aligned}$$

The  $\pi$ -parametrized value  $\hat{v}$  used here has the type `ProphAsn`  $\rightarrow$  `Unit`, which has only one value `const ()`. In modeling Rust, we can use the invalid-data type  $\zeta_n$  to model data invalidated after a move operation and also to model uninitialized data obtained by memory allocation.<sup>2</sup> In particular, an object of a type  $\tau$  can be invalidated into an invalid object of the type  $\zeta_{|\tau|}$ .

**Box Pointer Type** The *box pointer type* `box`  $\tau$ , which corresponds to `Box<T>` in Rust, represents a pointer that targets an object of the type  $\tau$  and has *full ownership* on the target object, i.e., can read, update and delete it. It can be defined as follows.

<sup>2</sup> This is actually a simplification of the actual behavior of the Rust compiler, which handles validity flags on each variable apart from type information.

$$|\text{box } \tau| := \lfloor \tau \rfloor \quad |\text{box } \tau| := 1 \quad \alpha \sqsubseteq \text{box } \tau := \triangleright (\alpha \sqsubseteq \tau)$$

$$\begin{aligned} [\mathbf{a}] \blacktriangleleft \text{box } \tau \{\hat{v}\}_d &:= d > 0 * \exists \mathbf{I} \text{ s.t. } \mathbf{I} = \mathbf{a}. \exists \bar{\mathbf{b}}. \\ &\quad \mathbf{I} \xrightarrow{1} \bar{\mathbf{b}} * \text{Free}_{|\tau|}(\mathbf{I}) * \triangleright (\bar{\mathbf{b}} \blacktriangleleft \tau \{\hat{v}\}_{d-1}) \end{aligned}$$

$$\begin{aligned} [\mathbf{a}] \triangleleft^\alpha \text{box } \tau \{\hat{v}\}_d &:= d > 0 * \exists \mathbf{I} \text{ s.t. } \mathbf{I} = \mathbf{a}. \exists \bar{\mathbf{b}}. \\ &\quad *_{i} \&_{\text{frac}}^\alpha (\lambda q. \mathbf{I} + i \xrightarrow{q} \bar{\mathbf{b}}[i]) * \triangleright (\bar{\mathbf{b}} \triangleleft^\alpha \tau \{\hat{v}\}_{d-1}) \end{aligned}$$

The value for a box pointer is set to the value for its target object, without the address information. The outliving predicate of  $\text{box } \tau$  is  $\triangleright (\alpha \sqsubseteq \tau)$ , the outliving predicate of the target type under the later modality.

The ownership predicate of  $\text{box } \tau$  consists mainly of (i) the *full points-to token*  $\mathbf{I} \xrightarrow{1} \bar{\mathbf{b}}$ , which owns  $|\tau|$  consecutive memory cells starting at  $\mathbf{I}$  with the cell values  $\bar{\mathbf{b}}$ , (ii) the *free-right token*  $\text{Free}_{|\tau|}(\mathbf{I})$ , which owns the right to free the memory block of the size  $|\tau|$  starting at  $\mathbf{I}$ , and (iii)  $\triangleright (\bar{\mathbf{b}} \blacktriangleleft \tau \{\hat{v}\}_{d-1})$ , the ownership predicate on the target object under the later modality.

On the other hand, the sharing predicate consists mainly of (i) *fractured borrows* on the points-to tokens to each address  $*_{i} \&_{\text{frac}}^\alpha (\lambda q. \mathbf{I} + i \xrightarrow{q} \bar{\mathbf{b}}[i])$ ,<sup>3</sup> from which we can temporarily take out  $\mathbf{I} \xrightarrow{q} \bar{\mathbf{b}}$  for some  $q$  while the lifetime  $\alpha$  is alive, and (ii) the target sharing predicate under the later modality  $\triangleright (\bar{\mathbf{b}} \triangleleft^\alpha \tau \{\hat{v}\}_{d-1})$ .

The *contractiveness* introduced by the later modality in the outliving, owning and sharing predicates is useful in defining *recursive types* with self reference under  $\text{box}$  (e.g., the standard singly linked list type; see §5.4). The depth of a box pointer is set to 1 plus the depth of its target object, which gives us one logical step to strip off the later modality on the outliving, ownership and sharing predicate in proving `SEM-TY-OWN-SHR`, `SEM-TY-OWN-PROPH-TOKEN` and `SEM-TY-SHR-PROPH-TOKEN` of  $\text{box } \tau$ . Here we elaborate the proof of `SEM-TY-OWN-SHR` of  $\text{box } \tau$ .

*Proof of SEM-TY-OWN-SHR of box  $\tau$ .* Assume that the input is a full borrow under  $\alpha$  of the ownership predicate of a box pointer,  $\&_{\text{full}}^\alpha [\mathbf{a}] \blacktriangleleft \text{box } \tau \{\hat{v}\}_d$ . We freeze the inner parameters  $\mathbf{I}, \bar{\mathbf{b}}$  of the full borrow (`FULLBOR-FREEZE`) and then split it (`FULLBOR-SPLIT`) to get two full borrows, (i) a full borrow of the points-to token  $P := \&_{\text{full}}^\alpha (\mathbf{I} \xrightarrow{1} \bar{\mathbf{b}})$  and (ii) a full borrow of the target ownership predicate under later  $Q := \&_{\text{full}}^\alpha (\triangleright \bar{\mathbf{b}} \triangleleft^\alpha \tau \{\hat{v}\}_{d-1})$ . In one logical step, we strip off the later modality of the full borrow  $Q$  (`FULLBOR-UNLATER`) and also the outliving predicate on the target object. Using `SEM-TY-OWN-SHR` on the target type  $\tau$  in  $d - 1$  logical steps, we get the target sharing predicate. The full borrow  $P$  can be transformed into fractured borrows  $*_{i} \&_{\text{frac}}^\alpha (\lambda q. \mathbf{I} + i \xrightarrow{q} \bar{\mathbf{b}}[i])$  (by `FULLBOR-SPLIT` and `FULLBOR-FRACBOR`).  $\square$

**Shared Reference Type** The *shared reference type*  $\&_{\text{shr}}^\alpha \tau$ , which corresponds to  $\&' \mathbf{a}$  T in Rust, can be defined fairly easily using the sharing predicate of  $\tau$ .

$$|\&_{\text{shr}}^\alpha \tau| := \lfloor \tau \rfloor \quad |\&_{\text{shr}}^\alpha \tau| := 1 \quad \beta \sqsubseteq \&_{\text{shr}}^\alpha \tau := \beta \sqsubseteq \alpha \wedge \triangleright (\beta \sqsubseteq \tau)$$

$$\begin{aligned} [\mathbf{a}] \blacktriangleleft \&_{\text{shr}}^\alpha \tau \{\hat{v}\}_d &= [\mathbf{a}] \triangleleft^\beta \&_{\text{shr}}^\alpha \tau \{\hat{v}\}_d := \\ &d > 0 * \exists \mathbf{I} \text{ s.t. } \mathbf{I} = \mathbf{a}. \exists \bar{\mathbf{b}}. *_{i} \&_{\text{frac}}^\alpha (\lambda q. \mathbf{I} + i \xrightarrow{q} \bar{\mathbf{b}}[i]) * \triangleright (\bar{\mathbf{b}} \triangleleft^\alpha \tau \{\hat{v}\}_{d-1}) \end{aligned}$$

The ownership and sharing predicates of  $\&_{\text{shr}}^\alpha \tau$  is the same as the sharing predicate of  $\text{box } \tau$  with the lifetime  $\alpha$ . Interestingly, we can ignore the lifetime argument  $\beta$  of the sharing predicate; we can still prove `SEM-TY-OWN-SHR` because the ownership predicate

<sup>3</sup>We need to use  $*_{i} \&_{\text{frac}}^\alpha (\lambda q. \mathbf{I} + i \xrightarrow{q} \bar{\mathbf{b}}[i])$  instead of simpler  $\&_{\text{frac}}^\alpha (\lambda q. \mathbf{I} \xrightarrow{q} \bar{\mathbf{b}})$  because the lifetime logic does not have the split rule like `FULLBOR-SPLIT` for fractured borrows.

is already persistent. Note that `SEM-TY-OWN-PROPH-TOKEN` of  $\&_{\text{shr}}^\alpha \tau$  can be proved using `SEM-TY-SHR-PROPH-TOKEN` on the target type  $\tau$ .

The later modality in the outliving, ownership and sharing predicates are for contractiveness, just as in the box pointer type `box`  $\tau$ .

**Plain Reference Type** To aid verification, we introduce the following *plain reference type*  $\&_q \tau$  for a type  $\tau$  and a fraction  $q$ .

$$\begin{aligned} \lfloor \&_q \tau \rfloor &:= \lfloor \tau \rfloor & |\&_q \tau| &:= 1 & \alpha \sqsubseteq \&_q \tau &:= \triangleright (\alpha \sqsubseteq \tau) \\ [\mathbf{a}] \blacktriangleleft \&_q \tau \{ \hat{v} \}_d &:= d > 0 * \exists I \text{ s.t. } I = \mathbf{a}. \exists \bar{\mathbf{b}}. I \xrightarrow{q} \bar{\mathbf{b}} * \triangleright (\bar{\mathbf{b}} \blacktriangleleft \tau \{ \hat{v} \}_{d-1}) \\ [\mathbf{a}] \triangleleft^\alpha \&_q \tau \{ \hat{v} \}_d &:= \text{True} \end{aligned}$$

It is a reference to an object typed  $\tau$ . The pure value of a plain reference is set to the pure value of its target object. The ownership predicate of a plain reference mainly consists of (i) the points-to token to the target values of the fraction  $q$ ,  $I \xrightarrow{q} \bar{\mathbf{b}}$ , and (ii) the ownership predicate on the target object under later,  $\triangleright (\bar{\mathbf{b}} \blacktriangleleft \tau \{ \hat{v} \}_{d-1})$ .

A plain reference type is intended to be used at the top level of a type and not inside other types. Since we do not need to use the plain reference type under a shared reference, the sharing predicate of the type is set to the trivial proposition `True`.

The later modality in the outliving, ownership and sharing predicates helps compatibility with other pointer types such as the box pointer type `box`  $\tau$ .

We call a plain reference with the fraction 1 of the type  $\&_1 \tau$  a *full plain reference* and call a plain reference with any fraction  $q$  a *fractional plain reference*.

**Pair Type** The *pair type*  $\tau \times \tau'$  of the semantic types  $\tau$  and  $\tau'$ , which corresponds to Rust's `(T, T')`, is defined as follows.

$$\begin{aligned} \lfloor \tau \times \tau' \rfloor &:= \lfloor \tau \rfloor \times \lfloor \tau' \rfloor & |\tau \times \tau'| &:= |\tau| + |\tau'| & \alpha \sqsubseteq \tau \times \tau' &:= \alpha \sqsubseteq \tau \wedge \alpha \sqsubseteq \tau' \\ \bar{\mathbf{a}} \blacktriangleleft \tau \times \tau' \{ \hat{v} \}_d &:= \bar{\mathbf{a}}[..|\tau|] \blacktriangleleft \tau \{ (.0) \circ \hat{v} \}_d * \bar{\mathbf{a}}[|\tau|..] \blacktriangleleft \tau' \{ (.1) \circ \hat{v} \}_d \\ \bar{\mathbf{a}} \triangleleft^\alpha \tau \times \tau' \{ \hat{v} \}_d &:= \bar{\mathbf{a}}[..|\tau|] \triangleleft^\alpha \tau \{ (.0) \circ \hat{v} \}_d * \bar{\mathbf{a}}[|\tau|..] \triangleleft^\alpha \tau' \{ (.1) \circ \hat{v} \}_d \end{aligned}$$

For  $i = 0, 1$ ,  $(.i)$  denotes the function  $\lambda v.v.i$  and thus  $(.i) \circ \hat{v}$  is equal to  $\lambda \pi.(\hat{v} \pi).i$ .

We piggyback the components  $(\lfloor \tau \rfloor, \alpha \sqsubseteq \tau, \text{etc.})$  of  $\tau$  and  $\tau'$  to construct the semantic type  $\tau \times \tau'$ . The list of cell values for data typed  $\tau \times \tau'$  is the concatenation of the list for  $\tau$  and the list for  $\tau'$ . The value for data typed  $\tau \times \tau'$  is set to the pair of the value for  $\lfloor \tau \rfloor$  and the value for  $\lfloor \tau' \rfloor$  (for each  $\pi$ ). The outliving predicate of  $\tau \times \tau'$  is the conjunction of those of  $\tau$  and  $\tau'$ . The ownership predicate for  $\tau \times \tau'$  is defined as the separating conjunction of those for  $\tau$  and  $\tau'$ , sharing the same depth. The sharing predicate is defined in a similar way. We can prove the required properties (`SEM-TY-OUTLV-PERS`, etc.) of  $\tau \times \tau'$  by using those of  $\tau$  and  $\tau'$ .

**Variant Type** The *variant type*  $\tau_0 + \tau_1$ , modeling Rust's `enum` type (e.g., `Result<T0, T1>`), is defined as follows.

$$\begin{aligned} \lfloor \tau_0 + \tau_1 \rfloor &:= \lfloor \tau_0 \rfloor + \lfloor \tau_1 \rfloor & |\tau_0 + \tau_1| &:= 1 + \max\{|\tau_0|, |\tau_1|\} \\ \alpha \sqsubseteq \tau_0 + \tau_1 &:= \alpha \sqsubseteq \tau_0 \wedge \alpha \sqsubseteq \tau_1 \\ \bar{\mathbf{a}} \blacktriangleleft \tau_0 + \tau_1 \{ \hat{v} \}_d &:= \exists i, \hat{w} \text{ s.t. } \bar{\mathbf{a}}[0] = i \wedge \hat{v} = \text{inj}_i \circ \hat{w}. \bar{\mathbf{a}}[1..1+|\tau_i|] \blacktriangleleft \tau_i \{ \hat{w} \}_d \\ \bar{\mathbf{a}} \triangleleft^\alpha \tau_0 + \tau_1 \{ \hat{v} \}_d &:= \exists i, \hat{w} \text{ s.t. } \bar{\mathbf{a}}[0] = i \wedge \hat{v} = \text{inj}_i \circ \hat{w}. \bar{\mathbf{a}}[1..1+|\tau_i|] \triangleleft^\alpha \tau_i \{ \hat{w} \}_d \end{aligned}$$

Here,  $i$  is either 0 or 1 and  $\hat{w}$  is a  $\pi$ -parametrized value of the type  $\text{ProphAsn} \rightarrow [\tau_i]$  (depending on  $i$ ). Note that  $\text{inj}_i \circ \hat{w}$  is equal to  $\lambda\pi. \text{inj}_i(\hat{w} \pi)$ .

**Vector Type** The *vector type*  $\text{vec } \tau$ , modeling Rust's `Vec<T>`, represents a dynamically allocated, growable array of objects typed  $\tau$ . It is defined as follows, extending the definition of the box pointer type  $\text{box } \tau$ .

$$[\text{vec } \tau] := \text{List } [\tau] \quad |\text{vec } \tau| := 3 \quad \alpha \sqsubseteq \text{vec } \tau := \triangleright (\alpha \sqsubseteq \tau)$$

$$\begin{aligned} \bar{\mathbf{a}} \blacktriangleleft \text{vec } \tau \{ \hat{v} \}_d &:= d > 0 * \exists \mathbf{I}, \text{cap}, \text{len} \text{ s.t. } [\mathbf{I}, \text{cap}, \text{len}] = \bar{\mathbf{a}} \wedge \text{cap} \geq \text{len}. \\ &\exists \bar{\mathbf{b}} \text{ s.t. } \text{len } \bar{\mathbf{b}} = \text{cap} \cdot |\tau|. \mathbf{I} \mapsto \bar{\mathbf{b}} * \text{Free}_{\text{cap} \cdot |\tau|}(\mathbf{I}) * (\forall \pi. \text{len}(\hat{v} \pi) = \text{len}) * \\ &*_{i < \text{len}} \triangleright (\bar{\mathbf{b}}[i \cdot |\tau| .. (i+1) \cdot |\tau|] \blacktriangleleft \tau \{ \lambda\pi. (\hat{v} \pi)[i] \}_{d-1}) \end{aligned}$$

$$\begin{aligned} \bar{\mathbf{a}} \triangleleft^\alpha \text{vec } \tau \{ \hat{v} \}_d &:= d > 0 * \exists \mathbf{I}, \text{cap}, \text{len} \text{ s.t. } [\mathbf{I}, \text{cap}, \text{len}] = \bar{\mathbf{a}} \wedge \text{cap} \geq \text{len}. \\ &\exists \bar{\mathbf{b}} \text{ s.t. } \text{len } \bar{\mathbf{b}} = \text{cap} \cdot |\tau|. *_{i \text{ \&frac{\alpha}{\text{frac}}}} (\lambda q. \mathbf{I} + i \mapsto \bar{\mathbf{b}}[i]) * (\forall \pi. \text{len}(\hat{v} \pi) = \text{len}) * \\ &*_{i < \text{len}} \triangleright (\bar{\mathbf{b}}[i \cdot |\tau| .. (i+1) \cdot |\tau|] \triangleleft^\alpha \tau \{ \lambda\pi. (\hat{v} \pi)[i] \}_{d-1}) \end{aligned}$$

At the low level, a vector consists of three cell values — the head location of data  $\mathbf{I}$ , the capacity size  $\text{cap}$ , and the content length  $\text{len}$ . The value for  $\text{vec } \tau$  is set to the list of  $\text{len}$  elements.

The ownership predicate of the vector type consists mainly of (i) the full points-to token  $\mathbf{I} \mapsto \bar{\mathbf{b}}$  that owns  $\text{cap} \cdot |\tau|$  ( $= \text{len } \bar{\mathbf{b}}$ ) consecutive memory cells starting at  $\mathbf{I}$ , (ii) the free-right token  $\text{Free}_{\text{cap} \cdot |\tau|}(\mathbf{I})$  on the memory block of the size  $\text{cap} \cdot |\tau|$  starting at  $\mathbf{I}$ , and (iii) the (separating) conjunction of the ownership predicate for the  $\text{len}$  elements. The sharing predicate is obtained by some modification to the ownership predicate, in an analogous way to the box pointer type.

The later modality in the outliving, ownership and sharing predicates are for contractiveness, just as in the box pointer type  $\text{box } \tau$ .

### 5.3 Modeling the Unique Reference Type with Prophecy

The unique reference type  $\&_{\text{uniq}}^\alpha \tau$  is the key semantic type in our framework RustHorn-Belt. A unique reference is modeled as the  $\pi$ -parametrized value  $\lambda\pi. (\hat{v} \pi, \pi x)$ , where  $\hat{v}$  models the current target value and  $x$  is the prophecy variable for the borrow.

We should design the semantic type  $\&_{\text{uniq}}^\alpha \tau$  so that we can support various operations on unique references. First, when we release a unique reference modeled  $\lambda\pi. (\hat{v} \pi, \pi x)$ , we should be able to resolve the prophecy variable  $x$  to  $\hat{v}$ . On the other hand, we should also be able to throw away a unique reference *without* resolving the prophecy variable, letting the borrowee resolve the prophecy variable when it reclaims the borrowed object; we need this feature because Rust's type system allows an object to be *leaked* out of the static control, especially when we have circular references (Klabnik et al., 2018, §15.6). Moreover, we should be able to subdivide a unique reference in various ways (e.g., take  $\&_{\text{uniq}}^\alpha \tau$  out of  $\&_{\text{uniq}}^\alpha (\tau \times \tau')$ ) and also support unique reborrows.

We already have the *lifetime logic* (introduced in § 2.3) developed by Jung et al. (2018a), which supports various operations for borrows. We set up some machinery on the side of prophecy.

**Prophecy Equalizer** First, we introduce the following predicate dubbed a *prophecy equalizer*  $\text{PE}(\hat{w}, \hat{v})$  for  $\hat{w}, \hat{v}: \text{ProphAsn} \rightarrow T$  on some type  $T$ .

$$\text{PE}(\hat{w}, \hat{v}) := \forall X \text{ s.t. } \text{Dep}(\hat{v}, X). \forall q. [X]_q \Rightarrow_{\mathcal{N}_{\text{proph}}} \langle \pi. \hat{w} \pi = \hat{v} \pi \rangle * [X]_q$$

It can be explained as a *delayed* prophecy observation  $\langle \pi. \hat{w} \pi = \hat{v} \pi \rangle$ ; the acquisition of the prophecy observation is delayed until we get a partial prophecy token  $[X]_q$  on the dependency  $X$  of  $\hat{v}$ . This notion is useful for supporting *flexible timing of resolution*; we need it for RustHornBelt, because the prophecy variable of a unique borrow can be resolved by the unique reference and also by the borrowee.

We can use a prophecy equalizer in the following way.

$$\frac{\text{PROPH EQZ-USE} \quad \text{Dep}(\hat{v}, X)}{\text{PE}(\hat{w}, \hat{v}) * [X]_q \Rightarrow_{\mathcal{N}_{\text{proph}}} \langle \pi. \hat{w} \pi = \hat{v} \pi \rangle * [X]_q}$$

We can construct prophecy equalizers using the following lemmas.

$$\begin{array}{c} \text{PROPH EQZ-PROPH OBS} \quad \text{PROPH EQZ-MODIFY} \\ \langle \pi. \hat{w} \pi = \hat{v} \pi \rangle \quad \langle \pi. \hat{w}' \pi = \hat{w} \pi \rangle \\ \hline \text{PROPH EQZ-PROPH TOKEN} \quad \text{PE}(\hat{w}, \hat{v}) \Rightarrow \text{PE}(\hat{w}', \hat{v}) \\ [x]_1 \Rightarrow \text{PE}(\lambda \pi. \pi x, \hat{v}) \end{array}$$

$$\frac{\text{PROPH EQZ-T RANSFORM} \quad f \text{ is injective}}{*_{i < n} \text{PE}(\hat{w}_i, \hat{v}_i) \Rightarrow \text{PE}(\lambda \pi. f(\hat{w}_0 \pi, \dots, \hat{w}_{n-1} \pi), \lambda \pi. f(\hat{v}_0 \pi, \dots, \hat{v}_{n-1} \pi))}$$

**PROPH EQZ-PROPH TOKEN** says that, out of a full prophecy token  $[x]_1$ , we can create a prophecy equalizer for  $\lambda \pi. \pi x$  and any  $\pi$ -parametrized value  $\hat{v}$ ; it follows from **PROPH-RESOLVE**. **PROPH EQZ-PROPH OBS** says that a prophecy equalizer can be trivially obtained from a prophecy observation on the expected equality. **PROPH EQZ-MODIFY** allows us to modify the first argument of a prophecy equalizer using a prophecy observation. **PROPH EQZ-T RANSFORM** says that we can get a new prophecy equalizer out of prophecy equalizers, where both for the first and second arguments we construct the new value for the output by applying the same injective function  $f$  (which does not depend on  $\pi$ ) to the values of the inputs.

*Proof of PROPH EQZ-T RANSFORM.* We satisfy the dependency precondition of each of the input prophecy equalizers by **DEP-DESTRUCT**. Also, we split the partial prophecy token into  $n$  pieces to feed the prophecy equalizers. Then we compose the view shifts of them in parallel. Finally, we merge the prophecy observations and the partial prophecy tokens that were returned by them.  $\square$

**Value Observer and Prophecy Control** Now we introduce extra machinery for the unique reference type using the prophecy equalizer. First, we register the following RA **UNQ** to the global camera.

$$\text{UNQ} := (x: \text{ProphVar}) \xrightarrow{\text{fin}} \text{FRACOWN}((\text{ProphAsn} \rightarrow x.\text{type}) \times \mathbb{N})$$

In this RA, we assign to a finite number of prophecy variables a value of the type  $(\text{ProphAsn} \rightarrow x.\text{type}) \times \mathbb{N}$ , under fractional ownership by **FRACOWN**. A value assigned to a prophecy variable is a pair  $(\hat{v}, d)$  of a  $\pi$ -parametrized pure value  $\hat{v}$  and a depth  $d$ . Using this RA, taking some fixed ghost name  $\gamma_{\text{unq}}$ , we introduce the following predicates, a *value observer*  $\text{VO}_x(\hat{v}, d)$  and a *prophecy control*  $\text{PC}(x, \hat{v}, d)$ .

$$\begin{aligned} \text{VO}_x(\hat{v}, d) &:= \left[ \left[ x \leftarrow \text{fown}_{1/2}(\hat{v}, d) \right] \right]_{\text{UNQ}}^{\gamma_{\text{unq}}} \\ \text{PC}(x, \hat{v}, d) &:= \left( \left[ \left[ x \leftarrow \text{fown}_{1/2}(\hat{v}, d) \right] \right]_{\text{UNQ}}^{\gamma_{\text{unq}}} * [x]_1 \right) \vee \\ &\quad \left( (\exists a. \left[ \left[ x \leftarrow \text{fown}_1 a \right] \right]_{\text{UNQ}}^{\gamma_{\text{unq}}}) * \text{PE}(\lambda \pi. \pi x, \hat{v}) \right) \end{aligned}$$

A value observer  $\text{VO}_x(\hat{v}, d)$  witnesses with half ownership that  $(\hat{v}, d)$  is assigned to  $x$  by the  $\text{UNQ}$  RA at  $\gamma_{\text{unq}}$ . A prophecy control  $\text{PC}(x, \hat{v}, d)$  either (i) witnesses with half ownership that  $(\hat{v}, d)$  is assigned to  $x$  and owns the full prophecy token  $[x]_1$  or (ii) has full ownership on the value assigned to  $x$  and owns a prophecy equalizer on  $\lambda\pi. \pi x$  and  $\hat{v}$ . They satisfy the following properties.

$$\begin{array}{c}
\text{VALOBS-PROPHCTRL-TIMELESS} \\
\text{timeless}(\text{VO}_x(\hat{v}, d)) \qquad \text{timeless}(\text{VO}_x(\hat{v}, d) * \text{PC}(x, \hat{v}, d)) \\
\\
\text{VALOBS-PROPHCTRL-INTRO} \\
\text{True} \Rightarrow_{\emptyset} \exists x \text{ s.t. } x.\text{type} = T. \text{VO}_x(\hat{v}, d) * \text{PC}(x, \hat{v}, d) \\
\\
\text{VALOBS-PROPHCTRL-AGREE} \\
\text{VO}_x(\hat{v}, d) * \text{PC}(x, \hat{v}', d') \Rightarrow \hat{v} = \hat{v}' \wedge d = d' \\
\\
\text{VALOBS-PROPHCTRL-UPDATE} \\
\text{VO}_x(\hat{v}, d) * \text{PC}(x, \hat{v}, d) \Rightarrow_{\mathcal{N}_{\text{proph}}} \text{VO}_x(\hat{v}', d') * \text{PC}(x, \hat{v}', d') \\
\\
\text{VALOBS-PROPHCTRL-PROPHTOKEN} \\
\text{VO}_x(\hat{v}, d) * \text{PC}(x, \hat{v}, d) \Rightarrow \text{VO}_x(\hat{v}, d) * [x]_1 * ([x]_1 \multimap \text{PC}(x, \hat{v}, d)) \\
\\
\text{VALOBS-PROPHCTRL-RESOLVE} \\
\frac{\text{Dep}(\hat{v}, Y)}{\text{VO}_x(\hat{v}, d) * \text{PC}(x, \hat{v}, d) * [Y]_q \Rightarrow_{\mathcal{N}_{\text{proph}}} \langle \pi. \pi x = \hat{v} \pi \rangle * \text{PC}(x, \hat{v}, d) * [Y]_q} \\
\\
\text{VALOBS-PROPHCTRL-PRERESOLVE} \\
\frac{\text{Dep}(\hat{w}, Y)}{\text{VO}_x(\hat{v}, d) * \text{PC}(x, \hat{v}, d) * [Y]_q \Rightarrow_{\mathcal{N}_{\text{proph}}} \langle \pi. \pi x = \hat{w} \pi \rangle * [Y]_q * (\forall \hat{v}', d'. \text{PE}(\hat{w}, \hat{v}') \multimap \text{PC}(x, \hat{v}', d'))} \\
\\
\text{PROPHCTRL-PROPHEQZ} \\
\text{PC}(x, \hat{v}, d) \Rightarrow \text{PE}(\lambda\pi. \pi x, \hat{v})
\end{array}$$

A value observer is timeless. A prophecy control itself is not timeless because of the view shift of the prophecy equalizer. Still, the separating conjunction of a prophecy control and a value observer is timeless ([VALOBS-PROPHCTRL-TIMELESS](#)).

[VALOBS-PROPHCTRL-INTRO](#) says that, taking a fresh prophecy variable  $x$ , we can create a value observer  $\text{VO}_x(\hat{v}, d)$  and a prophecy control  $\text{PC}(x, \hat{v}, d)$ . When we have a value observer and a prophecy control on the same prophecy variable, the values of the two agree ([VALOBS-PROPHCTRL-AGREE](#)) and the values can be simultaneously updated ([VALOBS-PROPHCTRL-UPDATE](#)). We can temporarily take out the full prophecy token out of a prophecy control as long as we have a value observer on the same prophecy variable ([VALOBS-PROPHCTRL-PROPHTOKEN](#)).

[VALOBS-PROPHCTRL-RESOLVE](#) says that, consuming a prophecy control and a value observer on a prophecy variable  $x$ , we can resolve  $x$  to a value  $\hat{v}$  and reclaim a prophecy control with the value  $\hat{v}$ , with help of a partial prophecy token on a dependency  $Y$  of  $\hat{v}$ . We use this rule to resolve the prophecy variable when we release a unique reference (see [SWKN-UNQREF-RELEASE](#) in §6.3.3). [VALOBS-PROPHCTRL-PRERESOLVE](#) is an advanced variant of the previous rule; in this lemma, we first resolve  $x$  to some value  $\hat{w}$  and, after we get a prophecy equalizer  $\text{PE}(\hat{w}, \hat{v}')$  for some  $\hat{v}'$ , we achieve a prophecy control with the value  $\hat{v}'$  and any depth  $d'$ . We use this rule for subdividing a unique reference (e.g., [RFN-SPLIT-UNQREF-PAIR-R](#) in §6.3.4). [PROPHCTRL-PROPHEQZ](#) says that we can obtain a prophecy equalizer out of a prophecy control. We use this rule to reclaim a borrowed variable (see [WKN-UNQBOR-PLNREF](#) in §6.3.1 and [SWKN-END-LOCALFT-RECLAIM](#) in §6.1.2).

*Proof of VALOBS-PROPHCTRL-TIMELESS.* The left-hand disjunct of the prophecy control, the separating conjunction of  $\llbracket [x \leftarrow \text{fown}_{1/2}(\hat{v}, d)] \rrbracket_{\text{UNQ}}^{\text{Y}_{\text{unq}}}$  and  $[x]_1$ , is timeless. The prophecy control cannot take the right-hand disjunct, because that side contains the full ownership token  $\llbracket [x \leftarrow \text{fown}_1(\hat{v}, d)] \rrbracket_{\text{UNQ}}^{\text{Y}_{\text{unq}}}$ , which is timeless and conflicts with the half ownership token of the value observer.  $\square$

*Proof of VALOBS-PROPHCTRL-INTRO.* We take a fresh prophecy variable  $x$  of the type  $T$  that has not been used *both* by the PROPH RA at  $\gamma_{\text{proph}}$  and by the UNQ RA at  $\gamma_{\text{unq}}$ . We obtain  $[x]_1$  just like PROPHTOKEN-INTRO. We know that the type  $T$  is inhabited because we have  $\hat{v}: \text{ProphAsn} \rightarrow T$ . We newly assign  $(\hat{v}, d)$  to  $x$  in the UNQ RA to get a full ownership token  $\llbracket [x \leftarrow \text{fown}_1(\hat{v}, d)] \rrbracket_{\text{UNQ}}^{\text{Y}_{\text{unq}}}$ , which can be halved into two half ownership tokens. Therefore, we finally obtain  $\text{VO}_x(\hat{v}, d)$  and the right-hand disjunct of  $\text{PC}(x, \hat{v}, d)$ . (We can use the rule GOWN-UPD-TWO for this update.)  $\square$

*Proof of PROPH-RESOLVE.* By agreement on the UNQ RA, since we have a value observer  $\text{VO}_x(\hat{v}, d)$ , we know that the input prophecy control  $\text{PC}(x, \hat{v}, d)$  takes the *left-hand* disjunct. By consuming the full prophecy token  $[x]_1$  of the prophecy control, by PROPH-RESOLVE, we resolve  $x$  into  $\hat{v}$  to obtain a prophecy observation  $\langle \pi. \pi x = \hat{v} \pi \rangle$ . We obtain a prophecy control  $\text{PC}(x, \hat{v}, d)$  by constructing the *right-hand* disjunct: by consuming the half ownership token of the value observer, in the UNQ RA we change the value assigned to  $x$  into none and get the half ownership token for that; we get a prophecy equalizer from the prophecy observation we have just obtained by PROPHSEQZ-PROPHOBS.  $\square$

*Proof of VALOBS-PROPHCTRL-PRERESOLVE.* First, like VALOBS-PROPHCTRL-RESOLVE, we consume the full prophecy token  $[x]_1$  of the prophecy control to obtain a prophecy observation  $\langle \pi. \pi x = \hat{w} \pi \rangle$ . Also, out of two half ownership tokens of the value observer and the left-hand disjunct of the prophecy control, we get a full ownership token for the right-hand disjunct of a prophecy control. After we get a prophecy equalizer  $\text{PE}(\hat{w}, \hat{v}')$ , by PROPHSEQZ-MODIFY, conjoining it with the prophecy observation we have just obtained, we get a prophecy equalizer  $\text{PE}(\lambda \pi. \pi x, \hat{v}')$ . Therefore, we can construct a prophecy control  $\text{PC}(x, \hat{v}', d')$ .  $\square$

*Proof of PROPHCTRL-PROPHSEQZ.* When the input prophecy control takes the left-hand disjunct, by PROPHSEQZ-PROPHTOKEN we turn the full prophecy token  $[x]_1$  into a prophecy equalizer. When the prophecy control takes the right-hand disjunct, we can just use the prophecy equalizer it has.  $\square$

**Unique Reference Type** Finally we define the semantic type of the unique reference type  $\&_{\text{unq}}^\alpha \tau$ .

$$\llbracket \&_{\text{unq}}^\alpha \tau \rrbracket := \llbracket \tau \rrbracket \times \llbracket \tau \rrbracket \quad \llbracket \&_{\text{unq}}^\alpha \tau \rrbracket := 1 \quad \beta \sqsubseteq \&_{\text{unq}}^\alpha \tau := \beta \sqsubseteq \alpha \wedge \triangleright \beta \sqsubseteq \tau$$

$$[\mathbf{a}] \blacktriangleleft \&_{\text{unq}}^\alpha \tau \{\hat{v}\}_d := d > 0 * \exists \mathbf{I} \text{ s.t. } \mathbf{I} = \mathbf{a}. \exists x \text{ s.t. } (.1) \circ \hat{v} = \lambda \pi. \pi x.$$

$$\text{VO}_x((.0) \circ \hat{v}, d-1) * \&_{\text{full}}^\alpha (\exists \bar{\mathbf{b}}, \hat{v}', d'. \mathbf{I} \mapsto \bar{\mathbf{b}} * \bar{\mathbf{b}} \blacktriangleleft \tau \{\hat{v}'\}_{d'} * \bar{\Delta}(d'+1) * \text{PC}(x, \hat{v}', d'))$$

$$[\mathbf{a}] \triangleleft^\beta \&_{\text{unq}}^\alpha \tau \{\hat{v}\}_d := d > 0 * \exists \mathbf{I} \text{ s.t. } \mathbf{I} = \mathbf{a}. \exists \bar{\mathbf{b}}. \exists x \text{ s.t. } \text{Dep}((.1) \circ \hat{v}, \{x\}).$$

$$\&_{\text{frac}}^{\beta \cap \alpha} (\lambda q. [x]_q) * *_{\mathbf{i}} \&_{\text{frac}}^{\beta \cap \alpha} (\lambda q. \mathbf{I} + i \mapsto \bar{\mathbf{b}}[i]) * \triangleright (\bar{\mathbf{b}} \triangleleft^{\beta \cap \alpha} \tau \{(\cdot 0) \circ \hat{v}\}_{d-1})$$

The ownership predicate  $[\mathbf{a}] \blacktriangleleft \&_{\text{unq}}^\alpha \tau \{\hat{v}\}_d$  of  $\&_{\text{unq}}^\alpha \tau$  is quite tricky. The inner predicate of the full borrow is *existentially quantified* over the cell values  $\bar{\mathbf{b}}$ , the  $\pi$ -parametrized value  $\hat{v}'$  and the depth  $d'$  of the target object, because a unique reference

should be able to update the target object. In order to fix the inner parameters  $\hat{v}'$  and  $d'$  to the expected values  $(.0) \circ \hat{v}$  and  $d - 1$ , the ownership predicate has a value observer  $\text{VO}_x((.0) \circ \hat{v}, d - 1)$ . The inner predicate of the full borrow has the points-to token  $I \mapsto \bar{\mathbf{b}}$ , the ownership predicate on the target  $\bar{\mathbf{b}} \triangleleft \tau \{\hat{v}'\}_{d'}$ , the persistent time receipt  $\bar{\chi}(d'+1)$ , and the prophecy control  $\text{PC}(x, \hat{v}', d')$ . These objects are passed to the borrowee when the borrowee that the unique reference originates from reclaims the object. The borrowee obtains a prophecy observation  $\langle \pi. \pi x = \hat{v}' \pi \rangle$  by consuming the returned prophecy control by `PROPHCTRL-PROPHQZ` and `PROPHQZ-USE`. The full borrow introduces contractiveness and thus we set the depth  $d$  to 1 plus the depth of the target object.

In the sharing predicate of  $\&_{\text{unq}}^\alpha \tau$ , the  $\pi$ -parametrized pure value and the depth of the target are simply fixed to  $(.0) \circ \hat{v}$  and  $d - 1$ , and the cell values  $\bar{\mathbf{b}}$  are existentially quantified outside borrows; the sharing predicate has fractured borrows on the prophecy token  $[x]_q$  and on the memory ownership  $I \mapsto^q \bar{\mathbf{b}}$  as well as the sharing predicate on the target object. The value  $(.1) \circ \hat{v}$  can be anything as long as it depends only on the prophecy variable  $x$ .

*Proof of SEMTY-OWN-SHR of  $\&_{\text{unq}}^\alpha \tau$ .* Assume that the input is a full borrow under the lifetime  $\beta$  of the ownership predicate of  $\&_{\text{unq}}^\alpha \tau$ , i.e.,  $\&_{\text{full}}^\beta([\mathbf{a}] \triangleleft \&_{\text{unq}}^\alpha \tau \{\hat{v}\}_d)$ . After freeze by `FULLBOR-FREEZE` and split by `FULLBOR-SPLIT`, we obtain a full borrow of the full borrow of the ownership predicate of the form  $\&_{\text{full}}^\beta \&_{\text{full}}^\alpha \dots$ , which we unnest it into  $\&_{\text{full}}^{\beta \cap \alpha} \dots$  in one logical step by `FULLBOR-UNNEST`. We freeze its inner parameters  $\bar{\mathbf{b}}$ ,  $\hat{v}'$  and  $d'$  by `FULLBOR-FREEZE` and split it by `FULLBOR-SPLIT` to get full borrows of the points-to token  $\&_{\text{full}}^{\beta \cap \alpha}(I \mapsto \bar{\mathbf{b}})$ , of the ownership predicate of the target object  $\&_{\text{full}}^{\beta \cap \alpha}(\bar{\mathbf{b}} \triangleleft \tau \{\hat{v}'\}_{d'})$ , and of the prophecy control  $\&_{\text{full}}^{\beta \cap \alpha} \text{PC}(x, \hat{v}', d')$ . We also have a full borrow on the value observer  $\&_{\text{full}}^\beta \text{VO}_x((.0) \circ \hat{v}, d - 1)$ . By agreement of the value observer and the prophecy control (`VALOBS-PROPHCTRL-AGREE`),  $\hat{v}'$  and  $d'$  turn out to be equal to  $(.0) \circ \hat{v}$  and  $d - 1$ .

We merge the full borrows on the value observer and the prophecy control by `FULLBOR-MERGE` and then turn it into a full borrow on the full prophecy token  $\&_{\text{full}}^{\beta \cap \alpha} [x]_1$  by `FULLBOR-SUBDIV` and `VALOBS-PROPHCTRL-PROPHTOKEN`, which is turned into a fractured borrow  $\&_{\text{frac}}^{\beta \cap \alpha}(\lambda q. [x]_q)$  by `FULLBOR-FRACBOR`. Also, the full borrow on the points-to token is turned into a fractured borrow  $*_i \&_{\text{frac}}^{\beta \cap \alpha}(\lambda q. I + i \mapsto^q \bar{\mathbf{b}}[i])$  by `FULLBOR-SPLIT FULLBOR-FRACBOR`. The full borrow on the ownership predicate of the target object is transformed into the sharing predicate  $\bar{\mathbf{b}} \triangleleft^{\beta \cap \alpha} \tau \{( .0 ) \circ \hat{v} \}_{d-1}$  in  $d - 1$  logical steps by `SEMTY-OWN-SHR` on the target type  $\tau$ .  $\square$

*Proof of SEMTY-OWN-PROPHTOKEN of  $\&_{\text{unq}}^\alpha \tau$ .* We temporarily access the content of the full borrow of the ownership predicate of  $\&_{\text{unq}}^\alpha \tau$  by `FULLBOR-ACCESS`, stripping off the later modality using one logical step. From the prophecy control  $\text{PC}(x, \hat{v}', d')$  we can temporarily get the full token  $[x]_1$  by `VALOBS-PROPHCTRL-PROPHTOKEN`, with help of the value observer. Using  $d - 1$  logical steps, we obtain partial prophecy tokens on dependencies of  $\hat{v}'$  out of the resource of the target  $\bar{\mathbf{b}} \triangleleft \tau \{\hat{v}'\}_{d-1}$  by `SEMTY-OWN-PROPHTOKEN` of  $\tau$ .  $\square$

*Proof of SEMTY-SHR-PROPHTOKEN on  $\&_{\text{unq}}^\alpha \tau$ .* Accessing the content of the fractured borrow  $\&_{\text{frac}}^{\beta \cap \alpha}(\lambda q. [x]_q)$  by `FRACBOR-ACCESS`, we temporarily get  $[x]_q$  for some  $q$ . We strip off the later modality of the target resource in one logical step and obtain expected partial prophecy tokens out of the resource  $\bar{\mathbf{b}} \triangleleft^{\beta \cap \alpha} \tau \{( .0 ) \circ \hat{v} \}_{d-1}$  in  $d - 1$  logical steps by `SEMTY-SHR-PROPHTOKEN` on  $\tau$ .  $\square$

## 5.4 Recursive Type

In this section, we introduce the *recursive type*. We support *recursion by contractiveness*, which includes contravariant and invariant recursion as well as covariant recursion.

We write  $\text{SemTy}_{T,n}$  for the subtype  $\{ \tau : \text{SemTy} \mid \lfloor \tau \rfloor = T \wedge |\tau| = n \}$  for the semantics types of the pure type  $T$  and the size  $n$ .

The *recursive type*  $\text{rec}_{T,n} \sigma$  is defined for a pure type  $T : \text{Type}$ , a size  $n : \mathbb{N}$ , and a *contractive* function over semantic types  $\sigma : \text{SemTy}_{T,n} \rightarrow \text{SemTy}_{T,n}$ . We introduce some notions for defining and using the recursive type.

**Injection Type** For convenience, we introduce the following semantic type  $\text{in}(\tau, f)$ , which we dub the *injection type*, for any semantic type  $\tau$ , pure type  $T : \text{Type}$  and injective function  $f : \lfloor \tau \rfloor \rightarrow T$ .

$$\begin{aligned} \lfloor \text{in}(\tau, f) \rfloor &:= T & |\text{in}(\tau, f)| &:= |\tau| & \alpha \sqsubseteq \text{in}(\tau, f) &:= \alpha \sqsubseteq \tau \\ \bar{\mathbf{a}} \blacktriangleleft \text{in}(\tau, f) \{ \hat{\nu} \}_d &:= \exists \hat{\nu}' \text{ s.t. } \hat{\nu} = f \circ \hat{\nu}'. \quad \bar{\mathbf{a}} \blacktriangleleft \tau \{ \hat{\nu}' \}_d \\ \bar{\mathbf{a}} \triangleleft^\alpha \text{in}(\tau, f) \{ \hat{\nu} \}_d &:= \exists \hat{\nu}' \text{ s.t. } \hat{\nu} = f \circ \hat{\nu}'. \quad \bar{\mathbf{a}} \triangleleft^\alpha \tau \{ \hat{\nu}' \}_d \end{aligned}$$

[SEM-TY-OWN-PROPH-TOKEN](#) and [SEM-TY-SHR-PROPH-TOKEN](#) on  $\text{in}(\tau, f)$  follow from the corresponding properties on  $\tau$ , the injectivity of  $f$ , and [DEP-DESTRUCT](#).

The injection type satisfies the following rules on *type equality*.

$$\begin{array}{c} \text{TYEQ-INJTY-ID} \\ \tau = \text{in}(\tau, \text{id}) \end{array} \quad \begin{array}{c} \text{TYEQ-INJTY-COMP} \\ \text{in}(\tau, g \circ f) = \text{in}(\text{in}(\tau, f), g) \end{array} \quad \begin{array}{c} \text{TYEQ-INJTY-COMP-ID} \\ \frac{g \circ f = \text{id}}{\tau = \text{in}(\text{in}(\tau, f), g)} \end{array}$$

**Non-expansiveness and Contractiveness** We say a function over semantic types  $\sigma : \text{SemTy}_{T,n} \rightarrow \text{SemTy}_{U,m}$  (for some  $T, n, U, m$ ) is *non-expansive* and write  $\text{Nonex}(\sigma)$  if it is non-expansive as a function over the triple of the outliving, ownership and sharing predicates. We say  $\sigma$  is *contractive* and write  $\text{Contr}(\sigma)$  if it is contractive from that perspective.

We have the following rules on non-expansiveness and contractiveness.

$$\begin{array}{c} \text{NONEX-ID} \\ \text{Nonex}(\text{id}) \end{array} \quad \begin{array}{c} \text{CONTR-CONST} \\ \text{Contr}(\text{const } \tau) \end{array} \quad \begin{array}{c} \text{CONTR-NONEX} \\ \frac{\text{Contr}(\sigma)}{\text{Nonex}(\sigma)} \end{array} \quad \begin{array}{c} \text{CONTR-BOXPTR} \\ \frac{\text{Nonex}(\sigma)}{\text{Contr}(\text{box } \circ \sigma)} \end{array}$$

$$\begin{array}{c} \text{CONTR-SHRREF} \\ \frac{\text{Nonex}(\sigma)}{\text{Contr}(\&_{\text{shr}}^\alpha \circ \sigma)} \end{array} \quad \begin{array}{c} \text{CONTR-UNQREF} \\ \frac{\text{Nonex}(\sigma)}{\text{Contr}(\&_{\text{unq}}^\alpha \circ \sigma)} \end{array} \quad \begin{array}{c} \text{CONTR-VEC} \\ \frac{\text{Nonex}(\sigma)}{\text{Contr}(\text{vec } \circ \sigma)} \end{array}$$

$$\begin{array}{c} \text{NONEX-PAIR} \\ \frac{\forall i. \text{Nonex}(\sigma_i)}{\text{Nonex}(\lambda \tau. \sigma_0 \tau \times \sigma_1 \tau)} \end{array} \quad \begin{array}{c} \text{CONTR-PAIR} \\ \frac{\forall i. \text{Contr}(\sigma_i)}{\text{Contr}(\lambda \tau. \sigma_0 \tau \times \sigma_1 \tau)} \end{array} \quad \begin{array}{c} \text{NONEX-VRNT} \\ \frac{\forall i. \text{Nonex}(\sigma_i)}{\text{Nonex}(\lambda \tau. \sigma_0 \tau + \sigma_1 \tau)} \end{array}$$

$$\begin{array}{c} \text{CONTR-VRNT} \\ \frac{\forall i. \text{Contr}(\sigma_i)}{\text{Contr}(\lambda \tau. \sigma_0 \tau + \sigma_1 \tau)} \end{array} \quad \begin{array}{c} \text{NONEX-INJTY} \\ \frac{\text{Nonex}(\sigma)}{\text{Nonex}(\lambda \tau. \text{in}(\sigma \tau, f))} \end{array} \quad \begin{array}{c} \text{CONTR-INJTY} \\ \frac{\text{Contr}(\sigma)}{\text{Contr}(\lambda \tau. \text{in}(\sigma \tau, f))} \end{array}$$

The pointer types  $\text{box } \tau$ ,  $\&_{\text{shr}}^\alpha \tau$ ,  $\&_{\text{unq}}^\alpha \tau$  and  $\text{vec } \tau$  are all contractive. For the unique reference type  $\&_{\text{unq}}^\alpha \tau$ , the contractiveness comes from the contractiveness of the full borrow  $\&_{\text{full}}^\alpha$ , as well as the later modality in the outliving and sharing predicates.

**Recursive Type** We define the *recursive type*  $\text{rec}_{T,n} \sigma$  for a pure type  $T: \text{Type}$ , a size  $n: \mathbb{N}$ , and a *contractive* function over semantic types  $\sigma: \text{SemTy}_{T,n} \rightarrow \text{SemTy}_{T,n}$  as the semantic type that satisfies the following recursive equation.

$$\text{rec}_{T,n} \sigma := \sigma(\text{rec}_{T,n} \sigma)$$

The existence and uniqueness of the solution is ensured by the contractiveness of  $\sigma$  (by Banach's fixed point theorem [Theorem 2.1](#)). Note that  $\text{rec}_{T,n} \sigma$  is *equal* to  $\sigma(\text{rec}_{T,n} \sigma)$ .

We introduce the following notation for convenience.

$$\text{rec}_{T,n}(\sigma, f) := \text{rec}_{T,n}(\lambda \tau. \text{in}(\sigma \tau, f))$$

Also, we omit the size (i.e., use the notation  $\text{rec}_T \sigma$  and  $\text{rec}_T(\sigma, f)$ ) when the size information is clear.

The recursive type satisfies the following properties on *type equality*.

$$\begin{array}{ll} \text{TYEQ-REC} & \text{TYEQ-REC-INJTY} \\ \text{rec}_{T,n} \sigma = \sigma(\text{rec}_{T,n} \sigma) & \text{rec}_{T,n}(\sigma, f) = \text{in}(\sigma(\text{rec}_{T,n}(\sigma, f)), f) \end{array}$$

The recursive type satisfies the following non-expansiveness and contractiveness rules.

$$\begin{array}{ll} \text{NONEX-REC} & \text{CONTR-REC} \\ \frac{\forall \tau'. \text{Nonex}(\lambda \tau. \sigma \tau \tau')}{\text{Nonex}(\lambda \tau. \text{rec}_{T,n}(\sigma \tau))} & \frac{\forall \tau'. \text{Contr}(\lambda \tau. \sigma \tau \tau')}{\text{Contr}(\lambda \tau. \text{rec}_{T,n}(\sigma \tau))} \end{array}$$

**Example Recursive Types** We introduce and discuss a few example recursive types.

*Example 5.1* (List Type). The semantic type of the *singly linked list type*  $\text{list } \tau$  can be defined as the following recursive type, using the list type  $\text{List } T$  of the metalogic.

$$\text{list } \tau := \text{rec}_{\text{List}[\tau]}(\lambda \tau'. \text{!}_0 + \tau \times \text{box } \tau', \text{in}_{\text{List}})$$

Here, we define  $\text{in}_{\text{List}}: \text{Unit} + T \times \text{List } T \rightarrow \text{List } T$  as follows.

$$\text{in}_{\text{List}}(\text{inj}_0()) := \text{nil} \quad \text{in}_{\text{List}}(\text{inj}_1(v, w)) := v :: w$$

The contractiveness condition is satisfied because the self reference is under the box pointer  $\text{box}$ . The function  $\text{in}_{\text{List}}$  has the following inverse  $\text{out}_{\text{List}}$ .

$$\text{out}_{\text{List}} \text{nil} := \text{inj}_0() \quad \text{out}_{\text{List}}(v :: w) := \text{inj}_1(v, w)$$

The following properties hold on  $\text{list } \tau$ .

$$\begin{array}{l} \text{TYEQ-LIST} \\ \text{list } \tau = \text{in}(\text{!}_0 + \tau \times \text{box list } \tau, \text{in}_{\text{List}}) \\ \text{[list } \tau] = \text{List } [\tau] \quad |\text{list } \tau| = |\tau| + 2 \quad \alpha \sqsubseteq \text{list } \tau \Leftrightarrow \alpha \sqsubseteq \tau \\ \bar{\mathbf{a}} \triangleleft \text{list } \tau \{ \hat{\nu} \}_d \Leftrightarrow \bar{\mathbf{a}} \triangleleft \text{!}_0 + \tau \times \text{box list } \tau \{ \text{out}_{\text{List}} \circ \hat{\nu} \}_d \\ \bar{\mathbf{a}} \triangleleft^\alpha \text{list } \tau \{ \hat{\nu} \}_d \Leftrightarrow \bar{\mathbf{a}} \triangleleft^\alpha \text{!}_0 + \tau \times \text{box list } \tau \{ \text{out}_{\text{List}} \circ \hat{\nu} \}_d \end{array}$$

The property  $\alpha \sqsubseteq \text{list } \tau \Leftrightarrow \alpha \sqsubseteq \tau$  on the outliving predicate holds by Löb induction: let  $P$  be  $\alpha \sqsubseteq \text{list } \tau$ ; by definition,  $P$  is equivalent to  $\alpha \sqsubseteq \tau \wedge \triangleright P$ ; so  $\alpha \sqsubseteq \tau$  is equivalent to  $\triangleright P \Rightarrow P$ , which is equivalent to  $P$  by Löb induction ([LÖB](#)).

*Example 5.2* (Recursion Under the Unique Reference Type). We can also construct a recursive type with self reference under the unique reference type. For example, we have the following irregular Peano number type  $\text{mnat}_\alpha$  whose self reference is under a unique reference instead of a box pointer.

$$\text{mnat}_\alpha := \text{rec}_{\text{BNat}} (\lambda \tau. \zeta_0 + \&_{\text{uniq}}^\alpha \tau, \text{in}_{\text{BNat}})$$

Here, the metalogic type  $\text{BNat}$  is the following inductive binary tree type without node data, having the data constructor  $\text{in}_{\text{BNat}}$ .

$$\text{inductive BNat} := \text{in}_{\text{BNat}} : \text{BNat} + \text{BNat} \times \text{BNat} \rightarrow \text{BNat}$$

Note that the unique reference type  $\&_{\text{uniq}}^\alpha \tau$  is *not monotone* over  $\tau$ , which makes it impossible to use the least fixed point of *covariant induction* to model recursive types like  $\text{mnat}_\alpha$ .

*Example 5.3* (Recursion Under the Recursive Type). We can construct a recursion type with *self reference under the recursive type*. For example, we can construct the following *rose tree* type  $\text{rose } \tau$ .

$$\text{rose } \tau := \text{rec}_{\text{Rose}[\tau]} (\lambda \tau'. \tau \times \text{box list } \tau', \text{in}_{\text{Rose}})$$

Here, the metalogic type  $\text{Rose } T$  is the following inductive type with the data constructor  $\text{in}_{\text{Rose}}$ .

$$\text{inductive Rose } T := \text{in}_{\text{Rose}} : T \times \text{List } \text{Rose } T \rightarrow \text{Rose } T$$

The contractiveness of  $\lambda \tau'. \tau \times \text{box list } \tau'$  follows from [NONEX-INJT<sub>Y</sub>](#) and [NONEX-REC](#) and [CONTR-BOXPTR](#).

## Chapter 6

# A Refined Type System for Verifying Rust Programs

In order to formalize and verify type-based translation of Rust programs into clean logic models in the style of RustHorn, we introduce and semantically model a type system for Rust that handles logic models for functional correctness, which we dub the *refined type system*. In the refined type system, we mainly use a typing judgment equipped with a logic model, which we call a *refined typing judgment*. The refined type system is designed as a composable and flexible deduction system. Although the typing rules are highly non-deterministic, once we get a deduction tree for typing, we can compute the logic model in a straightforward way. Thanks to the *semantic approach* built on the separation logic Iris, we can prove soundness of the refined type system just by checking each deduction rule semantically, instead of checking the global behavior of the system.

For simplicity and flexibility, we formalize the refined typing judgment *only semantically* as a persistent predicate and treat each deduction rule as a lemma on them, without introducing any syntactic structures for reasoning. We also do not provide algorithms for type checking. In modeling the refined typing judgments, we use *semantic types* (Chapter 5) to model type information, *lifetime logic* (Jung et al., 2018a) (explained in §2.3) to handle lifetime information, and *prophecy observations*  $\langle \pi. \phi_\pi \rangle$  to describe the precondition and postcondition handling prophecy information.

The logic model of each operation is formalized as a (*backward*) *predicate transformer*, a function that maps each postcondition to some precondition, which is a standard technique used by various verification tools such as F\* (Swamy et al., 2016). For example, for the following expression of incrementing the target integer of a unique reference and then releasing the unique reference, we get the following refined typing judgment, where the last part is the predicate transformer (Example 6.3).

$$\alpha \mid \mathbf{a}: \&_{\text{uniq}}^\alpha \text{int} \vdash \mathbf{a} \leftarrow * \mathbf{a} + 1; \frac{}{\lambda} \mid \mid \lambda \text{post}, ((n, n_o)). n_o = n + 1 \Rightarrow \text{post} ()$$

Although the *input-output relation* is also a common representation for the postcondition used for verification methods like reduction to constrained Horn clauses (Grebenshchikov et al., 2012; Bjørner et al., 2015), we can lift an input-output relation into a predicate transformer in a natural way, as we see later in §6.1.2. We can translate Rust programs into CHCs in a fairly straightforward way using our refined type system, although we do not present a specific algorithm for that in this thesis for brevity.

In §6.1, we introduce the judgments used for the refined type system and basic rules on them. We also introduce the function type in §6.1.4. In §6.2, we introduce refined typing rules for basic types (excluding the unique reference type). In §6.3, we introduce refined typing rules for the unique reference type, giving detailed proofs for some of them. In §6.4, we present some examples of verifying functional correctness of Rust programs using the refined typing rules. In §6.5 we discuss some related work.

## 6.1 Judgments for the Refined Type System

We introduce the judgments used for the refined type system – the refined typing judgment and the weakening, super weakening, copyability, subtyping, and access judgments. We also introduce some basic rules on the judgments. These judgments are *semantically* modeled as persistent predicates in the separation logic Iris. We also introduce here the function type, which depends on the refined typing judgment.

### 6.1.1 Refined Typing Judgment and Its Variants

**Variants of the Ownership Predicate** We introduce the following variant of the ownership predicate which *hides the depth information*.

$$\bar{a} \blacktriangleleft \tau \{ \hat{v} \} := \exists d. \bar{a} \blacktriangleleft \tau \{ \hat{v} \}_d * \bar{\chi} d$$

Here, we can use any depth  $d$  for the ownership predicate  $\bar{a} \blacktriangleleft \tau \{ \hat{v} \}_d$  as long as we have a persistent time receipt  $\bar{\chi} d$ , which witnesses that at least  $d$  physical steps have been passed.

We also introduce the following predicate that represents the *resource for an object which is borrowed under the lifetime  $\alpha$* .

$$\bar{a} \blacktriangleleft^{\dagger\alpha} \tau \{ \hat{v} \} := [\dagger\alpha] \Rightarrow_{\mathcal{N}_{\text{ft}}} \exists \hat{v}'. \bar{a} \blacktriangleleft \tau \{ \hat{v}' \} * \triangleright \text{PE}(\hat{v}, \hat{v}')$$

After we know that the lifetime  $\alpha$  has ended, with update in one logical step, we get the resource of the ownership predicate with some value  $\hat{v}'$  and a prophecy equalizer under later  $\triangleright \text{PE}(\hat{v}, \hat{v}')$ . If we can temporarily access a relevant partial prophecy token by `SEM-TY-OWN-PROPH-TOKEN`, consuming the prophecy equalizer we can get a prophecy observation  $\langle \pi. \hat{v} \pi = \hat{v}' \pi \rangle$  by `PROPH-EQZ-USE`.

We introduce the notion of an *activity mode act*: Act, which is of the form  $!$  (unborrowed) or  $\dagger\alpha$  (borrowed under the lifetime  $\alpha$ : Lft). We define the following variant of the ownership predicate which is parametrized over an activity mode.

$$\bar{a} \blacktriangleleft^{\text{act}} \tau \{ \hat{v} \} := \begin{cases} \bar{a} \blacktriangleleft \tau \{ \hat{v} \} & (\text{act} = !) \\ \bar{a} \blacktriangleleft^{\dagger\alpha} \tau \{ \hat{v} \} & (\text{act} = \dagger\alpha) \end{cases}$$

**Type and Lifetime Contexts** A *type context*  $\Gamma$  is a sequence of items of the form  $\mathbf{a} :^{\text{act}} \tau$ , which represents an object of a cell value  $\mathbf{a}$ : CellVal, an activity mode  $\text{act}$ : Act, and a type  $\tau$ : SemTy<sub>1</sub>. Here, SemTy<sub>1</sub> denotes the subtype  $\{ \tau : \text{SemTy} \mid |\tau| = 1 \}$ , the type for a semantic type whose size is 1. For an item of a type context, we just write  $\mathbf{a} : \tau$  for  $\mathbf{a} :^! \tau$  (the case  $\text{act} = !$ ). Roughly speaking, a type context represents a sequence of ‘local variables’, but each ‘local variable’ is not given an explicit name.

A *lifetime context*  $\Delta$  is simply a sequence of lifetimes  $\vec{\alpha}$ , each of which represents a *local lifetime*, a lifetime such that we know it is alive and we have the right to end it. The Iris proposition  $[\Delta]$  for a lifetime context  $\Delta$  is defined as follows.

$$[\vec{\alpha}] := *_i ([\alpha_i]_1 * ([\alpha_i]_1 \Rightarrow_{\mathcal{N}_{\text{ft}}} [\dagger\alpha_i]))$$

For each lifetime  $\alpha_i$ , it has (i) the full token  $[\alpha_i]_1$  and (ii) the step-taking view shift  $[\alpha_i]_1 \Rightarrow_{\mathcal{N}_{\text{ft}}} [\dagger\alpha_i]$  which can end the lifetime  $\alpha_i$  by consuming the full token (see also the rule `LFT-INTRO`).

For a lifetime  $\alpha$  and a lifetime context  $\Delta = \vec{\beta}$ , we write  $\sqcap(\alpha, \Delta)$  for the intersection of  $\alpha, \vec{\beta}$ .

**Refined Typing Judgment** Now we introduce the *refined typing judgment*, which has the following form.

$$\alpha \mid \Delta; \Gamma \vdash e : \tau \mid \Delta'; \Gamma' \mid pre$$

It specifies the effect of executing the expression  $e$ . The input objects, i.e., the objects we have before the execution of  $e$ , are represented by  $\Gamma$ . The immediate output object of the expression  $e$  has the type  $\tau$ . The remaining output objects are represented by  $\Gamma'$ .

The lifetime contexts  $\Delta, \Delta'$  represent the set of local lifetimes we have before and after the execution of  $e$ . The lifetime  $\alpha$  at the head is the lifetime that remains alive during the execution of  $e$  aside from the local lifetimes. Unlike RustBelt (Jung et al., 2018a,b), we omit information about the outliving order  $\alpha \sqsubseteq \beta$  over lifetimes from this judgment; for reasoning about the outliving order, we piggyback the inference rules of the lifetime logic and Iris.

The last part  $pre$  represents the specification of the operation in terms of a *predicate transformer*, whose first argument is the postcondition on the pure values of the output objects and whose second argument is the tuple of the pure values of the input objects. If we have  $\Gamma = \overrightarrow{\mathbf{a} : \mathit{act} \tau'}$  and  $\Gamma' = \overrightarrow{\mathbf{b} : \mathit{act} \tau''}$ , the predicate transformer  $pre$  has the following type.

$$pre : (\times([\tau], [\tau'']) \rightarrow \text{Prop}) \rightarrow \times([\tau']) \rightarrow \text{Prop}$$

Here,  $\times(T_0, \dots, T_{n-1})$  denotes the tuple type  $T_0 \times \dots \times T_{n-1}$ .

The refined typing judgment is defined as the following persistent predicate.

$$\begin{aligned} \alpha \mid \Delta; \overrightarrow{\mathbf{a} : \mathit{act} \tau'} \vdash e : \tau \mid \Delta'; \overrightarrow{\mathbf{b} : \mathit{act} \tau''} \mid pre & := \quad \forall post, q. \\ \left\{ \exists \vec{\hat{v}}. \langle \pi. pre \ post (\vec{\hat{v}} \pi) \rangle * \ast_i ([\mathbf{a}_i] \leftarrow \mathit{act}_i \tau'_i \{ \hat{v}_i \}) * [\Delta] * [\alpha]_q \right\} e & \left\{ \mathbf{c}. \exists \hat{w}, \vec{\hat{v}}'. \right. \\ \left. \langle \pi. post (\hat{w} \pi, \vec{\hat{v}}' \pi) \rangle * [\mathbf{c}] \leftarrow \tau \{ \hat{w} \} * \ast_j ([\mathbf{b}_j] \leftarrow \mathit{act}'_j \tau''_j \{ \hat{v}'_j \}) * [\Delta'] * [\alpha]_q \right\}_\top \end{aligned}$$

The Hoare triple is universally quantified over the postcondition  $post$ , as well as the fraction  $q$  for the lifetime  $\alpha$ . The input objects have the  $\pi$ -parametrized pure values  $\hat{v}_i$ ; we know that they satisfy the precondition  $pre \ post$  for *any valid*  $\pi$  through the *prophecy observation*  $\langle \pi. pre \ post (\vec{\hat{v}} \pi) \rangle$ . The output objects have the  $\pi$ -parametrized pure values  $\hat{w}, \vec{\hat{v}}'$ , where  $\hat{w}$  is for the immediate output and  $\vec{\hat{v}}'$  are for the rest; we know that they satisfy the postcondition for *any valid*  $\pi$  through the prophecy observation  $\langle \pi. post (\hat{w} \pi, \vec{\hat{v}}' \pi) \rangle$ .

We also use the following variant of the refined typing judgment that ignores the immediate output of the expression  $e$ .

$$\begin{aligned} \alpha \mid \Delta; \Gamma \vdash e \mid \Delta'; \Gamma' \mid pre & := \\ \alpha \mid \Delta; \Gamma \vdash e : \downarrow_1 \mid \Delta'; \Gamma' \mid pre \circ (\lambda post, (\vec{v}). post ((\cdot), \vec{v})) \end{aligned}$$

Here, the function  $\lambda post, (\vec{v}). post ((\cdot), \vec{v})$  modifies the postcondition.

For a refined typing judgment, we can omit the parts ' $\Delta$ ;' and ' $\Delta'$ ;' if the input and output lifetime contexts  $\Delta, \Delta'$  are empty, and we can omit the first part  $\alpha \mid$  if  $\alpha$  is  $\infty$ .

The refined typing judgment satisfies the following adequacy theorem on integers.

**Theorem 6.1** (Adequacy of the Refined Typing Judgment on Integers). *The following is a tautology for any expression  $e$ , predicate transformer  $pre$  and postcondition  $post$ .*

$$\begin{aligned} \overrightarrow{\mathbf{a} : \text{int}} \vdash e : \text{int} \mid \mid pre & \Rightarrow \\ \left\{ \exists \vec{m} \text{ s.t. } \forall i. m_i = \mathbf{a}_i. pre \ post (\vec{m}) \right\} e & \left\{ \mathbf{b}. \exists n \text{ s.t. } n = \mathbf{b}. post \ n \right\}_\top \end{aligned}$$

*Proof.* By the definition of the integer type `int` (§5.2) and the rule `PROPHOBS-FACT`.  $\square$

**Weakening and Super Weakening Judgments** We also introduce the following *weakening judgment*, which amounts to the refined typing judgment *without execution*.

$$\alpha \mid \Delta; \Gamma \vdash \Delta'; \Gamma' \mid pre$$

We use a *predicate transformer*, instead of just a function over pure values, which allows us to impose a *precondition* on the weakening operation; it is later effectively used in the rules like [WKN-PLNREF-VRNT-TRIPLE](#). The weakening judgment is defined as follows, in a way analogous to the definition of the refined typing judgment.

$$\begin{aligned} \alpha \mid \Delta; \overrightarrow{\mathbf{a} :^{act} \tau'} \vdash \Delta'; \overrightarrow{\mathbf{b} :^{act'} \tau''} \mid pre &:= \forall post, q. \\ (\exists \vec{\hat{v}}. \langle \pi. pre \ post (\vec{\hat{v}} \pi) \rangle * \ast_i ([\mathbf{a}_i] \leftarrow^{act_i} \tau_i \{ \hat{v}_i \}) * [\Delta] * [\alpha]_q) &\Rightarrow_{\top} \\ \exists \vec{\hat{w}}, \vec{\hat{v}}'. \langle \pi. post (\vec{\hat{w}} \pi, \vec{\hat{v}}' \pi) \rangle * \ast_j ([\mathbf{b}_j] \leftarrow^{act'_j} \tau''_j \{ \hat{v}'_j \}) * [\Delta'] * [\alpha]_q \end{aligned}$$

We also introduce the following *super weakening judgment*.

$$\alpha \mid \Delta; \Gamma \vdash^{\#} \Delta'; \Gamma' \mid pre$$

It is defined as follows.

$$\begin{aligned} \alpha \mid \Delta; \overrightarrow{\mathbf{a} :^{act} \tau'} \vdash \Delta'; \overrightarrow{\mathbf{b} :^{act'} \tau''} \mid pre &:= \forall post, q. \\ (\exists \vec{\hat{v}}. \langle \pi. pre \ post (\vec{\hat{v}} \pi) \rangle * \ast_i ([\mathbf{a}_i] \leftarrow^{act_i} \tau_i \{ \hat{v}_i \}) * [\Delta] * [\alpha]_q) &\Rightarrow_{\top}^{\#} \\ \exists \vec{\hat{w}}, \vec{\hat{v}}'. \langle \pi. post (\vec{\hat{w}} \pi, \vec{\hat{v}}' \pi) \rangle * \ast_j ([\mathbf{b}_j] \leftarrow^{act'_j} \tau''_j \{ \hat{v}'_j \}) * [\Delta'] * [\alpha]_q \end{aligned}$$

Unlike the weakening judgment, we can use a *super fancy update*. We can perform a super weakening judgment on let binding ([RFN-LET](#)) and sequential execution ([RFN-SEQ](#)).

For weakening and super weakening judgments, we can omit the parts ‘ $\Delta$ ,’ and ‘ $\Delta'$ ,’ if the input and output lifetime contexts  $\Delta, \Delta'$  are empty, and we can omit the first part  $\alpha \mid$  if  $\alpha$  is  $\infty$ .

## 6.1.2 Basic Deduction Rules

**Structural Rules** We can modify a refined typing judgment both on the input and output sides using weakening judgments.

$$\begin{array}{c} \text{RFN-WKN} \\ \frac{\alpha \mid \Delta_0; \Gamma_0 \vdash e : \tau \mid \Delta_1; \Gamma_1 \mid pre \quad \forall \mathbf{a}. \alpha \mid \Delta_1; \Gamma_1 \vdash \Delta'_1; \Gamma'_1 \mid pre_1}{\alpha \mid \Delta'_0; \Gamma'_0 \vdash e : \tau \mid \Delta'_1; \Gamma'_1 \mid pre_0 \circ pre \circ pre_1} \end{array}$$

We compose the predicate transformers of the three judgments. Likewise, we can modify a super weakening judgment by weakening judgments.

$$\begin{array}{c} \text{SWKN-WKN} \\ \frac{\alpha \mid \Delta_0; \Gamma_0 \vdash^{\#} \Delta_1; \Gamma_1 \mid pre \quad \forall \mathbf{a}. \alpha \mid \Delta_1; \Gamma_1 \vdash \Delta'_1; \Gamma'_1 \mid pre_1}{\alpha \mid \Delta'_0; \Gamma'_0 \vdash^{\#} \Delta'_1; \Gamma'_1 \mid pre_0 \circ pre \circ pre_1} \end{array}$$

Also, weakening judgments can be composed.

$$\begin{array}{c} \text{WKN-COMP} \\ \frac{\alpha \mid \Delta; \Gamma \vdash \Delta'; \Gamma' \mid pre \quad \alpha \mid \Delta'; \Gamma' \vdash \Delta''; \Gamma'' \mid pre'}{\alpha \mid \Delta; \Gamma \vdash \Delta''; \Gamma'' \mid pre \circ pre'} \end{array}$$

We have the *frame rules* on the refined typing, weakening and super weakening judgments.

$$\frac{\text{RFN-FRAME} \quad \alpha \mid \Delta; \Gamma \vdash e: \tau \mid \Delta'; \Gamma' \mid pre}{\alpha \mid \Delta; \mathbf{a}:^{act} \tau, \Gamma \vdash e: \tau \mid \Delta; \mathbf{a}:^{act} \tau, \Gamma' \mid \lambda post, (v, \vec{w}). pre (\lambda(w', \vec{w}')). post (w', v, \vec{w}') (\vec{w})}$$

$$\frac{\text{WKN-FRAME} \quad \alpha \mid \Delta; \Gamma \vdash \Delta'; \Gamma' \mid pre}{\alpha \mid \Delta; \mathbf{a}:^{act} \tau, \Gamma \vdash \Delta'; \mathbf{a}:^{act} \tau, \Gamma' \mid \lambda post, (v, \vec{w}). pre (\lambda(\vec{w}')). post (v, \vec{w}') (\vec{w})}$$

$$\frac{\text{SWKN-FRAME} \quad \alpha \mid \Delta; \Gamma \vdash \Delta'; \Gamma' \mid pre}{\alpha \mid \Delta; \mathbf{a}:^{act} \tau, \Gamma \vdash^{\#} \Delta'; \mathbf{a}:^{act} \tau, \Gamma' \mid \lambda post, (v, \vec{w}). pre (\lambda(\vec{w}')). post (v, \vec{w}') (\vec{w})}$$

We can move a local lifetime  $\alpha$  to the head lifetime if we do not end  $\alpha$  during the execution.

$$\frac{\text{RFN-LFT-LOCAL-HEAD} \quad \beta \sqcap \alpha \mid \Delta; \Gamma \vdash e: \tau \mid \Delta'; \Gamma' \mid pre}{\alpha \mid \beta, \Delta; \Gamma \vdash e: \tau \mid \beta, \Delta'; \Gamma' \mid pre} \quad \frac{\text{WKN-LFT-LOCAL-HEAD} \quad \beta \sqcap \alpha \mid \Delta; \Gamma \vdash \Delta'; \Gamma' \mid pre}{\alpha \mid \beta, \Delta; \Gamma \vdash \beta, \Delta'; \Gamma' \mid pre}$$

$$\frac{\text{SWKN-LFT-LOCAL-HEAD} \quad \beta \sqcap \alpha \mid \Delta; \Gamma \vdash^{\#} \Delta'; \Gamma' \mid pre}{\alpha \mid \beta, \Delta; \Gamma \vdash^{\#} \beta, \Delta'; \Gamma' \mid pre}$$

We can modify the head lifetime and the predicate transformer.

$$\frac{\text{RFN-MONO} \quad \alpha \mid \Delta; \Gamma \vdash e: \tau \mid \Delta'; \Gamma' \mid pre \quad pre \Rightarrow pre' \quad \beta \sqsubseteq \alpha}{\beta \mid \Delta; \Gamma \vdash e: \tau \mid \Delta'; \Gamma' \mid pre'}$$

$$\frac{\text{WKN-MONO} \quad \alpha \mid \Delta; \Gamma \vdash \Delta'; \Gamma' \mid pre \quad pre \Rightarrow pre' \quad \beta \sqsubseteq \alpha}{\beta \mid \Delta; \Gamma \vdash \Delta'; \Gamma' \mid pre'}$$

$$\frac{\text{SWKN-MONO} \quad \alpha \mid \Delta; \Gamma \vdash^{\#} \Delta'; \Gamma' \mid pre \quad pre \Rightarrow pre' \quad \beta \sqsubseteq \alpha}{\beta \mid \Delta; \Gamma \vdash^{\#} \Delta'; \Gamma' \mid pre'}$$

Here, the relation  $pre' \Rightarrow pre$  denotes the following pointwise implication.

$$pre' \Rightarrow pre := \forall post, (\vec{v}). pre' post (\vec{v}) \Rightarrow pre post (\vec{v})$$

We can freely permute the type context, accordingly modifying the predicate transformer, by using the following rule as well as [WKN-FRAME](#) and [WKN-COMP](#) a number of times.

$$\frac{\text{WKN-SWAP} \quad \mathbf{a}:^{act} \tau, \mathbf{b}:^{act'} \tau', \Gamma \vdash \mathbf{b}:^{act'} \tau', \mathbf{a}:^{act} \tau, \Gamma \mid \lambda post, (v, v', \vec{w}). post (v', v, \vec{w})}{\alpha \mid \Delta; \Gamma \vdash \Delta'; \Gamma' \mid pre}$$

We can also permute the lifetime context.

$$\frac{\text{WKN-LFT-PERMUTE} \quad \Delta' \text{ is a permutation of } \Delta}{\Delta; \vdash \Delta'; \mid id}$$

We can throw away an object in the type context.

$$\begin{array}{c} \text{WKN-LEAK} \\ \mathbf{a} :^{act} \tau \vdash \mid \lambda post, (v). post () \end{array}$$

We can modify the lifetime limit of a borrowed object in the type context.

$$\begin{array}{c} \text{WKN-MONO-BORLFT} \\ \alpha \sqsubseteq \beta \\ \hline \mathbf{a} :^{\dagger\alpha} \tau \vdash \mid \mathbf{a} :^{\dagger\beta} \tau \mid \text{id} \end{array}$$

We have the reflexive super weakening judgment.

$$\begin{array}{c} \text{SWKN-REFL} \\ \alpha \mid \Delta; \Gamma \vdash^{\#} \Delta; \Gamma \mid \text{id} \end{array}$$

**Binding** We have the following binding rule based on an evaluation context.

$$\begin{array}{c} \text{RFN-EVCTX} \\ \alpha \mid \Delta; \Gamma \vdash e : \tau \mid \Delta'; \Gamma' \mid pre \quad \forall \mathbf{a}. \alpha \mid \Delta'; \mathbf{a} : \tau, \Gamma' \vdash K[\mathbf{a}] : \tau' \mid \Delta''; \Gamma'' \mid pre' \\ \hline \alpha \mid \Delta; \Gamma \vdash K[e] : \tau' \mid \Delta''; \Gamma'' \mid pre \circ pre' \end{array}$$

The cell value we get from  $e$  is *universally quantified* under the name  $\mathbf{a}$ . Note that the output type context  $\Gamma''$  cannot depend on the value  $\mathbf{a}$ . The resulting predicate transformer  $pre \circ pre'$  is simply the composite of the given predicate transformers  $pre, pre'$ .

Based on [RFN-EVCTX](#), we can introduce the following binding rule for let binding  $\text{let } a = e \text{ in } e'$ .

$$\begin{array}{c} \text{RFN-LET} \\ \alpha \mid \Delta; \Gamma \vdash e : \tau \mid \Delta'; \Gamma' \mid pre \quad \forall \mathbf{a}. \alpha \mid \Delta'; \mathbf{a} : \tau, \Gamma' \vdash^{\#} \Delta^+; \Gamma_a^+ \mid pre^+ \\ \square (P \Rightarrow \forall \mathbf{a}. \alpha \mid \Delta^+; \Gamma_a^+ \vdash e'[\mathbf{a}/a] : \tau' \mid \Delta''; \Gamma'' \mid pre') \quad \triangleright P \\ \hline \alpha \mid \Delta; \Gamma \vdash \text{let } a = e \text{ in } e' : \tau' \mid \Delta''; \Gamma'' \mid pre \circ pre^+ \circ pre' \end{array}$$

We replace the program variable  $a$  in the expression  $e'$  with the evaluation result  $\mathbf{a}$  of  $e$ . Using the physical step for let binding, we can perform the *super fancy update* of a *super weakening judgment*. We can also strip off the later modality on any (persistent) assumption  $P$ , which can be used for strengthening the outliving assumptions on a function (e.g., strengthening  $\alpha \sqsubseteq \text{box } \tau$  into  $\alpha \sqsubseteq \tau$ ).

Since sequential execution  $e; e'$  is defined as  $\text{let } \_ = e \text{ in } e'$ , we can derive the following rule for sequential execution from [RFN-LET](#).

$$\begin{array}{c} \text{RFN-SEQ} \\ \alpha \mid \Delta; \Gamma \vdash e \mid \Delta'; \Gamma' \mid pre \quad \alpha \mid \Delta'; \Gamma' \vdash^{\#} \Delta^+; \Gamma^+ \mid pre^+ \\ \square (P \Rightarrow \alpha \mid \Delta^+; \Gamma^+ \vdash e' : \tau' \mid \Delta''; \Gamma'' \mid pre') \quad \triangleright P \\ \hline \alpha \mid \Delta; \Gamma \vdash e; e' : \tau' \mid \Delta''; \Gamma'' \mid pre \circ pre^+ \circ pre' \end{array}$$

**Concurrency** We have the following refined typing rule for the concurrent execution  $e \parallel e'$ .

$$\begin{array}{c} \text{RFN-CONCUR} \\ \alpha \mid \Delta; \Gamma \vdash e : \tau \mid \Delta_+; \Gamma_+ \mid pre \quad \alpha \mid \Delta'; \Gamma' \vdash e' \mid \Delta'_+; \Gamma'_+ \mid pre' \\ \hline \alpha \mid \Delta, \Delta'; \Gamma, \Gamma' \vdash e \parallel e' : \tau \mid \Delta_+, \Delta'_+; \Gamma_+, \Gamma'_+ \mid \lambda post, (\vec{v}, \vec{v}'). \\ pre (\lambda(\vec{w}). \forall \vec{w}'. post(\vec{w}, \vec{w}'))(\vec{v}) \wedge pre' (\lambda(\vec{w}'). \forall \vec{w}. post(\vec{w}, \vec{w}'))(\vec{v}') \end{array}$$

When we have the postcondition  $post$  on the outputs of both  $e$  and  $e'$ , we can let the postcondition on the outputs of  $e$  be  $\lambda(\vec{w}). \forall \vec{w}'. post(\vec{w}, \vec{w}')$  and of  $e'$  be  $\lambda(\vec{w}'). \forall \vec{w}. post(\vec{w}, \vec{w})$ . We can prove this rule using [HOARE-CONCUR](#).

**Manipulating Local Lifetimes** We have the following rule for *introducing a new local lifetime*.

$$\frac{\text{RFN-INTRO-LOCALLFT} \quad \forall \alpha. \beta \vdash \alpha, \Delta; \Gamma \vdash e: \tau \mid \Delta'; \Gamma' \mid \text{pre}}{\beta \vdash \Delta; \Gamma \vdash e: \tau \mid \Delta'; \Gamma' \mid \text{pre}}$$

This rule follows from **LFT-INTRO**, which acquires a full lifetime token  $[\alpha]_1$  and a step-taking view shift  $[\alpha]_1 \Rightarrow_{\top} [\dagger\alpha]$  for *some* lifetime  $\alpha$ . To eliminate existential quantification over  $\alpha$ , we need the universal quantifier over the lifetime  $\alpha$  in the assumption of the rule **RFN-INTRO-LOCALLFT**. The lifetime and type contexts  $\Delta, \Delta', \Gamma, \Gamma'$  and the type  $\tau$  cannot depend on the lifetime  $\alpha$ ; in particular,  $\Delta'$  cannot contain  $\alpha$ .

We can borrow objects using local lifetimes. Later we introduce the rules for borrowing (e.g., **SWKN-SHRBOR-BOXPTR**, **WKN-UNQBOR-BOXPTR**, **WKN-UNQBOR-UNQREF**).

For a local lifetime  $\alpha$ , we have a full lifetime token  $[\alpha]_1$  and a *step-taking* view shift  $[\alpha]_1 \Rightarrow_{\mathcal{N}_{\text{fit}}} [\dagger\alpha]$ .

After we get  $[\dagger\alpha]$ , we can reclaim a borrowed object using the following lemma.

$$\frac{\text{RAW-RECLAIM} \quad [\dagger\alpha] \quad \beta \sqsubseteq \tau}{[\mathbf{a}] \blacktriangleleft^{\dagger\alpha} \tau \{ \hat{v} \} * [\beta]_q \Rightarrow_{\mathcal{N}_{\text{fit}}} \exists d. \sum d * \Rightarrow_{\mathcal{N}_{\text{fit}} + \mathcal{N}_{\text{proph}}}^{d+1} (\exists \hat{v}'. [\mathbf{a}] \blacktriangleleft \tau \{ \hat{v}' \} * \langle \pi. \hat{v} \pi = \hat{v}' \pi \rangle)}$$

Although the  $\pi$ -parametrized value changes from  $\hat{v}$  to  $\hat{v}'$ , the prophecy observation  $\langle \pi. \hat{v} \pi = \hat{v}' \pi \rangle$  tells us that  $\hat{v}$  and  $\hat{v}'$  are the same for valid prophecy assignments.

*Proof.* From  $[\dagger\alpha]$  and  $\alpha' \sqsubseteq \alpha$ , we obtain the dead-lifetime token  $[\dagger\alpha']$  by **LFTINCL-DEADLFT**. Using it, we turn the borrowed ownership predicate  $\bar{\mathbf{a}} \blacktriangleleft^{\dagger\alpha'} \tau \{ \hat{v} \}$  (defined in §6.1.1) into an ownership predicate  $\bar{\mathbf{a}} \blacktriangleleft \tau \{ \hat{v}' \}$  and a prophecy equalizer under later  $\triangleright \text{PE}(\hat{v}, \hat{v}')$ , for some  $\pi$ -parametrized value  $\hat{v}'$ . The ownership predicate can be decomposed into  $\bar{\mathbf{a}} \blacktriangleleft \tau \{ \hat{v}' \}_d$  and a persistent time receipt  $\sum d$  for some depth  $d$ . We spend one logical step to strip off the later modality on the prophecy equalizer  $\text{PE}(\hat{v}, \hat{v}')$ .

By **SEMPTY-OWN-PROPH-TOKEN**, we temporarily take out from the ownership predicate a partial prophecy token on some dependency  $X$  of  $\hat{v}'$  in  $d$  logical steps, with help of a partial lifetime token on  $\tau$ . By consuming the prophecy equalizer with help of the partial prophecy token by **PROPH-EQZ-USE**, we get a prophecy observation  $\langle \pi. \hat{v} \pi = \hat{v}' \pi \rangle$ .  $\square$

Using **RAW-RECLAIM**, we can prove the following refined typing rule for ending a local lifetime  $\alpha$  and reclaiming objects borrowed under  $\alpha$ .

$$\frac{\text{SWKN-END-LOCALLFT-RECLAIM} \quad \forall i. \beta \sqsubseteq \tau_i}{\beta \vdash \alpha; \bar{\mathbf{a}}: \overrightarrow{\dagger\alpha} \tau \vdash^{\#} \bar{\mathbf{a}}: \overrightarrow{\tau} \mid \text{id}}$$

**Input-Output Relations and Predicate Transformers** An *input-output relation* is a common representation of the *postcondition* of functions and programs, which is used in verification methods like reduction to constrained Horn clauses (**Grebenshchikov et al., 2012; Bjørner et al., 2015**).

We can lift an input-output relation  $R: T \rightarrow U \rightarrow \text{Prop}$  (where  $T$  and  $U$  respectively represent the input and the output) into the following (backward) predicate transformer  $[R]: (U \rightarrow \text{Prop}) \rightarrow T \rightarrow \text{Prop}$ .

$$[R] := \lambda \text{post}. v. \forall w. R v w \Rightarrow \text{post } w$$

Various operations on predicate transformers correspond to natural operations on input-output relations through the lift function.

$$[R] \circ [R'] = [\lambda u, w. \exists v. R u v \wedge R' v w]$$

$$\lambda post, v. post(f v) = [\lambda v, w. w = f v]$$

$$\lambda post, v. \phi v \Rightarrow [R] post v = [\lambda v, w. \phi v \wedge R v w]$$

$$\lambda post, v. [R] post v \wedge [R'] post v = [\lambda v, w. R v w \vee R' v w]$$

Composition of predicate transformers amounts to relational composition of input-output relations. Passing a value transformed by a function  $f$  to the postcondition in the predicate transformer amounts to the relational graph of  $f$  used as the input-output relation. Adding an assumption  $\phi v$  on the input  $v$  in the predicate transformer amounts to adding the postcondition  $\phi v$  in the input-output relation. Conjunction of predicate transformers amounts to disjunction of input-output relations.

### 6.1.3 Other Judgments

**Copyability Judgment** We introduce the *copyability judgment*  $\tau \text{ copy}$ , which is defined as the following persistent predicate.

$$\tau \text{ copy} := \forall \bar{a}, \hat{v}, d, \alpha. \square (\bar{a} \blacktriangleleft \tau \{\hat{v}\}_d \Leftrightarrow \bar{a} \triangleleft^\alpha \tau \{\hat{v}\}_d)$$

It simply means that the the ownership and sharing predicates of the type  $\tau$  are equivalent.

By the persistence of the sharing predicate, it allows us to duplicate the ownership predicate, i.e., the following lemma holds.

$$\frac{\text{RAW-COPY} \quad \tau \text{ copy}}{\bar{a} \blacktriangleleft \tau \{\hat{v}\}_d \Rightarrow \bar{a} \blacktriangleleft \tau \{\hat{v}\}_d * \bar{a} \blacktriangleleft \tau \{\hat{v}\}_d}$$

Using this lemma, we can prove the following refined typing rule for copying an un-borrowed object of a copyable type.

$$\frac{\text{WKN-COPY} \quad \tau \text{ copy}}{\mathbf{a}: \tau \vdash \mathbf{a}: \tau, \mathbf{a}: \tau \mid \lambda post, (v). post(v, v)}$$

It follows from [RAW-COPY](#).

The integer type `int`, the invalid-data type  $\int_n$  and the shared reference type  $\&_{\text{shr}}^\alpha \tau$  are copyable ([COPY-INT](#), [COPY-INVALID](#), [COPY-SHRREF](#)). We have also structural copyability rules ([COPY-PAIR](#), [COPY-VRNT](#), [COPY-INJTYP](#)).

Types like the box pointer type `box`  $\tau$ , the vector type `vec`  $\tau$  and the unique reference type  $\&_{\text{unq}}^\alpha \tau$  are not copyable.

**Subtyping Judgment** We introduce the notion of an *access mode*  $acc$ : `Acc`, which is either `own` (ownership) or `shr`( $\alpha$ ) (sharing under the lifetime  $\alpha$ ).

The *subtyping judgment* has the following form.

$$\tau \sqsubseteq^{acc} \tau' \mid f$$

We transform an object typed  $\tau$  into an object typed  $\tau'$  using the *value transformer*  $f$ :  $[\tau] \rightarrow [\tau']$ . When we apply subtyping, letting  $v$  be the pure value of the input object,

the pure value of the output object is  $f v$ . We also require that the types  $\tau, \tau'$  should have the same size (i.e., satisfy  $|\tau| = |\tau'|$ ).

Now the subtyping judgment is defined as the follows.

$$\begin{aligned}\tau \sqsubseteq^{\text{own}} \tau' \mid f & := \forall \bar{a}, \hat{v}, d. \square (\bar{a} \blacktriangleleft \tau \{ \hat{v} \}_d \Rightarrow \bar{a} \blacktriangleleft \tau' \{ f \circ \hat{v} \}_d) \\ \tau \sqsubseteq^{\text{shr}(\alpha)} \tau' \mid f & := \forall \bar{a}, \hat{v}, d. \square (\bar{a} \blacktriangleleft^\alpha \tau \{ \hat{v} \}_d \Rightarrow \bar{a} \blacktriangleleft^\alpha \tau' \{ f \circ \hat{v} \}_d)\end{aligned}$$

When the access mode *acc* is *own*, we work on the ownership predicate, and when it is *shr*( $\alpha$ ), we work on the sharing predicate with the lifetime set to  $\alpha$ .

By applying the subtyping judgment, we can modify the type and value of an object in the type context.

$$\frac{\text{WKN-SUBTY} \quad \tau \sqsubseteq^{\text{own}} \tau' \mid f}{\mathbf{a} : {}^{\text{act}} \tau \vdash \mathbf{a} : {}^{\text{act}} \tau' \mid \lambda \text{post}, (v). \text{post} (f v)}$$

Any type is a subtype of itself under the identity transformer.

$$\frac{\text{SUBTY-ID}}{\tau \sqsubseteq^{\text{acc}} \tau \mid \text{id}}$$

We can compose subtyping judgments, accordingly composing the value transformers.

$$\frac{\text{SUBTY-COMP} \quad \tau \sqsubseteq^{\text{acc}} \tau' \mid f \quad \tau' \sqsubseteq^{\text{acc}} \tau'' \mid g}{\tau \sqsubseteq^{\text{acc}} \tau'' \mid g \circ f}$$

**Access Judgment** We introduce the following *access judgment*.

$$\alpha \mid \mathbf{a} : \tau_+ / \tau'_+ \vdash e : q. \tau^{(q)} / \tau'^{(q)} \mid \text{get}; \text{set}$$

The meaning of the expression is as follows. First we have an object of the cell value  $\mathbf{a}$  and the type  $\tau_+$ . Out of this object, by performing the expression  $e$ , we can take out a *sub-object* of the type  $\tau_q$  for some fraction  $q$ . (The part ' $q$ .' binds the variable  $q$ .) The fraction can be defined by the *plain reference type* introduced above. We can freely update the sub-object. After we get back the sub-object with a new type  $\tau'(q)$ , we retrieve the original object with the updated type  $\tau'_+$ . The initial value of the sub-object is specified by the *getter* function  $\text{get} : [\tau_+] \rightarrow [\tau(1)]$ , which takes the pure value of the original object. The value of the retrieved object is specified by the *setter* function  $\text{set} : [\tau_+] \rightarrow [\tau'(1)] \rightarrow [\tau'_+]$ , which takes the pure value of the original object and the returned sub-object. The use of the getter and setter functions here is closely related to the notion of *functional lenses* used in functional programming (Foster et al., 2007).

Here, we assume that all the types  $\tau^+, \tau'^+, \tau(q), \tau'(q)$  have the size 1 (regardless of  $q$ ) and that the pure types  $[\tau(q)]$  and  $[\tau'(q)]$  are constant over the fraction  $q$ . The lifetime  $\alpha$  is ensured to be alive during this access.

The access judgment is defined as the following persistent predicate.

$$\begin{aligned}\alpha \mid \mathbf{a} : \tau_+ / \tau'_+ \vdash e : q. \tau^{(q)} / \tau'^{(q)} \mid \text{get}; \text{set} & := \forall q'. \\ \{ [\mathbf{a}] \blacktriangleleft \tau_+ \{ \hat{v} \} * [\alpha]_{q'} \} e \{ \mathbf{b}. \exists q. [\mathbf{b}] \blacktriangleleft \tau(q) \{ \lambda \pi. \text{get} (\hat{v} \pi) \} * \\ \forall \hat{w}. ([\mathbf{b}] \blacktriangleleft \tau'(q) \{ \hat{w} \} \Rightarrow_{\top} [\mathbf{a}] \blacktriangleleft \tau'_+ \{ \lambda \pi. \text{set} (\hat{v} \pi) (\hat{w} \pi) \} * [\alpha]_{q'} ) \} \}_{\top}\end{aligned}$$

We introduce the following shorthand on the access judgment. We can omit the part ' $\alpha \mid$ ' when the lifetime  $\alpha$  is  $\infty$ . We can omit the part ' $q$ .' if the types  $\tau(q)$  and  $\tau'(q)$  do not depend on the fraction  $q$ .

We can temporarily take out a full plain reference out of a box pointer and a unique reference ([ACC-BOXPTR-PLNREF](#), [ACC-UNQREF-PLNREF](#)) and a fractional plain reference out of a shared reference ([ACC-SHRREF-PLNREF](#)). We can also subdivide plain references (e.g., [ACC-DEREF-PLNREF-BOXPTR](#), [ACC-PLNREF-PAIR-L](#)).

By the following rule, we can use an access judgment whose expression  $e$  is just the original address  $\mathbf{a}$ .

$$\begin{array}{c} \text{RFN-ACC-SAME} \\ \alpha \mid \mathbf{a} : \tau_+ / \tau'_+ \vdash \mathbf{a} : q. \tau^{(q)} / \tau'(q) \mid \text{get}; \text{set} \\ \forall q. \alpha \mid \Delta; \mathbf{a} : \tau(q), \Gamma \vdash e_* : \tau_* \mid \Delta'; \mathbf{a} : \tau'(q), \Gamma' \mid \text{pre} \\ \hline \alpha \mid \Delta; \mathbf{a} : \tau_+, \Gamma \vdash e_* : \tau_* \mid \Delta'; \mathbf{a} : \tau'_+, \Gamma' \mid \\ \lambda \text{post}, (v, \vec{w}). \text{pre} (\lambda (v', \vec{w}'). \text{post} (\text{set } v \ v', \vec{w}')) (\text{get } v, \vec{w}) \end{array}$$

The assumption refined typing judgment is universally quantified over the fraction  $q$  and has  $\mathbf{a} : \tau(q)$  in the inputs and  $\mathbf{a} : \tau'(q)$  in the outputs. The predicate transformer  $\text{pre}$  of the assumption judgment is modified using the getter and setter functions. Note that we can use the reverse direction of the rule [HOARE-VAL](#) to prove this rule.

The following rule is a variant of [RFN-ACC-SAME](#) for an access judgment whose expression may not be just the original value.

$$\begin{array}{c} \text{RFN-ACC-EXPR} \\ \alpha \mid \mathbf{a} : \tau_+ / \tau'_+ \vdash e : q. \tau^{(q)} / \tau'(q) \mid \text{get}; \text{set} \\ \forall \mathbf{b}, q. \alpha \mid \Delta; \mathbf{b} : \tau(q), \Gamma \vdash K_*[\mathbf{b}] : \tau_* \mid \Delta'; \mathbf{b} : \tau'(q), \Gamma' \mid \text{pre} \\ \hline \alpha \mid \Delta; \mathbf{a} : \tau_+, \Gamma \vdash K_*[e] : \tau_* \mid \Delta'; \mathbf{a} : \tau'_+, \Gamma' \mid \\ \lambda \text{post}, (v, \vec{w}). \text{pre} (\lambda (v', \vec{w}'). \text{post} (\text{set } v \ v', \vec{w}')) (\text{get } v, \vec{w}) \end{array}$$

We have the following reflexive access judgment.

$$\begin{array}{c} \text{Acc-ID} \\ \mathbf{a} : \tau / \tau' \vdash \mathbf{a} : \tau / \tau' \mid \text{id}; \lambda v, w. w \end{array}$$

We can modify the head lifetime of the access judgment.

$$\begin{array}{c} \text{ACC-MONO-LFT} \\ \alpha \mid \mathbf{a} : \tau_+ / \tau'_+ \vdash \mathbf{a} : q. \tau^{(q)} / \tau'(q) \mid \text{get}; \text{set} \quad \beta \sqsubseteq \alpha \\ \hline \beta \mid \mathbf{a} : \tau_+ / \tau'_+ \vdash \mathbf{a} : q. \tau^{(q)} / \tau'(q) \mid \text{get}; \text{set} \end{array}$$

We can compose access judgments.

$$\begin{array}{c} \text{ACC-COMP} \\ \alpha \mid \mathbf{a} : \tau_{++} / \tau'_{++} \vdash e : q. \tau^{(q)} / \tau'(q) \mid \text{get}; \text{set} \\ \forall \mathbf{b}, q. \alpha \mid \mathbf{b} : \tau^{(q)} / \tau'(q) \vdash K[\mathbf{b}] : \tau^{(q)} / \tau'(q) \mid \text{get}'; \text{set}' \\ \hline \alpha \mid \mathbf{a} : \tau_{++} / \tau'_{++} \vdash K[e] : q. \tau^{(q)} / \tau'(q) \mid \text{get}' \circ \text{get}; \lambda v, w. \text{set } v \ (\text{set}' (\text{get}' v) w) \end{array}$$

$$\begin{array}{c} \text{ACC-COMP-VAR} \\ \alpha \mid \mathbf{a} : \tau_{++} / \tau'_{++} \vdash e : q. \tau^{(q)} / \tau'(q) \mid \text{get}; \text{set} \\ \forall \mathbf{b}, q. \alpha \mid \mathbf{b} : \tau^{(q)} / \tau'(q) \vdash K[\mathbf{b}] : q'. \tau^{(q')} / \tau'(q') \mid \text{get}'; \text{set}' \\ \hline \alpha \mid \mathbf{a} : \tau_{++} / \tau'_{++} \vdash K[e] : q'. \tau^{(q')} / \tau'(q') \mid \text{get}' \circ \text{get}; \lambda v, w. \text{set } v \ (\text{set}' (\text{get}' v) w) \end{array}$$

The getter and setter functions are composed in the standard way. In [ACC-COMP](#) the second access judgment keeps the fraction  $q$  given by the first access judgment, whereas in [ACC-COMP-VAR](#) the second access judgment introduces a new fraction  $q'$ .

### 6.1.4 Function Type

Now that the refined typing judgment is defined, we can introduce the *function type*. It has the following form.

$$\forall \vec{\alpha}. \text{fn}(\vec{\tau}_{\vec{\alpha}}) \rightarrow \tau'_{\vec{\alpha}}$$

The function type is universally quantified over the lifetime parameters  $\vec{\alpha}$ . We omit the assumptions on inclusion among these lifetimes; we virtually don't lose expressivity because we can use conjunction over lifetimes. The input types are  $\vec{\tau}_{\vec{\alpha}}$  and the output type is  $\tau'_{\vec{\alpha}}$ , which depends on the lifetime parameters  $\vec{\alpha}$ . The input and output types should have the size 1 regardless of  $\vec{\alpha}$  and the pure types of the input and output types should be constant over  $\vec{\alpha}$ .

The function type is modeled as follows.

$$[\forall \vec{\alpha}. \text{fn}(\vec{\tau}_{\vec{\alpha}}) \rightarrow \tau'_{\vec{\alpha}}] := (\times([\tau'_{\infty}]) \rightarrow \text{Prop}) \rightarrow \times([\tau_{\infty}]) \rightarrow \text{Prop}$$

$$|\forall \vec{\alpha}. \text{fn}(\vec{\tau}_{\vec{\alpha}}) \rightarrow \tau'_{\vec{\alpha}}| := 1 \quad \beta \sqsubseteq \forall \vec{\alpha}. \text{fn}(\vec{\tau}_{\vec{\alpha}}) \rightarrow \tau'_{\vec{\alpha}} := \text{True}$$

$$[\mathbf{a}] \blacktriangleleft \forall \vec{\alpha}. \text{fn}(\vec{\tau}_{\vec{\alpha}}) \rightarrow \tau'_{\vec{\alpha}} \{ \hat{v} \}_d = [\mathbf{a}] \triangleleft^{\alpha'} \forall \vec{\alpha}. \text{fn}(\vec{\tau}_{\vec{\alpha}}) \rightarrow \tau'_{\vec{\alpha}} :=$$

$$\exists f, \vec{b}, e. \text{let } \mathfrak{f} = \text{fn } f(\vec{b}) \{ e \} \text{ in } \mathbf{a} = \mathfrak{f} * \exists \text{pre s.t. } \hat{v} = \text{const pre.}$$

$$\forall \vec{\alpha}, \vec{b}, \beta. \triangleright \square (\beta \sqsubseteq (\vec{\tau}_{\vec{\alpha}}) \Rightarrow \beta \mid \vec{b}: \tau_{\vec{\alpha}} \vdash e[\mathfrak{f}/f, \vec{b}/\vec{b}]: \tau'_{\vec{\alpha}} \mid \mid \text{pre})$$

Here, we introduce the following shorthand.

$$\alpha \sqsubseteq (\vec{\tau}) := \forall i. \alpha \sqsubseteq \tau_i$$

The  $\pi$ -parametrized pure value for a function is a constant function on a *predicate transformer pre*. The cell value of a function is some function value  $\mathfrak{f}$ . For any lifetime parameters  $\vec{\alpha}$ , a function call  $\mathfrak{f}(\vec{b})$  on any inputs  $\vec{b}$  typed  $\vec{\tau}_{\vec{\alpha}}$  should output an object typed  $\tau'_{\vec{\alpha}}$  under some lifetime  $\beta$  and the predicate transformer *pre*, which is specified by the refined typing judgment. The lifetime  $\beta$  is ensured to be outlived by all the input types. The refined typing judgment in the ownership and sharing predicates is under the later modality for *contractiveness* over the input and output types  $\vec{\tau}_{\vec{\alpha}}, \tau'_{\vec{\alpha}}$ . We can strip off the later modality at the physical step of the function call (**HOARE-FNCALL**).

The function type satisfies the following contractiveness rule.

$$\frac{\text{CONTR-FN} \quad \forall i, \vec{\alpha}. \text{Nonex}(\sigma_{\vec{\alpha}, i}) \quad \forall \vec{\alpha}. \text{Nonex}(\sigma'_{\vec{\alpha}})}{\text{Contr}(\lambda \tau. \forall \vec{\alpha}. \text{fn}(\vec{\sigma}_{\vec{\alpha}} \vec{\tau}) \rightarrow \sigma'_{\vec{\alpha}} \tau)}$$

For the function type  $\forall \vec{\alpha}. \text{fn}(\vec{\tau}_{\vec{\alpha}}) \rightarrow \tau'_{\vec{\alpha}}$ , we also use the following shorthand. If the sequence of lifetime parameters  $\vec{\alpha}$  is empty, we can omit the part ' $\forall \vec{\alpha}.$ ' of the function type. If the subscript lifetime  $\beta_{\vec{\alpha}}$  is constantly  $\infty$ , we can omit it. Also, if the output type  $\tau'_{\vec{\alpha}}$  is  $\downarrow_1$ , we can omit the part ' $\rightarrow \tau'_{\vec{\alpha}}$ '.

## 6.2 Specifying and Verifying Operations on Basic Types

We introduce refined typing rules on the judgments defined in §6.1 for the semantic types defined in §5.2, §5.4 and §6.1.4.

Ths refined typing rules for the unique reference type are introduced later in §6.3.

## 6.2.1 Value Types

**Integer Type** We have the following refined typing rules on the integer type `int` (§5.2).

$$\begin{array}{c}
\text{RFN-VAL-INT} \\
\vdash n : \text{int} \mid \mid \lambda \text{post}, (). \text{post}(n) \\
\\
\text{RFN-INTOP} \\
\mathbf{a} : \text{int}, \mathbf{b} : \text{int} \vdash \mathbf{a} \text{ iop } \mathbf{b} : \text{int} \mid \mid \lambda \text{post}, (m, n). \text{post}(m \text{ iop } n) \\
\\
\text{RFN-INTREL} \\
\mathbf{a} : \text{int}, \mathbf{b} : \text{int} \vdash \mathbf{a} \text{ irel } \mathbf{b} : \text{bool} \mid \mid \lambda \text{post}, (m, n). \text{post}(m \text{ irel}_{\text{bool}} n) \\
\\
\text{RFN-NDINT} \qquad \qquad \qquad \text{COPY-INT} \\
\vdash \text{ndint} : \text{int} \mid \mid \lambda \text{post}, (). \forall n. \text{post}(n) \qquad \qquad \text{int copy}
\end{array}$$

`RFN-VAL-INT`, `RFN-INTOP`, `RFN-INTREL`, and `RFN-NDINT` specify an expression on a constant integer, an integer operation, an integer relation, and a non-deterministic integer (*iop* can be `+`, `-`, `*` and *irel* can be `≤`, `<`, `=`, for example). `COPY-INT` enables copying of an integer.

Here, we elaborate the proof of `RFN-INTOP`.

*Proof of RFN-INTOP.* It suffices to prove that the following is a tautology.

$$\{ \exists \hat{v}, \hat{v}'. \langle \pi. \text{post}(\hat{v} \pi \text{ iop } \hat{v}' \pi) \rangle * [\mathbf{a}] \blacktriangleleft \text{int} \{ \hat{v} \} * [\mathbf{b}] \blacktriangleleft \text{int} \{ \hat{v}' \} \} \\
\mathbf{a} \text{ iop } \mathbf{b} \quad \{ \mathbf{c}. \exists \hat{w}. \langle \pi. \text{post}(\hat{w} \pi) \rangle * [\mathbf{c}] \blacktriangleleft \text{int} \{ \hat{w} \} \}$$

By the model of the integer type, the cell values `a` and `b` should be some integers  $m$  and  $n$ . Also, the  $\pi$ -parametrized pure values  $\hat{v}$  and  $\hat{v}'$  should be `const m` and `const n`. So by `HOARE-INTOP`, we can execute `a iop b` and obtain a cell value `c`, which is equal to the integer  $l := m \text{ iop } n$ . We can set  $\hat{w}$  to `const l`. Now we have `[c] ◀ int {w}` by the model of the integer type. We also get  $\langle \pi. \text{post}(\hat{w} \pi) \rangle$ , because it is equivalent to  $\langle \pi. \text{post}(\hat{v} \pi \text{ iop } \hat{v}' \pi) \rangle$  by `PROPHOBS-WKN`.  $\square$

**Boolean Type** We have the following rules on the boolean type `bool` (§5.2).

$$\begin{array}{c}
\text{RFN-VAL-BOOL} \qquad \qquad \qquad \text{COPY-BOOL} \\
\vdash \text{bl} : \text{bool} \mid \mid \lambda \text{post}, (). \text{post}(\text{bl}) \qquad \qquad \text{bool copy} \\
\\
\text{RFN-IF} \\
\frac{\forall \text{bl}. \alpha \mid \Delta; \Gamma \vdash e_{\text{bl}} : \tau \mid \Delta'; \Gamma' \mid \text{pre}_{\text{bl}}}{\alpha \mid \Delta; \mathbf{a} : \text{bool}, \Gamma \vdash \text{if } \mathbf{a} \{ e_{\text{tt}} \} \text{ else } \{ e_{\text{ff}} \} : \tau \mid \Delta'; \Gamma' \mid \lambda \text{post}, (\text{bl}, \vec{v}). \text{pre}_{\text{bl}} \text{post}(\vec{v})}
\end{array}$$

`RFN-VAL-BOOL` specifies a boolean-value expression. `COPY-BOOL` enables copying of a boolean value. `RFN-IF` specifies conditional branching by a boolean value; note that the resulting predicate transformer is equivalent to the following.

$$\lambda \text{post}, (\text{bl}, \vec{v}). (\text{bl} = \text{tt} \Rightarrow \text{pre}_{\text{tt}} \text{post}(\vec{v})) \wedge (\text{bl} = \text{ff} \Rightarrow \text{pre}_{\text{ff}} \text{post}(\vec{v}))$$

We can *reinterpret* a boolean value as an integer value.

$$\begin{array}{c}
\text{SUBTY-BOOL-INT} \\
\text{bool} \sqsubseteq^{\text{acc}} \text{int} \mid \lambda \text{bl}. \text{if } \text{bl} \{ 1 \} \text{ else } \{ 0 \}
\end{array}$$

**Invalid-Data Type** We have the following rules on the invalid-data type  $\downarrow_n$  (§5.2).

$$\begin{array}{c} \text{COPY-INVALID} \\ \downarrow_n \text{ copy} \end{array} \qquad \begin{array}{c} \text{SUBTY-INVALID} \\ \tau \sqsubseteq^{\text{acc}} \downarrow_{|\tau|} \mid \lambda v. () \end{array}$$

**COPY-INT** enables copying of invalid data. **SUBTY-INVALID** invalidates an object of any type, turning the pure value into  $()$  regardless of the original pure value  $v$ .

## 6.2.2 Basic Pointer Types

**Box Pointer Type** Now we introduce refined typing rules for the box pointer type  $\text{box } \tau$  (§5.2).

By allocating a new memory block of the size  $n$ , we can create a box pointer to uninitialized, invalid data of the size  $n$ .

$$\begin{array}{c} \text{RFN-ALLOC} \\ \vdash \text{ alloc } n : \text{box } \downarrow_n \mid \lambda \text{post}, (). \text{post } () \end{array}$$

We can free the memory block of a box pointer.

$$\begin{array}{c} \text{RFN-FREE} \\ \mathbf{a} : \text{box } \tau \vdash \text{ free } \mathbf{a} \mid \lambda \text{post } (v). \text{post } () \end{array}$$

We can temporarily take out a *full plain reference* from a box pointer. We may change the target type when we put back a full reference, as long as the size of the target type does not change.

$$\begin{array}{c} \text{ACC-BOXPTR-PLNREF} \\ \frac{|\tau| = |\tau'|}{\mathbf{a} : \text{box } \tau / \text{box } \tau' \vdash \mathbf{a} : \&_1 \tau / \&_1 \tau' \mid \text{id}; \lambda v, w. w} \end{array}$$

It follows from the following lemma.

$$\begin{array}{c} \text{RAW-BOXPTR-PLNREF} \\ \frac{|\tau| = |\tau'|}{[\mathbf{a}] \blacktriangleleft \text{box } \tau \{ \hat{v} \} \Rightarrow [\mathbf{a}] \blacktriangleleft \&_1 \tau \{ \hat{v} \} * \forall \hat{v}'. ([\mathbf{a}] \blacktriangleleft \&_1 \tau' \{ \hat{v}' \} \multimap [\mathbf{a}] \blacktriangleleft \text{box } \tau' \{ \hat{v}' \})} \end{array}$$

We can also subdivide a plain reference to a box pointer.

$$\begin{array}{c} \text{ACC-DEREF-PLNREF-BOXPTR} \\ \frac{|\tau| = |\tau'|}{\mathbf{a} : \&_q \text{box } \tau / \&_q \text{box } \tau' \vdash * \mathbf{a} : \&_1 \tau / \&_1 \tau' \mid \text{id}; \lambda v, w. w} \end{array}$$

It follows from the following lemma.

$$\begin{array}{c} \text{RAW-SUBDIV-PLNREF-BOXPTR} \\ \frac{|\tau| = |\tau'|}{[\mathbf{I}] \blacktriangleleft \&_q \text{box } \tau \{ \hat{v} \} \Rightarrow_{\emptyset} \exists \mathbf{b}. \mathbf{I} \xrightarrow{q} \mathbf{b} * [\mathbf{b}] \blacktriangleleft \&_1 \tau \{ \hat{v} \} * \forall \hat{v}', d'. (\bigwedge \mathbf{1} * \mathbf{I} \xrightarrow{q} \mathbf{b} * [\mathbf{b}] \blacktriangleleft \&_1 \tau \{ \hat{v}' \} \multimap [\mathbf{I}] \blacktriangleleft \&_q \text{box } \tau \{ \hat{v} \})} \end{array}$$

We use the cumulative time receipt  $\bigwedge \mathbf{1}$  to swell the persistent time receipt of the new target object (**CUMU-TIME-SWELL-PERS-TIME**).

The box pointer type admits the following subtyping rule.

$$\begin{array}{c} \text{SUBTY-BOXPTR} \\ \frac{\triangleright (\tau \sqsubseteq^{\text{acc}} \tau' \mid f)}{\text{box } \tau \sqsubseteq^{\text{acc}} \text{box } \tau' \mid f} \end{array}$$

The subtyping assumption on the target types can be *under the later modality*. This is important for proving subtyping on *recursive types* with self reference under the box pointer type; see also **Recursive Type** of this section.

**Shared Reference Type** We introduce refined typing rules for the shared reference type  $\&_{\text{shr}}^\alpha \tau$  (§5.2).

We can create a shared reference by borrowing a full plain reference.

$$\begin{array}{c} \text{SWKN-SHRBOR-PLNREF} \\ \mathbf{a} : \&_1 \tau \vdash^\# \mathbf{a} : \&_{\text{shr}}^\alpha \tau, \mathbf{a} : \dagger^\alpha \&_1 \tau \mid \lambda \text{post}, (v). \text{post}(v, v) \end{array}$$

It follows from the following lemma.

$$\begin{array}{c} \text{RAW-SHRBOR-PLNREF} \\ [\mathbf{I}] \blacktriangleleft \&_1 \tau \{\hat{v}\}_{d+1} \rightleftarrows_{\mathcal{N}_{\text{fit}}}^d [\mathbf{I}] \blacktriangleleft \&_{\text{shr}}^\alpha \tau \{\hat{v}\}_{d+1} * [\mathbf{I}] \blacktriangleleft \dagger^\alpha \&_1 \tau \{\hat{v}\} \end{array}$$

The ownership predicate under borrow is defined in §6.1.1.

*Proof of RAW-SHRBOR-PLNREF.* The plain reference consists of a full points-to token  $\mathbf{I} \mapsto \bar{\mathbf{b}}$  and a target object under the later modality  $\triangleright (\bar{\mathbf{b}} \blacktriangleleft \tau \{\hat{v}\}_d)$ .

We fully borrow the points-to token (**FULLBOR-INTRO**). We split the full borrow  $\&_{\text{full}}^\alpha (\mathbf{I} \mapsto \bar{\mathbf{b}})$  into  $\ast_i \&_{\text{full}}^\alpha (\mathbf{I} + i \mapsto \bar{\mathbf{b}}[i])$  (**FULLBOR-SPLIT**) and turn them into fractured borrows  $\ast_i \&_{\text{frac}}^\alpha (\lambda q. \mathbf{I} + i \mapsto^q \bar{\mathbf{b}}[i])$  (**FULLBOR-FRACBOR**). We also get a view shift  $[\dagger\alpha] \rightleftarrows_{\mathcal{N}_{\text{fit}}} \mathbf{I} \mapsto^q \bar{\mathbf{b}}$ .

We also fully borrow the target object under later (**FULLBOR-INTRO**). We can turn the resulting full borrow  $\&_{\text{full}}^\alpha (\bar{\mathbf{b}} \blacktriangleleft \tau \{\hat{v}\}_d)$  into the sharing predicate on the target object  $\bar{\mathbf{b}} \blacktriangleleft^\alpha \tau \{\hat{v}\}_d$  in  $d$  logical steps by **SEM-TY-OWN-SHR** on  $\tau$ . We also get a view shift  $[\dagger\alpha] \rightleftarrows_{\mathcal{N}_{\text{fit}}} \triangleright (\bar{\mathbf{b}} \blacktriangleleft \tau \{\hat{v}\}_d)$ .

Combining the fractured borrows and the sharing predicate, we get a shared reference  $[\mathbf{I}] \blacktriangleleft \&_{\text{shr}}^\alpha \tau \{\hat{v}\}_{d+1}$ . Combining the two view shifts with a trivial prophecy equalizer  $\text{PE}(\hat{v}, \hat{v})$  (made by **PROPH-SEQZ-PROPHOBS** and **PROPHOBS-FACT**), we get a plain reference under borrow  $\mathbf{I} \blacktriangleleft \dagger^\alpha \&_1 \tau \{\hat{v}\}_{d+1}$ .  $\square$

We can create a shared reference also by borrowing a box pointer.

$$\begin{array}{c} \text{SWKN-SHRBOR-BOXPTR} \\ \mathbf{a} : \text{box } \tau \vdash^\# \mathbf{a} : \&_{\text{shr}}^\alpha \tau, \mathbf{a} : \dagger^\alpha \text{box } \tau \mid \lambda \text{post}, (v). \text{post}(v, v) \end{array}$$

It follows from **RAW-SHRBOR-PLNREF** and **RAW-BOXPTR-PLNREF**.

The shared reference type is copyable, unlike the box pointer type and the unique reference type.

$$\begin{array}{c} \text{COPY-SHRREF} \\ \&_{\text{shr}}^\alpha \tau \text{ copy} \end{array}$$

We can temporarily take out a *fractional plain reference* from an shared reference if the target type is copyable.

$$\begin{array}{c} \text{ACC-SHRREF-PLNREF} \\ \frac{\tau \text{ copy} \quad \tau' \text{ copy} \quad |\tau| = |\tau'|}{\alpha \mid \mathbf{a} : \&_{\text{shr}}^\alpha \tau / \&_{\text{shr}}^\alpha \tau' \vdash \mathbf{a} : q. \&_q \tau / \&_q \tau' \mid \text{id}; \lambda v, w. w} \end{array}$$

It follows from the following lemma.

$$\begin{array}{c} \text{RAW-SHRREF-PLNREF} \\ \frac{\tau \text{ copy} \quad \tau' \text{ copy} \quad |\tau| = |\tau'|}{[\mathbf{a}] \blacktriangleleft \&_{\text{shr}}^\alpha \tau \{\hat{v}\} * [\alpha]_{q'} \rightleftarrows_{\mathcal{N}_{\text{fit}}} \exists q. [\mathbf{a}] \blacktriangleleft \&_q \tau \{\hat{v}\} * \forall \hat{v}'. ([\mathbf{a}] \blacktriangleleft \&_q \tau' \{\hat{v}'\} \rightleftarrows_{\mathcal{N}_{\text{fit}}} [\mathbf{a}] \blacktriangleleft \&_{\text{shr}}^\alpha \tau' \{\hat{v}'\} * [\alpha]_{q'}}} \end{array}$$

*Proof of RAW-SHRREF-PLNREF.* The cell value  $\mathbf{a}$  should be some address  $\mathbf{l}$  and the depth  $d$  should be positive.

Let  $\bar{\mathbf{b}}$  the target cell values of the shared reference. Out of the fractured borrows of the shared reference, we can temporarily take out a fractional points-to token  $\mathbf{l} \mapsto^q \bar{\mathbf{b}}$  for some fraction  $q$  by [FRACBOR-ACCESS](#), using a fractional lifetime token  $[\alpha]_{q'}$ . By  $\tau$  copy, we can turn the target sharing predicate under later  $\triangleright (\bar{\mathbf{b}} \triangleleft^\alpha \tau \{\hat{v}\}_{d-1})$  into the target ownership predicate under later  $\triangleright (\bar{\mathbf{b}} \triangleleft \tau \{\hat{v}\}_{d-1})$ . Therefore we can construct a plain reference of some fraction  $q$ ,  $[\mathbf{a}] \triangleleft \&_q \tau \{\hat{v}\}_d$ .

Assume that we get back a plain reference  $[\mathbf{a}] \triangleleft \&_q \tau' \{\hat{v}'\}_d$ . We get back a fractional points-to token of the fraction  $q$ ,  $\mathbf{l} \mapsto^q \bar{\mathbf{b}}'$ , and thus get back the lifetime token  $[\alpha]_{q'}$ . By  $\tau'$  copy, the target ownership predicate under later  $\triangleright (\bar{\mathbf{b}}' \triangleleft \tau' \{\hat{v}'\}_{d-1})$  of the plain reference can be turned into the target sharing predicate. So we can reconstruct the shared reference  $[\mathbf{a}] \triangleleft \&_{\text{shr}}^\alpha \tau' \{\hat{v}'\}_d$ .  $\square$

Combining [ACC-SHRREF-PLNREF](#) and [RFN-DEREF-PLNREF-COPY](#), we can derive the following dereference rule on a shared reference to a shared reference. The lifetime of the resulting reference can be that of the *inner* shared reference  $\beta$ .

$$\begin{array}{c} \text{RFN-DEREF-SHRREF-SHRREF} \\ \alpha \mid \mathbf{a} : \&_{\text{shr}}^\alpha \&_{\text{shr}}^\beta \tau \vdash * \mathbf{a} : \&_{\text{shr}}^\beta \tau \mid \mid \text{id} \end{array}$$

We can also subdivide a plain reference to an shared reference.

$$\begin{array}{c} \text{ACC-DEREF-PLNREF-SHRREF} \\ \frac{|\tau| = |\tau'| \quad \tau \text{ copy} \quad \tau' \text{ copy}}{\alpha \mid \mathbf{a} : \&_q \&_{\text{shr}}^\alpha \tau / \&_q \&_{\text{shr}}^\alpha \tau' \vdash * \mathbf{a} : q' . \&_{q'} \tau' / \&_{q'} \tau' \mid \text{id}; \lambda v, w. w} \end{array}$$

It can be proved using [RAW-SHRREF-PLNREF](#). Note that we get a cumulative time receipt  $\boxtimes 1$  by the physical step of the load operation.

The shared reference type admits the following subtyping rule.

$$\begin{array}{c} \text{SUBTY-SHRREF} \\ \frac{\triangleright (\tau \sqsubseteq^{\text{shr}(\alpha)} \tau' \mid f) \quad \beta \sqsubseteq \alpha}{\&_{\text{shr}}^\alpha \tau \sqsubseteq^{\text{acc}} \&_{\text{shr}}^\beta \tau' \mid f} \end{array}$$

It modifies the target type of the shared reference type; the subtyping assumption on the target type is discussed on the sharing access mode and can be under the later modality like [SUBTY-BOXPTR](#).

We can turn a box pointer into a shared reference under the access mode  $\text{shr}(\alpha)$ .

$$\begin{array}{c} \text{SUBTY-SHR-BOXPTR-SHRREF} \\ \text{box } \tau \sqsubseteq^{\text{shr}(\alpha)} \&_{\text{shr}}^\alpha \tau \mid \text{id} \end{array}$$

Combining this with [RFN-DEREF-SHRREF-SHRREF](#), we can derive the following dereference rule.

$$\begin{array}{c} \text{RFN-DEREF-SHRREF-BOXPTR} \\ \frac{\beta \sqsubseteq \alpha}{\beta \mid \mathbf{a} : \&_{\text{shr}}^\alpha \text{box } \tau \vdash * \mathbf{a} : \&_{\text{shr}}^\alpha \tau \mid \mid \text{id}} \end{array}$$

**Plain Reference Type** We introduce refined typing rules for the plain reference type  $\&_q \tau$  (§5.2). We read and update objects through plain references, using access rules like [ACC-BOXPTR-PLNREF](#) and [ACC-SHRREF-PLNREF](#).

By dereferencing a plain reference to a size-1 type  $\tau$ , we can move out the target object. If the target type is copyable, the plain reference can retain the target object.

$$\begin{array}{c}
\text{RFN-DEREF-PLNREF-MOVE} \\
\mathbf{a} : \&_q \tau \vdash * \mathbf{a} : \tau \mid \mathbf{a} : \&_q \not\downarrow_{|\tau|} \mid \lambda post, (v). post (v, ()) \\
\\
\text{RFN-DEREF-PLNREF-COPY} \\
\tau \text{ copy} \\
\hline
\mathbf{a} : \&_q \tau \vdash * \mathbf{a} : \tau \mid \mathbf{a} : \&_q \tau \mid \lambda post, (v). post (v, v)
\end{array}$$

We can strip off the later modality on the target object of a plain reference, using the physical step of the load operation.

We can use the following rule when the target type  $\tau, \tau'$  has the size 1.

$$\begin{array}{c}
\text{RFN-STORE-PLNREF-VAL} \\
|\tau| = |\tau'| = 1 \\
\hline
\mathbf{a} : \&_1 \tau, \mathbf{b} : \tau' \vdash \mathbf{a} \leftarrow \mathbf{b} \mid \mathbf{a} : \&_1 \tau' \mid \lambda post, (v, w). post (w)
\end{array}$$

When the new target object  $\mathbf{b}$  has the depth  $d$ , the updated plain reference  $\mathbf{a}$  should have the depth  $d + 1$  and thus have the persistent time receipt  $\bar{\chi}(d + 1)$ . We make  $\bar{\chi}(d + 1)$  out of the persistent time receipt  $\bar{\chi}d$  of  $\mathbf{b}$  and the cumulative time receipt  $\bar{\chi}1$  obtained by the physical step of the store operation, using [CUMU<sub>TIME</sub>-SWELL-PERS<sub>TIME</sub>](#).

We can also move or copy an object from a plain reference to a full plain reference.

$$\begin{array}{c}
\text{RFN-STORE-PLNREF-PLNREF-MOVE} \\
|\tau'| = |\tau| \\
\hline
\mathbf{a} : \&_1 \tau', \mathbf{b} : \&_q \tau \mid \mathbf{a} \leftarrow_{|\tau|}^* \mathbf{b} \mid \mathbf{a} : \&_1 \tau, \mathbf{b} : \&_q \not\downarrow_{|\tau|} \mid \lambda post, (w, v). post (v, ()) \\
\\
\text{RFN-STORE-PLNREF-PLNREF-COPY} \\
|\tau'| = |\tau| \quad \tau \text{ copy} \\
\hline
\mathbf{a} : \&_1 \tau', \mathbf{b} : \&_q \tau \mid \mathbf{a} \leftarrow_{|\tau|}^* \mathbf{b} \mid \mathbf{a} : \&_1 \tau, \mathbf{b} : \&_q \tau \mid \lambda post, (w, v). post (v, v)
\end{array}$$

The plain reference type satisfies the following structural subtyping rule.

$$\begin{array}{c}
\text{SUBTY-PLNREF} \\
\tau \sqsubseteq^{\text{own}} \tau' \mid f \\
\hline
\&_q \tau \sqsubseteq^{\text{own}} \&_q \tau' \mid f
\end{array}$$

### 6.2.3 Constructive Types

**Pair Type** We introduce refined typing rules for the pair type  $\tau \times \tau'$  (§5.2).

When we have a plain reference to a pair, we can take out a plain reference to each component of the pair.

$$\begin{array}{c}
\text{ACC-PLNREF-PAIR-L} \\
|\tau_0| = |\tau'_0| \\
\hline
\mathbf{a} : \&_q(\tau_0 \times \tau_1) / \&_q(\tau'_0 \times \tau_1) \vdash \mathbf{a} : \&_q \tau_0 / \&_q \tau'_0 \mid \lambda v. v.0; \lambda v, w. (w, v.1) \\
\\
\text{ACC-PLNREF-PAIR-R} \\
\mathbf{a} : \&_q(\tau_0 \times \tau_1) / \&_q(\tau_0 \times \tau'_1) \vdash \mathbf{a}.|\tau_0| : \&_q \tau_1 / \&_q \tau'_1 \mid \lambda v. v.1; \lambda v, w. (v.0, w)
\end{array}$$

We can also split a plain reference to a pair into plain references to each element of the pair.

$$\begin{array}{c}
\text{RFN-SPLIT-PLNREF-PAIR-R} \\
|\tau_0| = |\tau'_0| \\
\forall \mathbf{b}. \alpha \mid \Delta; \mathbf{a} : \&_q \tau_0, \mathbf{b} : \&_q \tau_1, \Gamma \vdash K[\mathbf{b}] : \tau \mid \Delta'; \mathbf{a} : \&_q \tau'_0, \mathbf{b} : \&_q \tau'_1, \Gamma' \mid pre \\
\hline
\alpha \mid \Delta; \mathbf{a} : \&_q(\tau_0 \times \tau_1), \Gamma \vdash K[\mathbf{a}.|\tau_0|] : \tau \mid \Delta'; \mathbf{a} : \&_q(\tau'_0 \times \tau'_1), \Gamma' \mid \\
\lambda post, ((v_0, v_1), \vec{w}). pre (\lambda (v'_0, v'_1, \vec{w}'). post ((v'_0, v'_1), \vec{w}')) (v_0, v_1, \vec{w})
\end{array}$$

These rules follow from the following simple lemma.

RAW-SPLIT-MERGE-PLNREF-PAIR

$$[I] \blacktriangleleft \&_q(\tau_0 \times \tau_1) \{ \lambda \pi. (\hat{v}_0 \pi, \hat{v}_1 \pi) \}_d \Leftrightarrow [I] \blacktriangleleft \&_q \tau_0 \{ \hat{v}_0 \}_d * [I + |\tau_0|] \blacktriangleleft \&_q \tau_1 \{ \hat{v}_1 \}_d$$

From a shared reference to a pair, we can take shared references to each component of the pair.

SUBTY-SUBDIV-SHREF-PAIR-L

$$\&_{\text{shr}}^\alpha(\tau \times \tau') \sqsubseteq \&_{\text{shr}}^\alpha \tau \mid \lambda v. v.0$$

RFN-SUBDIV-SHREF-PAIR-R

$$\mathbf{a} : \&_{\text{shr}}^\alpha(\tau \times \tau') \vdash \mathbf{a}.|\tau| : \&_{\text{shr}}^\alpha \tau' \mid \mid \lambda \text{post}, (v). \text{post}(v.1)$$

We have the following structural rules on copyability and subtyping on the pair type.

$$\begin{array}{c} \text{COPY-PAIR} \\ \frac{\tau \text{ copy} \quad \tau' \text{ copy}}{\tau \times \tau' \text{ copy}} \end{array} \qquad \begin{array}{c} \text{SUBTY-PAIR} \\ \frac{\tau_0 \sqsubseteq^{\text{acc}} \tau'_0 \mid f \quad \tau_1 \sqsubseteq^{\text{acc}} \tau'_1 \mid g}{\tau_0 \times \tau_1 \sqsubseteq^{\text{acc}} \tau'_0 \times \tau'_1 \mid f \times g} \end{array}$$

Here, the function  $f \times g$  is defined as follows.

$$f \times g := \lambda(v, w). (f v, g w)$$

We can also use the following subtyping rules for *reinterpreting* a nested pair in terms of associativity.

SUBTY-ASSOC-PAIR-L

$$(\tau \times \tau') \times \tau'' \simeq \tau \times (\tau' \times \tau'') \mid \lambda((v, v'), v''). (v, (v', v''))$$

SUBTY-ASSOC-PAIR-R

$$\tau \times (\tau' \times \tau'') \times \tau'' \simeq (\tau \times \tau') \mid \lambda(v, (v', v'')). ((v, v'), v'')$$

Using these rules, we can simplify the address shift operation for data access, i.e., we can write  $l$  instead of  $e.m.n$  where  $l = m + n$ .

We have the following subtyping rule for splitting an invalid object.

SUBTY-SPLIT-INVALID

$$\downarrow_{m+n} \sqsubseteq^{\text{acc}} \downarrow_m \times \downarrow_n \mid \text{const}(((), ()))$$

**Variant Type** We introduce refined typing rules for the variant type  $\tau + \tau'$  (§5.2).

We can convert between a plain reference to a variant and a plain reference to a triple of the tag, body and invalid object.

WKN-PLNREF-VRNT-TRIPLE

$$\mathbf{a} : \&_q(\tau_0 + \tau_1) \vdash \mathbf{a} : \&_q \left( (\text{int} \times \tau_i) \times \downarrow_{|\tau_0 + \tau_1| - 1 - |\tau_i|} \right) \mid \lambda \text{post}, (v). \exists w \text{ s.t. } v = \text{inj}_i w. \text{post}(((i, w), ()))$$

WKN-PLNREF-TRIPLE-VRNT

$$\mathbf{a} : \&_q \left( (\text{int} \times \tau_i) \times \downarrow_{|\tau_0 + \tau_1| - 1 - |\tau_i|} \right) \vdash \mathbf{a} : \&_q(\tau_0 + \tau_1) \mid \lambda \text{post}, ((n, v), ()). n = i \wedge \text{post}(\text{inj}_i v)$$

In the rule [WKN-PLNREF-VRNT-TRIPLE](#), the predicate transformer requires the precondition  $\exists w. v = \text{inj}_i w$  (i.e., the tag of  $v$  is  $i$ ) for any postcondition  $\text{post}$ , which ensures that the cell value at the address  $\mathbf{a}$  is  $i$  (by the definition of the variant type and [PROPHOBS-SAT](#)). In the rule [WKN-PLNREF-TRIPLE-VRNT](#), we have the precondition that the tag number  $n$  has the value  $i$ . These rules allow us to read and update a variant object using the rules for pairs.

Also, we also have the following rule for subdividing a shared reference to a variant into a shared reference to the body object.

RFN-SUBDIV-SHRREF-VRNT

$$\mathbf{a} : \&_{\text{shr}}^\alpha (\tau_0 + \tau_1) \vdash \mathbf{a}.1 : \&_{\text{shr}}^\alpha \tau_i \mid \mid \lambda \text{post}, (v). \exists w \text{ s.t. } \text{inj}_i w = v. \text{post} (w)$$

Like [WKN-PLNREF-VRNT-TRIPLE](#), the predicate transformer ensures the precondition  $\exists w. \text{inj}_i w = v$  for any postcondition  $\text{post}$ .

We have the following rule for conditional branching on the tag of a variant object.

RFN-CASE-VRNT-ACCESS

$$\frac{\begin{array}{c} \Pi(\beta, \Delta) \mid \mathbf{a} : \tau_+ / \tau_+ \vdash e : q. \&_{fq}(\tau_0 + \tau_1) / \&_{fq}(\tau_0 + \tau_1) \mid \text{get}; \text{set} \\ \forall i. \beta \mid \Delta; \mathbf{a} : \tau_+, \Gamma \vdash e_i : \tau' \mid \Delta'; \Gamma' \mid \text{pre}_i \end{array}}{\beta \mid \Delta; \mathbf{a} : \tau_+, \Gamma \vdash \text{case } *e \text{ of } \{ 0 \rightarrow e_0, 1 \rightarrow e_1 \} : \tau' \mid \Delta'; \Gamma' \mid \lambda \text{post}, (v, \vec{w}). \forall i. (\exists v'. \text{inj}_i v' = \text{get } v) \Rightarrow \text{pre}_i \text{post} (\text{set } v (\text{get } v), \vec{w})}$$

The access judgment says that, by performing  $e$ , we can access some plain reference to a variant object inside  $\mathbf{a} : \tau_+$  with the getter function  $\text{get}$ . We use a function on the fraction variable  $f$ , which is usually set to  $\text{id}$  or  $\text{const } 1$ . The value  $\text{set } v (\text{get } v)$  is usually equal to  $v$ . If we go to the branch of  $i$  (0 or 1), we get the information that the variant sub-object  $v$  has the tag  $i$  as a postcondition.

We have the following structural rules on copyability and subtyping on the variant type.

COPY-VRNT

$$\frac{\tau \text{ copy} \quad \tau' \text{ copy}}{\tau + \tau' \text{ copy}}$$

SUBTY-VRNT

$$\frac{\tau_0 \sqsubseteq^{\text{acc}} \tau'_0 \mid f \quad \tau_1 \sqsubseteq^{\text{acc}} \tau'_1 \mid g}{\tau_0 + \tau_1 \sqsubseteq^{\text{acc}} \tau'_0 + \tau'_1 \mid f + g}$$

Here, the function  $f + g$  is defined as follows.

$$f + g := \lambda v. \text{case } v \text{ of } \{ \text{inj}_0 w_0 \rightarrow f w_0, \text{inj}_1 w_1 \rightarrow g w_1 \}$$

**Injection Type** We introduce refined typing rules for the injection type  $\text{in}(\tau, f)$  (§5.4).

We can subdivide a plain reference to an injection type by the following rule.

WKN-SUBDIV-PLNREF-INJTY

$$\mathbf{a} : \&_q \text{in}(\tau, f) \vdash \mathbf{a} : \&_q \tau \mid \lambda \text{post}, (v). \forall w \text{ s.t. } f w = v. \text{post} (w)$$

Also, we can subdivide a shared reference to an injection type by the following rule.

WKN-SUBDIV-SHRREF-INJTY

$$\mathbf{a} : \&_{\text{shr}}^\alpha \text{in}(\tau, f) \vdash \mathbf{a} : \&_{\text{shr}}^\alpha \tau \mid \lambda \text{post}, (v). \forall w \text{ s.t. } f w = v. \text{post} (w)$$

The injection type satisfies the following structural copyability rule.

COPY-INJTY

$$\frac{\tau \text{ copy}}{\text{in}(\tau, f) \text{ copy}}$$

We have the following subtyping rules between  $\tau$  and the injection type  $\text{in}(\tau, f)$ .

$$\text{SUBTY-TY-INJTY} \quad \tau \sqsubseteq^{\text{acc}} \text{in}(\tau, f) \mid f \quad \text{SUBTY-INJTY-TY} \quad \frac{g \circ f = \text{id}}{\text{in}(\tau, f) \sqsubseteq^{\text{acc}} \tau \mid g}$$

Combining the two rules, we can derive the following structural subtyping rule.

$$\text{SUBTY-INJTY-INJTY} \quad \frac{\tau \sqsubseteq^{\text{acc}} \tau' \mid h \quad g \circ f = \text{id}}{\text{in}(\tau, f) \sqsubseteq^{\text{acc}} \text{in}(\tau', f') \mid f' \circ h \circ g}$$

**Recursive Type** We introduce refined typing rules for the recursive type  $\text{rec}(\sigma, f)$  (§5.4).

We can prove subtyping on recursive types using *Löb induction* (LÖB).

*Example 6.1* (Subtyping on the List Type). For example, for the list type  $\text{list } \tau$  defined in [Example 5.1](#), we can *derive* the following subtyping rule.

$$\frac{\tau \sqsubseteq^{acc} \tau' \mid f}{\text{list } \tau \sqsubseteq^{acc} \text{list } \tau' \mid \text{map } f}$$

We apply the value transformer  $f$  to each element of the list by the function  $\text{map } f$ .

*Proof.* Assume  $\tau \sqsubseteq^{acc} \tau' \mid f$ . By Löb induction (LÖB), we can also assume  $\triangleright (\text{list } \tau \sqsubseteq^{acc} \text{list } \tau' \mid \text{map } f)$ .

Unfolding  $\text{list } \tau$  by [TYEQ-LIST](#), by [SUBTY-INJTY-INJTY](#), [SUBTY-VRNT](#), [SUBTY-PAIR](#) and [SUBTY-BOXPTR](#), we have the following subtyping judgment.

$$\text{list } \tau \sqsubseteq^{acc} \text{list } \tau' \mid \text{in}_{\text{List}} \circ (\text{id} + f \times \text{map } f) \circ \text{out}_{\text{List}}$$

The value transformer of this judgment is equal to  $\text{map } f$ . Therefore, we obtain the expected judgment.  $\square$

## 6.2.4 Function Type

We introduce refined typing rules for the function type (§6.1.4). We also describe how we can specify and verify the while loop, based on the technique for recursive functions.

We introduce the following shorthand for specifying a constant value  $\mathbf{a}$ .

$$\mathbf{a} : \tau \mid v \quad := \quad \vdash \quad \mathbf{a} : \tau \mid \lambda \text{post}, (). \text{post } (v)$$

**Function Type** We can specify a non-recursive function by the following rule.

$$\frac{\text{RFN-FN-NONREC} \quad \forall \vec{\alpha}, \vec{\mathbf{a}}, \beta. \quad \square (\beta \sqsubseteq (\vec{\tau}_{\vec{\alpha}}) \Rightarrow \beta \mid \vec{\mathbf{a}} : \vec{\tau}_{\vec{\alpha}} \vdash e[\vec{\mathbf{a}}/\vec{\mathbf{a}}] : \tau'_{\vec{\alpha}} \mid \text{pre})}{\text{fn } (\vec{\mathbf{a}}) \{e\} : \forall \vec{\alpha}. \text{fn } (\vec{\tau}_{\vec{\alpha}}) \rightarrow \tau'_{\vec{\alpha}} \mid \text{pre}}$$

If we want to set the pure value of the function to the predicate transformer  $\text{pre}$ , we just need to specify the body expression  $e$  (under arbitrary substitution) with the predicate transformer  $\text{pre}$ , using the refined typing judgment.

Extending [RFN-FN-NONREC](#), we get the following rule for specifying a *recursive* function.

$$\frac{\text{RFN-FN-REC} \quad \tau_+ = \forall \vec{\alpha}. \text{fn } (\vec{\tau}_{\vec{\alpha}}) \rightarrow \tau'_{\vec{\alpha}} \quad \forall \vec{\alpha}, \mathbf{b}, \vec{\mathbf{a}}, \beta. \quad \square (\beta \sqsubseteq (\vec{\tau}_{\vec{\alpha}}) \Rightarrow \beta \mid \mathbf{b} : \tau_+, \vec{\mathbf{a}} : \vec{\tau}_{\vec{\alpha}} \vdash e[\mathbf{b}/f, \vec{\mathbf{a}}/\vec{\mathbf{a}}] : \tau'_{\vec{\alpha}} \mid \text{pre}') \quad \text{pre} \Rightarrow \lambda \text{post}, (\vec{v}). \text{pre}' \text{post } (\text{pre}, \vec{v})}{\text{fn } f(\vec{\mathbf{a}}) \{e\} : \tau_+ \mid \text{pre}}$$

We want to prove that the function  $\mathbf{f}$  satisfies the predicate transformer  $\text{pre}$ . First, we specify the body expression  $e$  under arbitrary substitution with a predicate transformer  $\text{pre}'$ . When we verify the substituted expression, we can use the assumption that  $\mathbf{f}$  satisfies  $\text{pre}$ , because we can use that assumption under later by Löb induction (LÖB) and strip off the later on the physical step of the function call ([HOARE-FNCALL](#)). Therefore, it suffices to prove that  $\text{pre}$  is stronger than  $\text{pre}'$  with the function pure value set to  $\text{pre}$ , i.e.,  $\lambda \text{post}, (\vec{v}). \text{pre}' \text{post } (\text{pre}, \vec{v})$ . If  $\text{pre}'$  satisfies monotonicity condition on the function

pure value, this is equivalent to the condition that  $pre$  is stronger than the *greatest fixed point* of  $pre'$  with a modified argument order  $\lambda pre_*. post, (\vec{v}). pre' post (pre_*, \vec{v})$  (whose existence is ensured by Tarski's fixed point theorem).

We can specialize the lifetime parameters of the function type by the following subtyping rule.

$$\text{SUBTY-FN-LFT} \\ \forall \vec{\alpha}. \text{fn}(\vec{\tau}_{\vec{\alpha}}) \rightarrow \tau'_{\vec{\alpha}} \sqsubseteq \forall \vec{\beta}. \text{fn}(\vec{\tau}_{\vec{\alpha}'_{\vec{\beta}}}) \rightarrow \tau'_{\vec{\alpha}'_{\vec{\beta}}} \mid \text{id}$$

Here, each of the old lifetime parameters  $\alpha_i$  is specialized into some lifetime  $\alpha'_{\vec{\beta},i}$  that depends on the new lifetime parameters  $\vec{\beta}$ . In particular, we can use an empty sequence for  $\vec{\beta}$ .

A function call is specified by the following rule (we can specialize the lifetime parameters of the function beforehand by [SUBTY-FN-LFT](#)).

$$\text{RFN-FN-CALL} \\ \frac{\alpha \sqsubseteq (\vec{\tau})}{\alpha \mid \mathbf{b}: \text{fn}(\vec{\tau}) \rightarrow \tau', \vec{\mathbf{a}}: \vec{\tau} \vdash \mathbf{b}(\vec{\mathbf{a}}): \tau' \mid \mid \lambda post, (pre, \vec{v}). pre post (\vec{v})}$$

We can use any lifetime arguments  $\vec{\beta}$  for the function call. The precondition required for the function call is the pure value of the function  $pre$  (predicate transformer) applied to the postcondition on the output  $post$  and the arguments  $\vec{v}$ . Note that we can derive the following rule from [RFN-FN-CALL](#) and [RFN-EVCTX](#).

$$\text{RFN-FN-CALL-CONST} \\ \frac{\mathbf{f}: \text{fn}(\vec{\tau}) \rightarrow \tau' \quad \alpha \sqsubseteq (\vec{\tau})}{\alpha \mid \vec{\mathbf{a}}: \vec{\tau} \vdash \mathbf{f}(\vec{\mathbf{a}}): \tau' \mid \mid \lambda post, (pre, \vec{v}). pre post (\vec{v})}$$

A function is copyable.

$$\text{COPY-FN} \\ \forall \vec{\alpha}. \text{fn}(\vec{\tau}_{\vec{\alpha}}) \rightarrow \tau'_{\vec{\alpha}} \text{ copy}$$

The function type satisfies the following structural subtyping rule.

$$\text{SUBTY-FN-TY} \\ \frac{\forall i, \vec{\alpha}. \triangleright (\tau'_{\vec{\alpha},i} \sqsubseteq^{\text{own}} \tau_{\vec{\alpha},i} \mid f_i) \quad \forall \vec{\alpha}. \triangleright (\tau_{\vec{\alpha}}^+ \sqsubseteq^{\text{own}} \tau_{\vec{\alpha}}^{+'} \mid g)}{\forall \vec{\alpha}. \text{fn}(\vec{\tau}_{\vec{\alpha}}) \rightarrow \tau_{\vec{\alpha}}^+ \sqsubseteq \forall \vec{\alpha}. \text{fn}(\vec{\tau}_{\vec{\alpha}}^{+'}) \rightarrow \tau_{\vec{\alpha}}^{+'} \mid \mid \lambda pre, post, (\vec{v}). pre (post \circ g) (f \vec{v})}$$

We can modify the input and output objects of the function by subtyping. Like [SUBTY-BOXPTR](#), the assumption subtyping judgments on the input and output types are under the later modality.

**While Loop** The while loop  $\text{while } e \{ e' \}$  is defined as follows.

$$\text{while } e \{ e' \} := (\text{fn } \text{while}() \{ \text{if } e \{ e'; \text{while}() \} \}) ()$$

We have the following refined typing rule on the while loop.

$$\text{RFN-WHILE} \\ \frac{\alpha \mid \Delta; \Gamma \vdash e: \text{bool} \mid \Delta; \Gamma \mid pre \quad \alpha \mid \Delta; \Gamma \vdash e' \mid \Delta; \Gamma \mid pre'}{pre_+ \Rightarrow \lambda post, (\vec{v}). pre (\lambda (bl, \vec{v}'). \text{if } bl \{ pre' (pre_+ post) (\vec{v}') \} \text{ else } \{ pre_+ post (\vec{v}') \}) (\vec{v})} \\ \alpha \mid \Delta; \Gamma \vdash \text{while } e \{ e' \} \mid \Delta; \Gamma \mid pre_+$$

Just like [RFN-FN-REC](#), we can use the predicate transformer  $pre_+$  if  $pre_+$  is stronger than the predicate transformer of the while loop with the predicate transformer for the recursive call set to  $pre_+$ .

We can freely add new control flows to this refined type system like this.

### 6.2.5 Vector Type

We introduce refined typing rules for the vector type  $\text{vec } \tau$  (§5.2).

We introduce the following function  $\text{new}_{\text{vec}}$  for newly allocating an empty vector of some capacity.

$$\text{new}_{\text{vec}} := \text{fn}(a) \{ \text{let } b = \text{alloc } 3 \text{ in } b \leftarrow \text{alloc}(a \times |\tau|); b.1 \leftarrow a; b.2 \leftarrow 0; b \}$$

It satisfies the following specification.

$$\begin{array}{c} \text{RFN-NEW-VEC} \\ \text{new}_{\text{vec}} : \text{fn}(\text{int}) \rightarrow \text{box } \text{vec } \tau \mid \lambda \text{post}, (n). \text{post } \text{nil} \end{array}$$

We introduce the following function for taking out of the reference to a vector  $a$  a reference to the  $i$ -th element of the vector.

$$\text{idx}_{\text{vec}}^\tau := \text{fn}(a, i) ( (*a).(i \times |\tau|) )$$

Note that this function does not have *dynamic bounds checking* on the index input  $i$ .

The function  $\text{idx}_{\text{vec}}^\tau$  satisfies the following specification for shared references.

$$\begin{array}{c} \text{RFN-IDX-VEC-SHRREF} \\ \text{idx}_{\text{vec}}^\tau : \forall \alpha. \text{fn}(\&_{\text{shr}}^\alpha \text{vec } \tau, \text{int}) \rightarrow \&_{\text{shr}}^\alpha \tau \mid \lambda \text{post}, (v, i). 0 \leq i < \text{len } v \wedge \text{post } (v[i]) \end{array}$$

We can safely perform  $\text{idx}_{\text{vec}}^\tau$  *without dynamic bounds checking*, because the precondition ensures that the index  $i$  is in the suitable range (the part  $0 \leq i < \text{len } v$ ).

Update operations on the vector type are performed through *unique references* and discussed later in §6.3.6.

We have the following structural subtyping rule on the vector type.

$$\begin{array}{c} \text{SUBTY-VEC} \\ \triangleright (\tau \sqsubseteq^{\text{acc}} \tau' \mid f) \\ \hline \text{vec } \tau \sqsubseteq^{\text{acc}} \text{vec } \tau' \mid \text{map } f \end{array}$$

Like [SUBTY-BOXPTR](#), the subtyping assumption on the target types can be under the later modality. We apply the value transformer  $f$  to each element of the vector by the function  $\text{map } f$ . The function  $\text{map} : (T \rightarrow U) \rightarrow \text{List } T \rightarrow \text{List } U$  is defined as follows.

$$\text{map } f \text{ nil} := \text{nil} \qquad \text{map } f (v :: w) := f v :: \text{map } f w$$

## 6.3 Specifying and Verifying Operations on Unique References

Now we introduce refined typing rules on the unique reference type  $\&_{\text{uniq}}^\alpha \tau$  (§5.3). We present detailed proofs on some of them, which contain interesting interactions of full borrow and prophecy.

### 6.3.1 Unique Borrow and Reborrow

We can borrow a full plain reference to get a unique reference by the following rule.

$$\begin{array}{c} \text{WKN-UNQBOR-PLNREF} \\ \mathbf{a} : \&_1 \tau \vdash \mathbf{a} : \&_{\text{uniq}}^\alpha \tau, \mathbf{a} : \dagger^\alpha \&_1 \tau \mid \lambda \text{post}, (v). \forall v_o. \text{post } ((v, v_o), v_o) \end{array}$$

We take a prophecy value  $v_o$  for the unique borrow and set the values of the unique reference and the borrowed plain reference to  $(v, v_o)$  and  $v_o$ , respectively. When we reclaim the plain reference after the lifetime  $\alpha$  ends ([SWKN-END-LOCALLEFT-RECLAIM](#)), it is ensured that  $v_o$  has been resolved to the new value of the plain reference.

[WKN-UNQBOR-PLNREF](#) follows from the following lemma.

$$\begin{array}{c} \text{RAW-UNQBOR-PLNREF} \\ [\mathbb{I}] \blacktriangleleft \&_1 \tau \{ \hat{v} \} \Rightarrow_{\mathcal{N}_{\text{fit}} + \mathcal{N}_{\text{proph}}} \exists x. \\ [\mathbb{I}] \blacktriangleleft \&_{\text{uniq}}^\alpha \tau \{ \lambda \pi. (\hat{v} \pi, \pi x) \} * [\mathbb{I}] \blacktriangleleft^{\dagger\alpha} \&_1 \tau \{ \lambda \pi. \pi x \} \end{array}$$

*Proof of [RAW-UNQBOR-PLNREF](#).* The plain reference is decomposed into (i) a full points-to token  $\mathbb{I} \mapsto \bar{\mathbf{b}}$ , (ii) the target object under later  $\triangleright \bar{\mathbf{b}} \blacktriangleleft \tau \{ \hat{v} \}_d$ , and (iii) a persistent time receipt  $\bar{\chi}(d+1)$ , for some target cell values  $\bar{\mathbf{b}}$  and depth  $d$ .

By [VALOBS-PROPHCTRL-INTRO](#), for a fresh prophecy variable  $x$  satisfying  $x.\text{type} = \lfloor \tau \rfloor$ , we get a value observer  $\text{VO}_x(\hat{v}, d)$  and a prophecy control  $\text{PC}(\hat{v}, d)$ . We name the following proposition  $P$ , which is a separating conjunction of a full points-to token, a target object, a persistent time receipt and a prophecy control that is existentially quantified over the list of cell values, target value and depth.

$$P := \exists \bar{\mathbf{b}}', \hat{v}', d'. \mathbb{I} \mapsto \bar{\mathbf{b}}' * \bar{\mathbf{b}}' \blacktriangleleft \tau \{ \hat{v}' \}_{d'} * \bar{\chi}(d'+1) * \text{PC}(x, \hat{v}', d')$$

We make  $\triangleright P$  consuming the full points-to token, the target object *under later*, and the prophecy control, which is then turned by [FULLBOR-INTRO](#) into a full borrow  $\&_{\text{full}}^\alpha P$  and a view shift  $[\dagger\alpha] \Rightarrow_{\mathcal{N}_{\text{fit}}} * \triangleright P$ . We turn the separating conjunction of the full borrow and the value observer into a unique reference  $[\mathbb{I}] \blacktriangleleft \&_{\text{uniq}}^\alpha \tau \{ \lambda \pi. (\hat{v} \pi, \pi x) \}$ .

Now let us construct the full plain reference under borrow  $[\mathbb{I}] \blacktriangleleft^{\dagger\alpha} \&_1 \tau \{ \lambda \pi. \pi x \}$  (defined in [§6.1.1](#)). Now assume we have  $[\dagger\alpha]$ . Consuming the view shift  $[\dagger\alpha] \Rightarrow_{\mathcal{N}_{\text{fit}}} * \triangleright P$ , we get  $\triangleright P$ . By consuming  $\triangleright P$ , we get a new full reference  $[\mathbb{I}] \blacktriangleleft \&_1 \tau \{ \hat{v}' \}$  and a prophecy control under later  $\triangleright \text{PC}(x, \hat{v}', d')$ . The latter can be transformed into a prophecy equalizer under later  $\triangleright \text{PE}(\lambda \pi. \pi x, \hat{v}')$  by [PROPHCTRL-PROPHEQZ](#). Therefore, we finally get  $[\mathbb{I}] \blacktriangleleft^{\dagger\alpha} \&_1 \tau \{ \lambda \pi. \pi x \}$ .  $\square$

Note that we can borrow only some component of a (possibly nested) pair by combining [WKN-UNQBOR-PLNREF](#) and [RFN-SPLIT-UNQREF-PAIR-R](#).

We can uniquely borrow a box pointer in an indirect way by combining [ACC-BOXPTR-PLNREF](#), [RFN-ACC-SAME](#) and [WKN-UNQBOR-PLNREF](#). Still, we can also have the following rule for uniquely borrowing a box pointer in a direct way.

$$\begin{array}{c} \text{WKN-UNQBOR-BOXPTR} \\ \mathbf{a} : \text{box } \tau \vdash \mathbf{a} : \&_{\text{uniq}}^\alpha \tau, \mathbf{a} : \dagger^\alpha \text{ box } \tau \mid \lambda \text{post}, (v). \forall v_o. \text{post}((v, v_o), v_o) \end{array}$$

This rule follows from [RAW-UNQBOR-PLNREF](#) and [RAW-BOXPTR-PLNREF](#).

We can also create a unique reference by *reborrowing* from an existing *unique reference*.

$$\begin{array}{c} \text{WKN-UNQBOR-UNQREF} \\ \mathbf{a} : \&_{\text{uniq}}^\beta \tau \vdash \mathbf{a} : \&_{\text{uniq}}^\alpha \tau, \mathbf{a} : \dagger^\alpha \&_{\text{uniq}}^\beta \tau \mid \lambda \text{post}, ((v, v'_o)). \forall v'_o. \text{post}((v, v_o), (v_o, v'_o)) \end{array}$$

Like [WKN-UNQBOR-BOXPTR](#), we take a prophecy value  $v_o$  for this borrow. The unique reference under borrow  $\mathbb{I} : \dagger^\alpha \&_{\text{uniq}}^\beta \tau$  takes the value  $(v_o, v'_o)$ , where the current-value part is set to the new prophecy value  $v_o$  and the prophecy-value part remains the same as before the reborrow  $v'_o$ .

[WKN-UNQBOR-UNQREF](#) follows from the following lemma.

$$\begin{array}{c} \text{RAW-UNQBOR-UNQREF} \\ \alpha \sqsubseteq \beta \\ \hline [\mathbb{I}] \blacktriangleleft \&_{\text{uniq}}^\beta \tau \{ \lambda \pi. (\hat{v} \pi, \pi y) \} \Rightarrow_{\mathcal{N}_{\text{fit}} + \mathcal{N}_{\text{proph}}} \exists x. \\ [\mathbb{I}] \blacktriangleleft \&_{\text{uniq}}^\alpha \tau \{ \lambda \pi. (\hat{v} \pi, \pi x) \} * [\mathbb{I}] \blacktriangleleft^{\dagger\alpha} \&_{\text{uniq}}^\beta \tau \{ \lambda \pi. (\pi x, \pi y) \} \end{array}$$

Proof of [RAW-UNQBOR-UNQREF](#). We name the following proposition  $P'$ .

$$P' := \exists \bar{b}'', \hat{v}'', d''. \text{I} \mapsto \bar{b}'' * \bar{b}'' \triangleleft \tau \{ \hat{v}'' \}_{d''} * \bar{\chi}(d''+1) * \text{PC}(y, \hat{v}'', d'')$$

Out of the unique reference, we get a value observer  $\text{VO}_y(\hat{v}, d)$ , a full borrow  $\&_{\text{full}}^\beta P'$  and a persistent time receipt  $\bar{\chi}(d+1)$  for some depth  $d$ .

The full borrow of the existing unique reference  $\&_{\text{full}}^\beta P'$  is turned by [FULLBOR-REBOR](#) into a new full borrow  $\&_{\text{full}}^\alpha P'$  and a view shift  $[\alpha] \Rightarrow_{\mathcal{N}_{\text{fit}}} \&_{\text{full}}^\beta P'$ . By [VALOBS-PROPHCTRL-INTRO](#), for a fresh prophecy variable  $x$  satisfying  $x.\text{type} = \lfloor \tau \rfloor$ , we get a value observer  $\text{VO}_x(\hat{v}, d)$  and a prophecy control  $\text{PC}(x, \hat{v}, d)$ . We name the following proposition  $Q$ .

$$Q := \exists \hat{v}', d'. \text{VO}_y(\hat{v}', d') * \bar{\chi}(d'+1) * \text{PC}(x, \hat{v}', d')$$

We can create  $Q$  out of the value observer  $\text{VO}_y(\hat{v}, d)$  of the existing unique reference, the prophecy control we created  $\text{PC}(x, \hat{v}, d)$  and the persistent time receipt  $\bar{\chi}(d+1)$ . Then  $Q$  is turned into a full borrow  $\&_{\text{full}}^\alpha Q$  and a view shift  $[\alpha] \Rightarrow_{\mathcal{N}_{\text{fit}}} \triangleright Q$  ([FULLBOR-INTRO](#)).

The two obtained view shifts are merged into  $[\alpha] \Rightarrow_{\mathcal{N}_{\text{fit}}} \&_{\text{full}}^\beta P' * \triangleright Q$ , which turns into the following.

$$[\alpha] \Rightarrow_{\mathcal{N}_{\text{fit}}} \exists \hat{v}', d'. [\text{I}] \triangleleft \&_{\text{unq}}^\beta \tau \{ \lambda \pi. (\hat{v}' \pi, \pi y) \}_{d'+1} * \bar{\chi}(d'+1) * \triangleright \text{PC}(x, \hat{v}', d')$$

By [PROPHCTRL-PROPHEQZ](#), we can turn a prophecy control under later  $\triangleright \text{PC}(x, \hat{v}', d')$  into a prophecy equalizer under later  $\triangleright \text{PE}(\lambda \pi. \pi x, \hat{v}')$ . Therefore, we get a unique reference under borrow  $[\text{I}] \triangleleft \dagger^\alpha \&_{\text{unq}}^\beta \tau$ .

The two obtained full borrows are merged by [FULLBOR-MERGE](#) into  $\&_{\text{full}}^\alpha (P' * Q)$ . We temporarily access the content  $\triangleright (P' * Q)$  of the full borrow by [FULLBOR-SUBDIV](#), with help of a fractional lifetime token on  $\alpha$ . The values  $\hat{v}'', d''$  inside  $P'$  turn out to be equal to  $\hat{v}', d'$  inside  $Q$  by agreement of the prophecy control of  $P'$  and the value observer of  $Q$  ([VALOBS-PROPHCTRL-AGREE](#)). Out of  $\triangleright (P' * Q)$ , aside from the value observer and the prophecy control, we can take out  $\triangleright P$ , where  $P$  is defined as follows.

$$P := \exists \bar{b}', \hat{v}', d'. \text{I} \mapsto \bar{b}' * \bar{b}' \triangleleft \tau \{ \hat{v}' \}_{d'} * \bar{\chi}(d'+1) * \text{PC}(x, \hat{v}', d')$$

After we get back  $\triangleright P$ , by simultaneously updating the arguments of the value observer and the prophecy control ([VALOBS-PROPHCTRL-UPDATE](#)) into the new internal values  $\hat{v}', d'$  of  $P$ , we can reconstruct  $\triangleright (P' * Q)$ .

Therefore, by the power of [FULLBOR-SUBDIV](#), we finally get a new full borrow  $\&_{\text{full}}^\alpha P$ . Combining the it with the value observer that we created  $\text{VO}_x(\hat{v}, d)$ , we get the expected unique reference  $[\text{I}] \triangleleft \&_{\text{unq}}^\alpha \tau \{ \lambda \pi. (\hat{v} \pi, \pi x) \}$ .  $\square$

### 6.3.2 Access on Unique References

We can temporarily take out a full plain reference from a unique reference.

$$\begin{array}{l} \text{ACC-UNQREF-PLNREF} \\ \alpha \mid \mathbf{a} : \&_{\text{unq}}^\alpha \tau / \&_{\text{unq}}^\alpha \tau \vdash \mathbf{a} : \&_1 \tau / \&_1 \tau \mid \lambda(v, v_o).v; \lambda(v, v_o), w. (w, v_o) \end{array}$$

The prophecy value  $v_o$  is retained through this access. Unlike [ACC-BOXPTR-PLNREF](#), the target type should be retained.

[ACC-UNQREF-PLNREF](#) follows from the following lemma.

$$\begin{array}{l} \text{RAW-UNQREF-PLNREF} \\ [\text{I}] \triangleleft \&_{\text{unq}}^\alpha \tau \{ \lambda \pi. (\hat{v} \pi, \pi x) \} * [\alpha]_q \Rightarrow_{\mathcal{N}_{\text{fit}}} [\text{I}] \triangleleft \&_1 \tau \{ \hat{v} \} * \\ \forall \hat{w}. ([\text{I}] \triangleleft \&_1 \tau \{ \hat{w} \} \Rightarrow_{\mathcal{N}_{\text{fit}}} [\text{I}] \triangleleft \&_{\text{unq}}^\alpha \tau \{ \lambda \pi. (\hat{w} \pi, \pi x) \} * [\alpha]_q) \end{array}$$

*Proof of RAW-UNQREF-PLNREF.* We temporarily access the content of the full borrow of the unique reference using a partial lifetime token on  $\alpha$  (FULLBOR-ACCESS). By agreement of the value observer and the prophecy control (VALOBS-PROPHCTRL-AGREE), the target value inside the full borrow turns out to be  $\hat{v}$ . Therefore, we can take out the full plain reference  $[I] \blacktriangleleft \&_1 \tau \{ \hat{v} \}$ .

Assume that we have got back  $[I] \blacktriangleleft \&_1 \tau \{ \hat{w} \}$ . We simultaneously update the arguments of the value observer and the prophecy control (VALOBS-PROPHCTRL-UPDATE) into the new pure value  $\hat{w}$  and depth. The full plain reference has the persistent time receipt we need for the full borrow and the unique reference itself. Therefore, we can close the full borrow, getting back  $[\alpha]_q$ , and reconstruct the unique reference with the new target value  $[I] \blacktriangleleft \&_{\text{unq}}^\alpha \tau \{ \lambda \pi. (\hat{w} \pi, \pi x) \}$ .  $\square$

We can also subdivide a plain reference to a unique reference.

$$\text{ACC-DEREF-PLNREF-UNQREF} \\ \alpha \mid \mathbf{a} : \&_q \&_{\text{unq}}^\alpha \tau / \&_q \&_{\text{unq}}^\alpha \tau \vdash * \mathbf{a} : \&_1 \tau / \&_1 \tau \mid \lambda (v, v_o). v; \lambda (v, v_o), w. (w, v_o)$$

It can be proved using RAW-UNQREF-PLNREF.

*Example 6.2* (Increment on a Unique Reference). For a simple example, we can increment the target integer of a unique reference.

$$\alpha \mid \mathbf{a} : \&_{\text{unq}}^\alpha \text{int} \vdash \mathbf{a} \leftarrow * \mathbf{a} + 1 \mid \mathbf{a} : \&_{\text{unq}}^\alpha \text{int} \mid \lambda \text{post}, ((n, n_o)). \text{post} ((n + 1, n_o))$$

*Proof.* Let  $(n, n_o)$  be the original pure value of the unique reference.

For the load  $* \mathbf{a}$ , we temporarily get a plain reference by ACC-UNQREF-PLNREF and acquire an integer value by RFN-DEREF-PLNREF-COPY. We increment the integer value and bind it to some cell value  $\mathbf{b}$ , which has the pure value  $n + 1$ .

Then for the store  $\mathbf{a} \leftarrow \mathbf{b}$ , we temporarily get a full plain reference again by ACC-UNQREF-PLNREF. Finally the pure value of the unique reference is updated into  $(n + 1, n_o)$ .  $\square$

### 6.3.3 Resolution of Unique References

When we release a unique reference, we can resolve its prophecy value into its current value.

$$\text{SWKN-UNQREF-RELEASE} \\ \frac{\beta \sqsubseteq \&_{\text{unq}}^\alpha \tau}{\beta \mid \mathbf{a} : \&_{\text{unq}}^\alpha \tau \vdash^\# \mid \lambda \text{post}, ((v, v_o)). v_o = v \Rightarrow \text{post} ()}$$

We get the information  $v_o = v$  that the prophecy value  $v_o$  is equal to the current value  $v$  in the postcondition. Note that we need to ensure that both the lifetime  $\alpha$  and the lifetime of the type  $\tau$  are alive.

SWKN-UNQREF-RELEASE follows from the following lemma.

$$\text{RAW-UNQREF-RELEASE} \\ \frac{\beta \sqsubseteq \&_{\text{unq}}^\alpha \tau}{[I] \blacktriangleleft \&_{\text{unq}}^\alpha \tau \{ \lambda \pi. (\hat{v} \pi, \pi x) \}_{d+1} * [\beta]_q \stackrel{d}{\Rightarrow}_{\mathcal{N}_{\text{fit}} + \mathcal{N}_{\text{proph}}} \langle \pi. \pi x = \hat{v} \pi \rangle * [\beta]_q}$$

*Proof of RAW-UNQREF-RELEASE.* Out of the unique reference, we obtain a value observer  $\text{VO}_x(\hat{v}, d)$  (for some depth  $d$ ) and a full borrow  $\&_{\text{full}}^\alpha P$ , where  $P$  is defined as follows.

$$P := \exists \bar{\mathbf{b}}, \hat{v}', d'. I \mapsto \bar{\mathbf{b}} * \bar{\mathbf{b}} \blacktriangleleft \tau \{ \hat{v}' \}_{d'} * \bar{\mathbb{X}}(d' + 1) * \text{PC}(x, \hat{v}', d')$$

By **FULLBOR-ACCESS**, we temporarily access  $\triangleright P$  using a partial lifetime token on  $\alpha$ . We open the internal values  $\bar{\mathbf{b}}, \hat{v}', d'$  of  $P$ . We get  $\triangleright \text{PC}(x, \hat{v}', d')$ , whose later modality can be removed because we also have the value observer  $\text{VO}_x(\hat{v}, d)$  (**VALOBS-PROPHCTRL-TIMELESS**). By agreement with the value observer (**VALOBS-PROPHCTRL-AGREE**), we know  $\hat{v}' = \hat{v}$  and  $d' = d$ . We also get the target object  $\bar{\mathbf{b}} \blacktriangleleft \tau \{\hat{v}\}_d$ .

From the target object, we temporarily take out a partial prophecy token  $[Y]_q$  on a dependency  $Y$  of the target value  $\hat{v}$  in  $d$  logical steps, with help of a partial lifetime token on  $\tau$  (**SEMTY-OWN-PROPHTOKEN**). By **VALOBS-PROPHCTRL-RESOLVE**, consuming the value observer and the prophecy control, with help of the partial prophecy token  $[Y]_q$ , we get a prophecy observation  $\langle \pi. \pi x = \hat{v} \pi \rangle$  and also a new prophecy control. Using the prophecy control, we can reconstruct  $\triangleright P$ , which is put back to the full borrow  $\&_{\text{full}}^\alpha P$ . The prophecy observation  $\langle \pi. \pi x = \hat{v} \pi \rangle$  serves as the postcondition  $v_\circ = v$  for the unique reference  $(v, v_\circ)$ .  $\square$

*Example 6.3* (Release of a Unique Reference after Increment). If we increment the target value of a unique reference  $(n, n_\circ)$  and release the unique reference after that, we get the postcondition that the prophecy value  $n_\circ$  is equal to  $n + 1$ .

$$\alpha \vdash \mathbf{a} : \&_{\text{uniq}}^\alpha \text{int} \vdash \mathbf{a} \leftarrow * \mathbf{a} + 1; \not\downarrow \mid \mid \lambda \text{post}, ((n, n_\circ)). n_\circ = n + 1 \Rightarrow \text{post}()$$

*Proof.* From **SWKN-UNQREF-RELEASE** and the refined typing judgment of [Example 6.2](#). We can use the super weakening judgment for **SWKN-UNQREF-RELEASE** using the additional part ‘;  $\not\downarrow$ ’ by **RFN-SEQ**.  $\square$

Also, when we turn a unique reference into a shared reference, we can resolve the prophecy value.

$$\frac{\text{RFN-UNQREF-SHRREF} \quad \beta \sqsubseteq \&_{\text{uniq}}^\alpha \tau}{\beta \vdash \mathbf{a} : \&_{\text{uniq}}^\alpha \tau \vdash^\# \mathbf{a} : \&_{\text{shr}}^\alpha \tau \mid \lambda \text{post}, ((v, v_\circ)). v_\circ = v \Rightarrow \text{post}(v)}$$

It follows from the following lemma.

$$\frac{\text{RAW-UNQREF-SHRREF} \quad \beta \sqsubseteq \&_{\text{uniq}}^\alpha \tau}{[\mathbf{a}] \blacktriangleleft \&_{\text{uniq}}^\alpha \tau \{ \lambda \pi. (\hat{v} \pi, \pi x) \}_{d+1} * [\beta]_q \xRightarrow{\mathcal{N}_{\text{fit}} + \mathcal{N}_{\text{proph}}^{2d}} \langle \pi. \pi x = \hat{v} \pi \rangle * [\mathbf{a}] \blacktriangleleft \&_{\text{shr}}^\alpha \tau \{ \hat{v} \}_{d+1} * [\beta]_q}$$

*Proof of **RAW-UNQREF-SHRREF**.* Just like the proof of **RAW-UNQREF-RELEASE**, we consume the value observer and get a prophecy observation  $\langle \pi. \pi x = \hat{v} \pi \rangle$ , by temporarily accessing the content of the full borrow and also temporarily accessing the partial prophecy tokens of the target object spending  $d$  logical steps.

We can freeze the inner values (the cell values  $\bar{\mathbf{b}}$ , pure value  $\hat{v}'$  and depth  $d'$ ) of the full borrow (**FULLBOR-FREEZE**). Temporarily accessing the full borrow, we know  $\hat{v}' = \hat{v}$  and  $d' = d$  by agreement of **VALOBS-PROPHCTRL-AGREE**. Splitting the full borrow (**FULLBOR-SPLIT**), we get full borrows of the full points-to token and the target ownership predicate. The full borrow of the full points-to token is turned into a fractured borrow of a fractional points-to token (**FULLBOR-FRACBOR**). The full borrow of the target ownership predicate is turned into the sharing predicate on the target object in  $d$  logical steps by **SEMTY-OWN-SHR**. Therefore, we finally get the expected shared reference.  $\square$

### 6.3.4 Subdivision of Unique References

By consuming a unique reference to an object, we can take unique references to some parts of the object. We call this operation *subdivision* of a unique reference.

A unique reference to a box pointer can be dereferenced into a unique reference without changing the pure value. In this simple case, we can just reuse the prophecy variable of the original unique reference.

$$\begin{array}{c} \text{RFN-DEREF-UNQREF-BOXPTR} \\ \alpha \mid \mathbf{a} : \&_{\text{unq}}^{\alpha} \text{box } \tau \vdash * \mathbf{a} : \&_{\text{unq}}^{\alpha} \tau \mid \mid \text{id} \end{array}$$

It follows from the following lemma.

$$\begin{array}{c} \text{RAW-DEREF-UNQREF-BOXPTR} \\ [I] \triangleleft \&_{\text{unq}}^{\alpha} \text{box } \tau \{ \lambda \pi. (\hat{v} \pi, \pi x) \}_{d+1} * [\alpha]_q \Rightarrow_{\mathcal{N}_{\text{Lft}}} \exists I'. \\ I \mapsto I' * (I \mapsto I' * \bar{\Delta} 1 \Rightarrow_{\mathcal{N}_{\text{Lft}}} [I'] \triangleleft \&_{\text{unq}}^{\alpha} \tau \{ \lambda \pi. (\hat{v} \pi, \pi x) \}_d * [\alpha]_q) \end{array}$$

*Proof of RAW-DEREF-UNQREF-BOXPTR.* The given unique reference has a full borrow  $\&_{\text{full}}^{\alpha} P$ , where  $P$  can be written as follows.

$$P := \exists I', \hat{v}', d'. I \mapsto I' * [I'] \triangleleft \text{box } \tau \{ \hat{v}' \}_{d'+1} * \bar{\Delta}(d'+2) * \text{PC}(x, \hat{v}', d'+1)$$

We temporarily access the content  $\triangleright P$  of the full borrow by **FULLBOR-SUBDIV**, with help of a fractional lifetime token on  $\alpha$ . In one logical step, we strip off the later modality and get  $P$ . We open the variables  $I', \hat{v}', d'$  inside  $P$ . By agreement (**VALOBS-PROPHCTRL-AGREE**), we know  $\hat{v}' = \hat{v}$  and  $d' = d - 1$ . We temporarily take out  $I \mapsto I'$ .

Assume that we have retrieved  $I \mapsto I'$  and also obtained a cumulative time receipt  $\bar{\Delta} 1$ . From the target box pointer  $[I'] \triangleleft \text{box } \tau \{ \hat{v}' \}_{d'+1}$ , we can take out the full points-to token to the inner target  $I' \mapsto \bar{\mathbf{b}}$  and the inner target object under later  $\triangleright \bar{\mathbf{b}} \triangleleft \text{box } \tau \{ \hat{v}' \}_{d'+1}$ , for some  $\bar{\mathbf{b}}$ . We also update the depth of the value observer and prophecy control from  $d$  into  $d - 1$  (**VALOBS-PROPHCTRL-UPDATE**). Now we can take out  $\triangleright Q$ , where we define  $Q$  as follows.

$$Q := \exists \bar{\mathbf{b}}, \hat{v}', d'. I' \mapsto \bar{\mathbf{b}} * \triangleright (\bar{\mathbf{b}} \triangleleft \tau \{ \hat{v}' \}_{d'}) * \bar{\Delta}(d'+1) * \text{PC}(x, \hat{v}', d')$$

After we get back  $\triangleright Q$ , we can reconstruct  $\triangleright P$  in the following way. We update the pure value and depth of the value observer and prophecy control into  $\hat{v}'$  and  $d' + 1$ , based on the variables  $\hat{v}', d'$  inside  $Q$ . We bump up the persistent time receipt  $\bar{\Delta}(d'+1)$  of  $Q$  into the persistent time receipt  $\bar{\Delta}(d'+2)$  required by  $P$  using  $\bar{\Delta} 1$  (**CUMUTIME-SWELL-PERTIME**).

Therefore, by the power of **FULLBOR-SUBDIV**, we get the full borrow  $\&_{\text{full}}^{\alpha} Q$ . Combining it with the value observer  $\text{VO}_x(\hat{v}, d)$ , we finally get the expected unique reference  $[I'] \triangleleft \&_{\text{unq}}^{\alpha} \tau \{ \hat{v} \}_{d+1}$ .  $\square$

A unique reference to a pair can be subdivided into unique references to the elements of the pair.

$$\begin{array}{c} \text{RFN-SPLIT-UNQREF-PAIR-R} \\ \alpha \mid \mathbf{a} : \&_{\text{unq}}^{\alpha} (\tau_0 \times \tau_1) \vdash \mathbf{a}.|\tau_0| : \&_{\text{unq}}^{\alpha} \tau_1 \mid \mathbf{a} : \&_{\text{unq}}^{\alpha} \tau_0 \mid \\ \lambda \text{post}, ((v, v_0)). \text{post}((v.1, v_0.1), (v.0, v_0.0)) \end{array}$$

In order to take two unique references of the addresses  $\mathbf{a}$  and  $\mathbf{a}.|\tau|$ , we use a slightly tricky way here.

**RFN-SPLIT-UNQREF-PAIR-R** follows from the following lemma.

$$\begin{array}{c} \text{RAW-SPLIT-UNQREF-PAIR} \\ [I] \triangleleft \&_{\text{unq}}^{\alpha} (\tau_0 \times \tau_1) \{ \lambda \pi. ((\hat{v} \pi), \pi x) \}_d * [\alpha]_q \Rightarrow_{\mathcal{N}_{\text{Lft}} + \mathcal{N}_{\text{proph}}} \exists y_0, y_1. \\ \langle \pi. \pi x = (\pi y_0, \pi y_1) \rangle * *_{\mathbf{i}} ([I + \mathbf{i}.|\tau_0|] \triangleleft \&_{\text{unq}}^{\alpha} \tau_i \{ \lambda \pi. ((\hat{v} \pi).i, \pi y_i) \}_d) * [\alpha]_q \end{array}$$

We take new prophecy variables  $y_0, y_1$  for the new unique references and *resolve* the original prophecy variable  $x$  into  $\lambda\pi. (\pi y_0, \pi y_1)$ .

*Proof of RAW-SPLIT-UNQREF-PAIR.* The given unique reference consists of a value observer  $\text{VO}_x(\lambda\pi. (\hat{v} \pi), d-1)$  and a full borrow  $\&_{\text{full}}^\alpha P$ , where we define  $P$  as follows.

$$P := \exists \bar{b}_0, \bar{b}_1, \hat{v}', d'. \\ \ast_i ( \mathbb{I} + i \cdot |\tau_0| \mapsto \bar{b}_i \ast \bar{b}_i \blacktriangleleft \tau_i \{ (\cdot i) \circ \hat{v}' \}_{d'} ) \ast \bar{\Delta}(d'+1) \ast \text{PC}(x, \lambda\pi. (\hat{v}' \pi), d')$$

We temporarily access the content  $\triangleright P$  of the full borrow using **FULLBOR-SUBDIV**. We open the variables  $\bar{b}_0, \bar{b}_1, \hat{v}', d'$  inside  $P$ . By agreement (**VALOBS-PROPHCTRL-AGREE**), we know  $\hat{v}' = \hat{v}$  and  $d' = d-1$ .

For each  $i$ , taking a fresh prophecy variable  $y_i$ , we create a value observer  $\text{VO}_{y_i}((\cdot i) \circ \hat{v}, d-1)$  and a prophecy control  $\text{PC}(y_i, (\cdot i) \circ \hat{v}, d-1)$  (**VALOBS-PROPHCTRL-INTRO**). Now, by **VALOBS-PROPHCTRL-PRERESOLVE**, consuming the value observer and prophecy control of the given unique reference, with help of the prophecy tokens  $[y_i]_1$  (taken by **VALOBS-PROPHCTRL-PROPHTOKEN**), we acquire the prophecy observation  $\langle \pi. \pi x = (\pi y_0, \pi y_1) \rangle$  and the following proposition  $R$ .

$$R := \forall \hat{v}', d'. \text{PE}(\lambda\pi. (\pi y_0, \pi y_1), \hat{v}') \ast \text{PC}(x, \hat{v}', d')$$

Using the remaining parts of  $\triangleright P$  and the prophecy controls we created, we get  $\triangleright (Q_0 \ast Q_1)$ , where  $Q_i$  is defined as follows for  $i = 0, 1$ .

$$Q_i := \exists \bar{b}_i, \hat{v}'_i, d'_i. \mathbb{I} + i \cdot |\tau_0| \mapsto \bar{b}_i \ast \bar{b}_i \blacktriangleleft \tau_i \{ \hat{v}'_i \}_{d'_i} \ast \bar{\Delta}(d'_i+1) \ast \text{PC}(y_i, \hat{v}'_i, d'_i)$$

Assume that we have retrieved  $\triangleright (Q_0 \ast Q_1)$ . We can reconstruct  $\triangleright P$  in the following way. We open the variables  $\bar{b}_i, \hat{v}'_i, d'_i$  inside  $Q_i$  for  $i = 0, 1$ . For each  $i$ , we turn the prophecy control on  $y_i$  (under later)  $\triangleright \text{PC}(y_i, \hat{v}'_i, d'_i)$  owned by  $\triangleright Q_i$  into a prophecy equalizer  $\triangleright \text{PE}(\lambda\pi. \pi y_i, \hat{v}'_i)$  by **PROPHCTRL-PROPHEQZ**. We then merge the two prophecy equalizers into  $\triangleright \text{PE}(\lambda\pi. (\pi y_0, \pi y_1), \lambda\pi. (\hat{v}'_0 \pi, \hat{v}'_1 \pi))$  by **PROPHEQZ-TRANSFORM**. By applying it to  $R$ , we get the prophecy control on  $x$ ,  $\triangleright \text{PC}(x, \lambda\pi. (\hat{v}'_0 \pi, \hat{v}'_1 \pi), d')$ , letting  $d'$  be  $\max\{d'_0, d'_1\}$ . Therefore, separatingly conjoining this with the remaining parts of  $Q_0, Q_1$ , we can reconstruct  $\triangleright P$ .

Therefore, by the power of **FULLBOR-SUBDIV**, we get the full borrow  $\&_{\text{full}}^\alpha (Q_0 \ast Q_1)$ , which can be split into  $\&_{\text{full}}^\alpha Q_0$  and  $\&_{\text{full}}^\alpha Q_1$  (**FULLBOR-SPLIT**). Combining the full borrows with the value observers we have created, we finally get the expected unique references.  $\square$

A unique reference to a variant can be subdivided into a unique reference to the body of the variant.

$$\text{RFN-SUBDIV-UNQREF-VRNT} \\ \alpha \vdash \mathbf{a}: \&_{\text{unq}}^\alpha (\tau_0 + \tau_1) \vdash \mathbf{a}.1: \&_{\text{unq}}^\alpha \tau_i \vdash \vdash \\ \lambda \text{post}, ((v, v_\circ)). \exists w \text{ s.t. } \text{inj}_i w = v. \forall w_\circ \text{ s.t. } \text{inj}_i w_\circ = v_\circ. \text{post}((w, w_\circ))$$

It follows from the following lemma.

$$\text{RAW-SUBDIV-UNQREF-VRNT} \\ [\mathbb{I}] \blacktriangleleft \&_{\text{unq}}^\alpha (\tau_0 + \tau_1) \{ \lambda\pi. (\text{inj}_i(\hat{v} \pi), \pi x) \}_d \ast [\alpha]_q \Rightarrow_{\mathcal{N}_{\text{fit}} + \mathcal{N}_{\text{proph}}} \exists y. \\ \langle \pi. \pi x = \text{inj}_i(\pi y) \rangle \ast [\mathbb{I} + 1] \blacktriangleleft \&_{\text{unq}}^\alpha \tau \{ \lambda\pi. (\hat{v} \pi, \pi y) \}_d \ast [\alpha]_q$$

We take a new prophecy variable  $y$  for the new unique reference and *resolve* the original prophecy variable  $x$  into  $\lambda\pi. \text{inj}_i(\pi y)$ .

*Proof of RAW-SUBDIV-UNQREF-VRNT.* Similar to the proof of RAW-SPLIT-UNQREF-PAIR. Note that we use the injectivity of  $\text{inj}_i$  when we apply [PROPHQZ-TRANSFORM](#).  $\square$

Also, a unique reference to an injection type can be subdivided as follows, in a similar way to [RFN-SUBDIV-UNQREF-VRNT](#).

$$\begin{array}{l} \text{WKN-SUBDIV-UNQREF-INJTY} \\ \alpha \mid \mathbf{a} : \&_{\text{unq}}^\alpha \text{in}(\tau, f) \vdash \mathbf{a} : \&_{\text{unq}}^\alpha \tau \mid \\ \lambda \text{post}, ((v, v_\circ)). \forall w \text{ s.t. } f w = v. \forall w_\circ \text{ s.t. } f w_\circ = v_\circ. \text{post}((w, w_\circ)) \end{array}$$

Note that we can use universal quantification instead of existential quantification in the part ' $\forall w \text{ s.t. } f w = v.$ ', unlike [RFN-SUBDIV-UNQREF-VRNT](#). It follows from the following lemma.

$$\begin{array}{l} \text{RAW-SUBDIV-UNQREF-INJTY} \\ [\mathbf{I}] \blacktriangleleft \&_{\text{unq}}^\alpha \text{in}(\tau, f) \{ \lambda \pi. (f(\hat{v} \pi), \pi x) \}_d * [\alpha]_q \Rightarrow_{\mathcal{N}_{\text{fit}} + \mathcal{N}_{\text{proph}}} \exists y. \\ \langle \pi. \pi x = f(\pi y) \rangle * [\mathbf{I}] \blacktriangleleft \&_{\text{unq}}^\alpha \tau \{ \lambda \pi. (\hat{v} \pi, \pi y) \}_d * [\alpha]_q \end{array}$$

*Proof of RAW-SUBDIV-UNQREF-INJTY.* Just analogous to the proof of [RAW-SUBDIV-UNQREF-VRNT](#). Note that the function  $f$  is injective.  $\square$

The rule [WKN-SUBDIV-UNQREF-INJTY](#) is useful for modifying a unique reference to a recursive type.

A unique reference to a unique reference (which we call a double unique reference) can be dereferenced into a unique reference.

$$\begin{array}{l} \text{RFN-DEREF-UNQREF-UNQREF} \\ \alpha \sqcap \beta \mid \mathbf{a} : \&_{\text{unq}}^\alpha \&_{\text{unq}}^\beta \tau \vdash * \mathbf{a} : \&_{\text{unq}}^{\alpha \sqcap \beta} \tau \mid \mid \\ \lambda \text{post}, (((v, v^\circ), (v_\circ, v_\circ^\circ))). v_\circ^\circ = v^\circ \Rightarrow \text{post}((v, v_\circ)) \end{array}$$

We handle prophecy values in an interesting way. The current and prophecy values  $v, v_\circ$  of the resulting unique reference are the current values of the current and prophecy values of the given double unique reference, respectively. We get the postcondition  $v_\circ^\circ = v^\circ$  that the prophecy values  $v^\circ, v_\circ^\circ$  of the current and prophecy values are equal.

[RAW-DEREF-UNQREF-UNQREF](#) follows from the following lemma.

$$\begin{array}{l} \text{RAW-DEREF-UNQREF-UNQREF} \\ [\mathbf{I}] \blacktriangleleft \&_{\text{unq}}^\alpha \&_{\text{unq}}^\beta \tau \{ \lambda \pi. ((\hat{v} \pi, \pi y), \pi x) \}_{d+1} * [\alpha \sqcap \beta]_q \Rightarrow_{\mathcal{N}_{\text{fit}} + \mathcal{N}_{\text{proph}}} \exists I', z. \\ \mathbf{I} \xrightarrow{1} I' * \langle \pi. \pi x = (\pi z, \pi y) \rangle * \\ (\mathbf{I} \xrightarrow{1} I' * \exists 1 \Rightarrow_{\mathcal{N}_{\text{fit}}} [\mathbf{I}'] \blacktriangleleft \&_{\text{unq}}^{\alpha \sqcap \beta} \tau \{ \lambda \pi. (\hat{v} \pi, \pi z) \}_d * [\alpha \sqcap \beta]_q) \end{array}$$

*Proof of RAW-DEREF-UNQREF-UNQREF.* The given double unique reference is decomposed into the value observer  $\text{VO}_x(\lambda \pi. (\hat{v} \pi, \pi y), d)$  and the full borrow that can be written as follows.

$$\&_{\text{full}}^\alpha (\exists I', \hat{v}', y, d'. \mathbf{I} \xrightarrow{1} I' * [\mathbf{I}'] \blacktriangleleft \&_{\text{unq}}^\beta \tau \{ \lambda \pi. (\hat{v}' \pi, \pi y) \}_{d'+1} * \exists (d'+2) * \text{PC}(y, \hat{v}', d'+1))$$

We freeze the inner variables  $I', y$  of the full borrow ([FULLBOR-FREEZE](#)) and split the full borrow (by [FULLBOR-SPLIT](#)) into the following two.

$$\begin{array}{l} \&_{\text{full}}^\alpha (\exists \hat{v}', d'. \mathbf{I} \xrightarrow{1} I' * \text{VO}_y(\hat{v}', d') * \text{PC}(x, \lambda \pi. (\hat{v}' \pi, \pi y), d'+1) * \exists (d'+2)) \\ \&_{\text{full}}^\alpha \&_{\text{full}}^\beta (\exists \bar{\mathbf{b}}, \hat{v}', d'. \mathbf{I} \xrightarrow{1} \bar{\mathbf{b}} * \bar{\mathbf{b}} \blacktriangleleft \tau \{ \hat{v}' \}_{d'} * \exists (d'+1) * \text{PC}(y, \hat{v}', d')) \end{array}$$

We unnest the latter in one logical step (**FULLBOR-UNNEST**) and then merge it with the former (**FULLBOR-MERGE**) to get the full borrow  $\&_{\text{full}}^{\alpha\Gamma\beta} P$ , where we define  $P$  as follows.

$$P := \exists \bar{b}, \hat{v}', d'. I \mapsto I' * I' \mapsto \bar{b} * \bar{b} \blacktriangleleft \tau \{ \hat{v}' \}_{d'} * \bar{\Delta}(d' + 2) \\ \text{VO}_y(\hat{v}', d') * \text{PC}(y, \hat{v}', d') * \text{PC}(x, \lambda\pi.(\hat{v}'\pi, \pi y), d' + 1)$$

We access the content  $\triangleright P$  of this full borrow by **FULLBOR-SUBDIV**. We can temporarily take out the points-to token  $I \mapsto I'$ .

Assume we have retrieved  $I \mapsto I'$  and also obtained a cumulative time receipt  $\bar{\Delta} 1$ . By agreement of the value observer and prophecy control on  $x$  (**VALOBS-PROPHCTRL-AGREE**), we know that  $d$  is positive. Taking a fresh prophecy variable  $z$ , we create a value observer  $\text{VO}_z(\hat{v}, d - 1)$  and a prophecy control  $\text{PC}(z, \hat{v}, d - 1)$ . By **VALOBS-PROPHCTRL-PRERESOLVE**, consuming the value observer and prophecy control of the given double unique reference, with help of prophecy tokens  $[y]_1$  and  $[z]_1$  (taken by **VALOBS-PROPHCTRL-PROPHTOKEN**), we obtain the prophecy observation  $\langle \pi. \pi x = (\pi z, \pi y) \rangle$  and the following proposition  $R$ .

$$R := \forall \hat{v}', d'. \text{PE}(\lambda\pi.(\pi z, \pi y), \hat{v}') -* \text{PC}(x, \hat{v}', d')$$

From the remaining parts of  $\triangleright P$  and the created prophecy control  $\text{PC}(z, \hat{v}, d - 1)$ , we can take out  $\triangleright Q$ , where we define  $Q$  as follows.

$$Q := \exists \bar{b}, \hat{v}', d'. I' \mapsto \bar{b} * \bar{b} \blacktriangleleft \tau \{ \hat{v}' \}_{d'} * \bar{\Delta}(d' + 1) * \text{PC}(z, \hat{v}', d')$$

When we get back  $\triangleright Q$ , we can reconstruct  $\triangleright P$  in the following way. We open the variables  $\bar{b}, \hat{v}', d'$  of  $Q$ . We turn the prophecy control on  $z$  (under later)  $\triangleright \text{PC}(z, \hat{v}', d')$  of  $\triangleright Q$  into the prophecy equalizer  $\triangleright \text{PE}(\lambda\pi. \pi z, \hat{v}')$  (**PROPHCTRL-PROPEQZ**) and then turn it into  $\triangleright \text{PE}(\lambda\pi. (\pi z, \pi y), \lambda\pi. (\hat{v}'\pi, \pi y))$  (**PROPEQZ-TRANSFORM**); applying it to  $R$ , we get the prophecy control on  $x$ ,  $\triangleright \text{PC}(x, \hat{v}, d + 2)$ . We update the arguments of the value observer and prophecy control on  $y$  into  $\hat{v}'$  and  $d'$  (**VALOBS-PROPHCTRL-UPDATE**). We bump up the persistent time receipt  $\bar{\Delta}(d' + 1)$  of  $Q$  into  $\bar{\Delta}(d' + 2)$  required by  $P$  by consuming  $\bar{\Delta} 1$  (**CUMUTIME-SWELL-PERTIME**).

Therefore, by the power of **FULLBOR-SUBDIV**, we get the full borrow  $\&_{\text{full}}^{\alpha\Gamma\beta} Q$ . Combining it with the value observer that we created  $\text{VO}_z(\hat{v}, d - 1)$ , we finally get the expected unique reference  $\&_{\text{unq}}^{\alpha\Gamma\beta} \tau \{ \lambda\pi. (\hat{v}'\pi, \pi z) \}_d$ .  $\square$

### 6.3.5 Subtyping on Unique References

We have the following structural subtyping rule on the unique reference type under the access mode own.

$$\frac{\text{SUBTY-OWN-UNQREF} \\ \triangleright (\tau \sqsubseteq^{\text{own}} \tau' \mid \text{id}) \quad \triangleright (\tau' \sqsubseteq^{\text{own}} \tau \mid \text{id}) \quad \beta \sqsubseteq \alpha}{\&_{\text{unq}}^{\alpha} \tau \sqsubseteq^{\text{own}} \&_{\text{unq}}^{\beta} \tau' \mid \text{id}}$$

Again, the subtyping assumptions on the target types can be under the later modality.

*Proof.* From **FULLBOR-IFF** and **FULLBOR-MONO-LFT**.  $\square$

Using the later modality in the assumption of **SUBTY-OWN-UNQREF**, with Löb induction (**LÖB**), we can prove subtyping rules (under the access mode own) for recursive types with self reference under the unique reference type.

*Example 6.4* (Subtyping for Recursion Under Unique References). For example, we can derive the following subtyping rule on the type  $\text{mnat}_\alpha$  (defined in [Example 5.2](#)), a Peano number type with self reference under the unique reference type.

$$\frac{\alpha \sqsubseteq \beta \quad \beta \sqsubseteq \alpha}{\text{mnat}_\alpha \sqsubseteq^{\text{own}} \text{mnat}_\beta \mid \text{id}}$$

*Proof.* Assume  $\alpha \sqsubseteq \beta$  and  $\beta \sqsubseteq \alpha$ . We prove both  $P := \text{mnat}_\alpha \sqsubseteq^{\text{own}} \text{mnat}_\beta \mid \text{id}$  and  $Q := \text{mnat}_\beta \sqsubseteq^{\text{own}} \text{mnat}_\alpha \mid \text{id}$ . By Löb induction ([LÖB](#)), we can also assume  $\triangleright(P \wedge Q)$ , which entails  $\triangleright P$  and  $\triangleright Q$ .

By [SUBTY-OWN-UNQREF](#), we can prove  $P$  from  $\triangleright P$ ,  $\triangleright Q$  and  $\beta \sqsubseteq \alpha$ . Similarly we can prove  $Q$ . Therefore, we have proved both  $P$  and  $Q$ .  $\square$

Under the sharing access mode, we have the following more liberated subtyping rule on the unique reference type.

$$\frac{\text{SUBTY-SHR-UNQREF} \quad \triangleright (\tau \sqsubseteq^{\text{shr}(\alpha \sqcap \beta)} \tau' \mid f) \quad \alpha' \sqsubseteq \alpha}{\&_{\text{unq}}^\alpha \tau \sqsubseteq^{\text{shr}(\beta)} \&_{\text{unq}}^{\alpha'} \tau' \mid \lambda(v, v_0). (f v, g v_0)}$$

We can modify the prophecy part  $v_0$  using any function  $g$ , which follows from [DEP-CONSTRUCT](#) (see the model of the unique reference in [§5.3](#)).

Also, a unique reference can be transformed into a shared reference under the sharing access mode.

$$\text{SUBTY-SHR-UNQREF-SHRREF} \quad \&_{\text{unq}}^\beta \tau \sqsubseteq^{\text{shr}(\alpha)} \&_{\text{shr}}^{\alpha \sqcap \beta} \tau \mid \lambda(v, v_0). v$$

Combining this with [RFN-DEREF-SHRREF-SHRREF](#), we can derive the following dereference rule.

$$\frac{\text{RFN-DEREF-SHRREF-UNQREF} \quad \beta \sqsubseteq \alpha}{\beta \mid \mathbf{a} : \&_{\text{shr}}^\alpha \&_{\text{unq}}^{\alpha'} \tau \vdash * \mathbf{a} : \&_{\text{shr}}^{\alpha \sqcap \alpha'} \tau \mid \mid \lambda \text{post}, ((v, v_0)). \text{post}(v)}$$

### 6.3.6 Manipulating Vectors via Unique References

Now we introduce refined typing rules for combination of the vector type  $\text{vec } \tau$  ([§5.2](#)) and the unique reference type  $\&_{\text{unq}}^\alpha \tau$ .

A function that takes a unique reference to a vector, updates the vector, and just releases the unique reference can be verified using [RAW-UNQREF-PLNREF](#) and [RAW-UNQREF-RELEASE](#). Although we need to work on the model of the vector type, we need no other rules on unique references.

For example, we can introduce the following function  $\text{pop}_{\text{vec}}^\tau$  that takes a unique reference to a vector  $a$  and moves out the last element from the vector making a new box pointer.

$$\text{pop}_{\text{vec}}^\tau := \text{fn}(a) \{ \text{let } b = \text{alloc } |\tau| \text{ in} \\ \text{let } l = *(a.2) - 1 \text{ in } a.2 \leftarrow l; b \leftarrow_{|\tau|}^* (*a).(l \times |\tau|); b \}$$

This function satisfies the following specification.

$$\text{RFN-POP-VEC} \quad \text{pop}_{\text{vec}}^\tau : \forall \alpha. \text{fn}(\&_{\text{unq}}^\alpha \text{vec } \tau) \rightarrow \text{box } \tau \mid \\ \lambda \text{post}, ((v, v_0)). \exists v', v'' \text{ s.t. } v' \# [v''] = v. v_0 = v' \Rightarrow \text{post}(v'')$$

Note that the precondition requires that the vector is non-empty. In order to prove this specification, we simply take out a full plain reference from the input unique reference by [RAW-UNQREF-PLNREF](#), verify the operations based on the model of the type, and then resolve the prophecy value by [RAW-UNQREF-RELEASE](#).

Also, we can introduce the following function  $\text{push}_{\text{vec}}^\tau$  that takes a unique reference to a vector  $a$  and an ownership  $b$  and pushes the target object of  $b$  to the end of the vector.

$$\begin{aligned} \text{push}_{\text{vec}}^\tau &:= \text{fn}(a, b) \{ \text{let } c = *(a.1) \text{ in let } l = *(a.2) \text{ in} \\ &\quad \text{if } c = l \{ \text{let } a' = \text{alloc}((c + 1) \times |\tau|) \text{ in } a' \leftarrow_{c \times |\tau|}^* *a; a \leftarrow a'; a.1 \leftarrow c + 1 \}; \\ &\quad (*a).(l \times |\tau|) \leftarrow_{|\tau|}^* b; \text{free } b; a.2 \leftarrow l + 1 \} \end{aligned}$$

If we have no room in the capacity  $c = l$ , we increment the capacity size and reallocate the memory block. We move the target object to the end address of the vector, free the memory block of  $b$ , and increment the length of the vector. The function  $\text{push}_{\text{vec}}^\tau$  satisfies the following specification.

$$\begin{aligned} &\text{RFN-PUSH-VEC} \\ \text{push}_{\text{vec}}^\tau &: \forall \alpha. \text{fn}(\&_{\text{unq}}^\alpha \text{vec } \tau, \text{box } \tau) \mid \lambda \text{post}, ((v, v_o), v'). v_o = v \# [v'] \Rightarrow \text{post}() \end{aligned}$$

Again, the rules we need about the unique reference are just [RAW-UNQREF-PLNREF](#) and [RAW-UNQREF-RELEASE](#), although we need to work on the big expression of  $\text{push}_{\text{vec}}^\tau$  and the model of the vector type.

We can also *subdivide* a unique reference to a vector into unique references to each element of the vector, which requires some new lemma on unique references. We can introduce the following lemma for that subdivision.

$$\begin{aligned} &\text{RAW-SPLIT-UNQREF-VEC} \\ \mathbb{I} \blacktriangleleft \&_{\text{unq}}^\alpha \text{vec } \tau \{ \lambda \pi. ([\hat{v} \pi], \pi x) \}_{d+1} * [\alpha]_q &\Rightarrow_{\mathcal{N}_{\text{fit}} + \mathcal{N}_{\text{proph}}} \exists \mathbb{I}', \vec{y}. \\ \mathbb{I} \xrightarrow{1} \mathbb{I}' * \langle \pi. \pi x = [\vec{\pi} \vec{y}] \rangle * & \\ (\mathbb{I} \xrightarrow{1} \mathbb{I}' * \bigwedge 1 \Rightarrow_{\mathcal{N}_{\text{fit}}} *_{i} ([\mathbb{I}' + i \cdot |\tau|] \blacktriangleleft \&_{\text{unq}}^\alpha \tau \{ \lambda \pi. (\hat{v}_i \pi, \pi y_i) \}_d) * [\alpha]_q) & \end{aligned}$$

*Proof.* Similar to [RAW-DEREF-UNQREF-BOXPTR](#) and [RAW-SPLIT-UNQREF-PAIR](#).  $\square$

Combining [RAW-SPLIT-UNQREF-VEC](#) and [RAW-UNQREF-RELEASE](#), we can prove the following lemma for taking out of a unique reference to a vector a unique reference to a specific element of the vector. The prophecy variables of the remaining unique references are resolved.

$$\begin{aligned} &\text{RAW-VEC-IDX-UNQREF} \\ \vec{v} = \hat{v}_0, \dots, \hat{v}_{n-1} \quad 0 \leq i < n \quad \beta \sqsubseteq \&_{\text{unq}}^\alpha \text{vec } \tau & \\ \hline \mathbb{I} \blacktriangleleft \&_{\text{unq}}^\alpha \text{vec } \tau \{ \lambda \pi. ([\hat{v} \pi], \pi x) \}_{d+1} * [\alpha]_q &\Rightarrow_{\mathcal{N}_{\text{fit}} + \mathcal{N}_{\text{proph}}} \exists \mathbb{I}'. \mathbb{I} \xrightarrow{1} \mathbb{I}' * \\ (\mathbb{I} \xrightarrow{1} \mathbb{I}' * \bigwedge 1 \Rightarrow_{\mathcal{N}_{\text{fit}}} *_{i}^d \exists y. & \\ \langle \pi. \pi x = [\hat{v} \pi] \{ i \leftarrow \pi y \} \rangle * [\mathbb{I}' + i \cdot |\tau|] \blacktriangleleft \&_{\text{unq}}^\alpha \tau \{ \lambda \pi. (\hat{v}_i \pi, \pi y) \}_d * [\alpha]_q) & \end{aligned}$$

Note that we can obtain  $\beta \sqsubseteq \alpha$  out of  $\beta \sqsubseteq \&_{\text{unq}}^\alpha \text{vec } \tau$  in two logical steps. This lemma leads to the following refined typing rule for the function  $\text{idx}_{\text{vec}}^\tau$  (defined in [§6.2.5](#)).

$$\begin{aligned} &\text{RFN-IDX-VEC-UNQREF} \\ \text{idx}_{\text{vec}}^\tau &: \forall \alpha. \text{fn}(\&_{\text{unq}}^\alpha \text{vec } \tau, \text{int}) \rightarrow \&_{\text{unq}}^\alpha \tau \mid \\ &\quad \lambda \text{post}, ((v, v_o), i). 0 \leq i < \text{len } v \wedge \forall w_o. v_o = v \{ i \leftarrow w_o \} \Rightarrow \text{post}((v[i], w_o)) \end{aligned}$$

The precondition ensures that the index  $i$  is within the bound of the current list  $v$  ( $0 \leq i < \text{len } v$ ). We take a prophecy value  $w_o$  for the newly taken unique reference to the

$i$ -th element. The new unique reference has the value  $(v[i], w_\circ)$ . We set the prophecy value  $v_\circ$  on the vector to the current value  $v$  with the  $i$ -th element updated to  $w_\circ$ .

We can also write a following *higher-order function* that inputs a reference  $a$  to a vector and a function  $f$  and calls  $f$  on a reference to each element of the vector.

$$\text{iter}_{\text{vec}}^\tau := \text{fn}(a, f) \{ \text{let } a' = *a \text{ in for } i \leftarrow 0 .. *(a.2) \{ f(a'.(i \times |\tau|)) \} \}$$

Using [RAW-SPLIT-UNQREF-VEC](#) and [RAW-UNQREF-RELEASE](#), we can prove the following refined typing rule on this function.

$$\begin{aligned} & \text{RFN-ITER-VEC-UNQREF} \\ \text{iter}_{\text{vec}}^\tau : \forall \alpha. \text{fn}(\&_{\text{unq}}^\alpha \text{vec } \tau, \text{fn}(\&_{\text{unq}}^\alpha \tau)) \vdash \lambda \text{post}, ((v, v_\circ), \text{pre}). \\ & \text{len } v_\circ = \text{len } v \Rightarrow \text{foldr}(\lambda w, \text{post}', (())). \text{pre } \text{post}'(w) (\text{zip } v v_\circ) \text{post}((\ )) \end{aligned}$$

By [RAW-SPLIT-UNQREF-VEC](#), we get the postcondition  $\text{len } v_\circ = \text{len } v$  and have the unique references  $(v[i], v_\circ[i])$  (for  $0 \leq i < \text{len } v$ ). In order to compose the predicate transformers on the unique reference, we use here  $\text{foldr}: (T \rightarrow U \rightarrow U) \rightarrow \text{List } T \rightarrow U \rightarrow U$  and  $\text{zip}: \text{List } T \rightarrow \text{List } T \rightarrow \text{List}(T \times T)$ , which are defined as follows.

$$\begin{aligned} \text{foldr } f \text{ nil } a &:= a & \text{foldr } f (v :: w) a &:= f v (\text{foldr } f w a) \\ \text{zip nil } w = \text{zip } v \text{ nil} &:= \text{nil} & \text{zip } (v :: v') (w :: w') &:= (v, w) :: \text{zip } v' w' \end{aligned}$$

## 6.4 Examples of Verification by the Refined Type System

In this section, we verify some Rust programs using the refined typing rules introduced in the previous sections, [§6.1](#), [§6.2](#) and [§6.3](#).

*Example 6.5* (Dynamic Decision of the Address of a Unique Reference). The verification problem [Example 1.9](#) verified by RustHorn can also be verified in our refined type system.

The function `take_max` can be formalized as the following function `take_max`.

$$\text{take\_max} := \text{fn}(a, b) \{ \text{if } *a \geq *b \{ \downarrow; a \} \text{ else } \{ \downarrow; b \} \}$$

It takes two unique references to an integer  $a, b$  and returns the one with the greater value, releasing the other. The point is that the returned unique reference ( $a$  or  $b$ ) is determined by a dynamic condition  $*a \geq *b$ . This function satisfies the following specification.

$$\begin{aligned} \text{take\_max} : \forall \alpha. \text{fn}(\&_{\text{unq}}^\alpha \text{int}, \&_{\text{unq}}^\alpha \text{int}) \rightarrow \&_{\text{unq}}^\alpha \text{int} \vdash \lambda \text{post}, ((n, n_\circ), (m, m_\circ)). \\ (n \geq m \Rightarrow m_\circ = m \Rightarrow \text{post}((n, n_\circ))) \wedge (n < m \Rightarrow n_\circ = n \Rightarrow \text{post}((m, m_\circ))) \end{aligned}$$

It follows from [ACC-UNQREF-PLNREF](#), [RFN-DEREF-PLNREF-COPY](#), [RFN-INTREL](#), [RFN-IF](#), and [SWKN-UNQREF-RELEASE](#).

The test function `test` can be formalized as the following function `inc_max`.

$$\begin{aligned} \text{inc\_max} := \text{fn}(a, b) \{ \text{let } a' = \text{alloc } a \text{ in let } b' = \text{alloc } b \text{ in } a' \leftarrow a; b' \leftarrow b; \\ \text{let } c' = \text{take\_max}(a', b') \text{ in } c' \leftarrow *c' + 1; \downarrow; \text{let } r = *a' \neq *b' \text{ in free } a'; \text{free } b'; r \} \end{aligned}$$

It allocates two integer values  $a, b$ , uniquely borrows the box pointers, increments the target of the unique reference with the larger value taken by `take_max`, and checks that the two integer values have become different. We can verify that this function always return `tt`, i.e., the following specification holds.

$$\text{inc\_max} : \text{fn}(\text{int}, \text{int}) \rightarrow \text{bool} \vdash \lambda \text{post}, (m, n). \text{post}(\text{tt})$$

The verification goes as follows. By [RFN-ALLOC](#), [RFN-INTRO-LOCALFT](#), [WKN-UNQBOR-BOXPTR](#), [RFN-FN-CALL](#), [SWKN-UNQREF-RELEASE](#), [SWKN-END-LOCALFT-RECLAIM](#), and [RFN-FREE](#), we first obtain the following predicate transformer.

$$\lambda post, (n, m). \forall n_o, m_o. (n \geq m \Rightarrow m_o = m \Rightarrow n_o = n + 1 \Rightarrow post (n_o \neq m_o)) \wedge \\ (n < m \Rightarrow n_o = n \Rightarrow m_o = m + 1 \Rightarrow post (n_o \neq m_o))$$

Now we can easily check that this is equivalent to  $\lambda post, (m, n). post (tt)$ .

*Example 6.6* (Update of a List via a Unique Reference). The verification problem [Example 1.10](#) verified by RustHorn can also be verified in our refined type system.

First, the function `sum` can be formalized as the following function `sumlist`.

$$\text{sum}_{\text{list}} := \text{fn } self(a) \{ \text{case } *a \text{ of } \{ \\ 0 \rightarrow 0, \quad 1 \rightarrow \text{let } a' = a.1 \text{ in } *a' + self(*a'.1) \\ \} \}$$

It satisfies the following specification using the function `sum : List ℤ → ℤ`.

$$\text{sum}_{\text{list}} : \forall \alpha. \text{fn}(\&_{\text{shr}}^{\alpha} \text{ list int}) \rightarrow \text{int} \mid \lambda post, (v). post (\text{sum } v)$$

We elaborate its proof here.

*Proof.* We use [RFN-FN-REC](#). We can assume that `self` is modeled as the predicate transformer  $\lambda post, (v). post (\text{sum } v)$ . The variable `a` has the type  $\&_{\text{shr}}^{\alpha} \text{ list int}$ . We know that the lifetime  $\alpha$  is alive during the function call. Let  $v : \text{List } \mathbb{Z}$  the pure value of `a`. By [SUBTY-TY-INJTY](#), we transform `a` into the type  $\&_{\text{shr}}^{\alpha} (\zeta_0 + \text{int} \times \text{box list int})$  and the value `outList v : Unit + ℤ × List ℤ`. By [ACC-SHRREF-PLNREF](#) and [RFN-CASE-VRNT-ACCESS](#), we temporarily access the tag of `a`.

Let us discuss the case 0. By the effect of [RFN-CASE-VRNT-ACCESS](#), we know  $v = \text{nil}$ . Since `sum nil = 0` holds and we return 0 here, the postcondition is satisfied.

Let us discuss the case 1. We know that  $v$  is of the form  $n :: w$ . By [RFN-SUBDIV-SHRREF-VRNT](#), the variable `a'` defined as `a.1` has the type  $\&_{\text{shr}}^{\alpha} (\text{int} \times \text{box list int})$  and the value  $(n, w)$ . By [SUBTY-SUBDIV-SHRREF-PAIR-L](#), `a'` can be used as a shared reference of the type  $\&_{\text{shr}}^{\alpha} \text{ int}$  and the value  $n$ . By [ACC-SHRREF-PLNREF](#) and [RFN-DEREF-PLNREF-COPY](#), we get an integer of the value  $n$  from `*a'`. By [RFN-SUBDIV-SHRREF-PAIR-R](#) and [RFN-DEREF-SHRREF-BOXPTR](#), `*(a'.1)` is a shared reference the type  $\&_{\text{shr}}^{\alpha} \text{ list int}$  and the value  $w$ . By passing it to `self`, we get an integer of the value `sum w`. By [HOARE-INTOP](#), we finally get the value  $n + \text{sum } w$ , which is equal to `sum(n :: w)`, i.e., `sum v`. So the postcondition is satisfied.  $\square$

The function `take_some` can be formalized as the following function `take_some`.

$$\text{take\_some} := \text{fn } self(a) \{ \text{case } *a \text{ of } \{ \\ 0 \rightarrow self(a), \\ 1 \rightarrow \text{let } a' = a.1 \text{ in let } a'' = *(a'.1) \text{ in if } \text{ndint} \geq 0 \{ \zeta; a' \} \text{ else } \{ \zeta; self(a'') \} \\ \} \}$$

We can prove that it satisfies the following specification, which is enough for verification of the test function `test`.

$$\text{take\_some} : \forall \alpha. \text{fn}(\&_{\text{unq}}^{\alpha} \text{ list int}) \rightarrow \&_{\text{unq}}^{\alpha} \text{ int} \mid \\ \lambda post, ((v, v_o)). \forall n, n_o. (n_o = n + 1 \Rightarrow \text{sum } v_o = \text{sum } v + 1) \Rightarrow post ((n, n_o))$$

The predicate transformer means that we get the postcondition  $n_o = n + 1 \Rightarrow \text{sum } v_o = \text{sum } v + 1$  on the result  $(n, n_o)$ .

*Proof.* We use [RFN-FN-REC](#). We can assume that *self* satisfies the expected specification. The variable *a* has the type  $\&_{\text{uniq}}^\alpha \text{ list int}$ . Let  $(v, v_o)$  be the pure value of *a*. We update the type of *a* into  $\&_{\text{uniq}}^\alpha (\zeta_0 + \text{int} \times \text{box list int})$  by [SUBTY-TY-INJTY](#) and temporarily access the tag by [ACC-UNQREF-PLNREF](#) and [RFN-CASE-VRNT-ACCESS](#). In the case 0, we revert the type of *a*, and the postcondition is trivially satisfied by the recursive call on *self*.

Let us discuss the case 1. We know that *v* has the form  $n :: v'$ . By [RFN-SUBDIV-UNQREF-VRNT](#), the variable  $a' = a.1$  has the type  $\&_{\text{uniq}}^\alpha (\text{int} \times \text{box list int})$  and the value  $((n, w), (n_o, w_o))$  for some  $n_o, w_o$ . By [RFN-SPLIT-UNQREF-PAIR-R](#), we split  $a'$  into two unique references, (i)  $a'.1$ , which has the type  $\&_{\text{uniq}}^\alpha \text{ box list int}$  and the value  $(w, w_o)$  and (ii)  $a'$ , which has the type  $\&_{\text{uniq}}^\alpha \text{ int}$  and the value  $(n, n_o)$ . Then by [RFN-DEREF-UNQREF-BOXPTR](#), we subdivide the former into the variable  $a'' = *(a'.1)$ , which has the type  $\&_{\text{uniq}}^\alpha \text{ list int}$  and the value  $(w, w_o)$ . Now by [RFN-VAL-BOOL](#), it suffices to verify the then case and the else case.

In the then case, we release  $a''$  by [SWKN-UNQREF-RELEASE](#) and get the postcondition  $w_o = w$ . We return  $a'$ , which has the value  $(n, n_o)$ . When  $n_o = n + 1$  holds, we have  $\text{sum } v_o = \text{sum}(n_o :: w_o) = n_o + \text{sum } w_o = n + 1 + \text{sum } w = \text{sum}(n :: w) + 1 = \text{sum } v + 1$ . Therefore, the postcondition is satisfied.

In the else case, we release  $a'$  by [SWKN-UNQREF-RELEASE](#) and get the postcondition  $n_o = n$ . We recursively call *self* with  $a''$ . Let  $(m, m_o)$  be the finally returned value. Assume that  $m_o = m + 1$  holds. By the assumption on *self*, we know  $\text{sum } w_o = \text{sum } w + 1$ . Therefore, we have  $\text{sum } v_o = \text{sum}(n_o :: w_o) = n_o + \text{sum } w_o = n + \text{sum } w + 1 = \text{sum}(n :: w) + 1 = \text{sum } v + 1$ . Therefore, the postcondition is satisfied.  $\square$

The test function `test` can be formalized as the following function `test_take_some` (we omit the structural deallocation of the list).

$$\begin{aligned} \text{test\_take\_some} &:= \text{fn}(a) \{ \text{let } n = \text{sum}_{\text{list}}(a) \text{ in} \\ &\quad \text{let } b = \text{take\_some}(a) \text{ in } b \leftarrow *b + 1; \zeta; \text{sum}_{\text{list}}(a) = n + 1 \} \end{aligned}$$

It satisfies the following specification.

$$\text{test\_take\_some} : \text{box list int} \rightarrow \text{bool} \mid \lambda \text{post}, (v). \text{post}(\text{tt})$$

It can be proved using the specifications on  $\text{sum}_{\text{list}}$  and `take_some`, as well as [WKN-UNQBOR-BOXPTR](#) and [SWKN-UNQREF-RELEASE](#).

*Example 6.7* (Splitting a Unique Reference to a List). The following function splits the unique reference to a list *a* into a box pointer to a list of unique references to each element of the list.

$$\begin{aligned} \text{split}_{\text{list}}^\tau &:= \text{fn } \text{self}(a) \{ \text{let } b = \text{alloc } 3 \text{ in case } *a \text{ of } \{ \\ &\quad 0 \rightarrow b \leftarrow 0; b, \\ &\quad 1 \rightarrow \text{let } a' = a.1 \text{ in } (b \leftarrow 1; b.2 \leftarrow \text{self}(*(a'.|\tau|)); b.1 \leftarrow a'); b \\ &\quad \} \} \end{aligned}$$

We can derive the following specification on this function.

$$\begin{aligned} \text{split}_{\text{list}}^\tau &: \forall \alpha. \text{fn}(\&_{\text{uniq}}^\alpha \text{ list } \tau) \rightarrow \text{box list } \&_{\text{uniq}}^\alpha \tau \mid \\ &\quad \lambda \text{post}, ((v, v_o)). \text{len } v_o = \text{len } v \Rightarrow \text{post}(\text{zip } v \ v_o) \end{aligned}$$

*Proof.* We use [RFN-FN-REC](#). We can assume that *self* satisfies the expected specification.

After the allocation `alloc 3`, *b* is typed  $\text{box } \zeta_3$  by [RFN-ALLOC](#). We convert the type of the unique reference *a* into  $\&_{\text{uniq}}^\alpha (\zeta_0 + \tau \times \text{box list } \tau)$  by [WKN-SUBDIV-UNQREF-INJTY](#). We perform the case operation using [RFN-CASE-VRNT-ACCESS](#).

Let us discuss the case where the tag is 0. We convert the type of the unique reference  $a$  back to  $\&_{\text{uniq}}^\alpha$  list  $\tau$  by `WKN-SUBDIV-UNQREF-INJTY`. We know that the current target value of  $a$  is nil by the effect of `RFN-CASE-VRNT-ACCESS`. By `ACC-BOXPTR-PLNREF` and `SUBTY-SPLIT-INVALID`, we can temporarily take out a full plain reference  $\&_1((\&_1 \times \&_0) \times \&_2)$  out of  $b$ . By `RFN-STORE-PLNREF-PLNREF-COPY` and `WKN-PLNREF-TRIPLE-VRNT`, we can store  $\text{inj}_0()$  to the box pointer  $b$ , which is turned into nil by `SUBTY-TY-INJTY`. Then we release the unique reference  $a$  by `SWKN-UNQREF-RELEASE`, which resolves the prophecy value of  $a$  into nil. Since  $\text{len nil} = \text{nil}$  and  $\text{zip nil nil} = \text{nil}$  hold, the specification is satisfied in this case.

Let us discuss the case where the tag is 1. We convert the type of the unique reference  $a$  back to  $\&_{\text{uniq}}^\alpha$  list  $\tau$ . We know that the pure value of  $a$  is of the form  $(v :: w, w'_o)$ . By `RAW-SUBDIV-UNQREF-VRNT`, the unique reference  $a.1$  named  $a'$  has the type  $\&_{\text{uniq}}^\alpha(\tau \times \text{box list } \tau)$  and the pure value  $((v, w), (v_o, w_o))$ , where  $w'_o$  has been resolved into  $v_o :: w_o$ . Inside the parenthesis, we access  $b$  by `ACC-BOXPTR-PLNREF` and `SUBTY-SPLIT-INVALID`. By `RFN-SPLIT-UNQREF-PAIR-R` and `RFN-DEREF-UNQREF-BOXPTR`, we can take out the unique reference  $*(a'.|\tau|)$ , which has the pure value  $(w, w_o)$ . By passing it to `self`, which satisfies the expected specification by the assumption, we get the postcondition  $\text{len } w_o = \text{len } w$  and acquire as the returned object a box pointer of the type box list  $\&_{\text{uniq}}^\alpha \tau$  and the pure value  $\text{zip } w \ w_o$ , which is stored to  $b.2$ . Now we have  $\text{len}(v_o :: w_o) = 1 + \text{len } w_o = 1 + \text{len } w = \text{len}(v :: w)$ . We still have a unique reference  $a'$  of the type  $\&_{\text{uniq}}^\alpha \tau$  and the pure value  $(v, v_o)$ , which is stored to  $b.1$ . By `WKN-PLNREF-TRIPLE-VRNT`,  $b$  finally has the type box list  $\&_{\text{uniq}}^\alpha \tau$  and the pure value  $(v, v_o) :: \text{zip } w \ w_o$ , which is equal to  $\text{zip}(v :: w) (v_o :: w_o)$ . So in this case the specification is satisfied, which concludes the proof.  $\square$

*Example 6.8* (Iterative Update of Vector by Increment). The following function `incr` increments the target integer of the unique reference  $a$  and then release  $a$ .

$$\text{incr} := \text{fn}(a) \{ a \leftarrow *a + 1; \& \}$$

It satisfies the following specification, which follows from the specification of [Example 6.3](#).

$$\text{incr} : \forall \alpha. \text{fn}(\&_{\text{uniq}}^\alpha \text{int}) \vdash \lambda \text{post}. ((n, n_o). n_o = n + 1 \Rightarrow \text{post}())$$

Combining this and `RFN-ITER-VEC-UNQREF`, we can derive the following specification.

$$\alpha \vdash \mathbf{a} : \&_{\text{uniq}}^\alpha \text{int} \vdash \text{iter}_{\text{vec}}^{\text{int}}(\mathbf{a}, \text{incr}) \vdash \lambda \text{post}. ((v, v_o). v_o = \text{map}(\lambda n. n+1) v \Rightarrow \text{post}())$$

## 6.5 Related Work

**RustHorn** As explained in [§1.3](#), the idea of using *prophecy* for modeling unique references comes from RustHorn ([Matsushita et al., 2020a,b](#)). They also presented the proof of *soundness and completeness* of the prophecy-based reduction from Rust programs to CHCs.

As a basis for the soundness and completeness proof, they made a core calculus of Rust dubbed Calculus of Ownership and Reference (COR), which models some basic features of Rust. In COR, a function consists of a set of very simple statements, each of which is associated with a program point and finally either jumps to another program point or returns from the function. Because of this design of COR, their translation of COR programs into CHCs is rather simple; a statement of each program point in a program is translated into one CHC (or two CHCs, for conditional branching).

As explained in § 3.1, their soundness and completeness proof is based on construction of bisimulation between the execution on the COR side and some deduction algorithm (called SLDC resolution) on the CHC side. On this construction, they associate each object in COR with a multiset of ownership tokens on memory cells of the heap memory and manage the safety invariant on the multiset of ownership tokens over all the objects at each program point, which has a flavor of separation logic but is performed rather in an ad-hoc way. A prophecy variable of each unique reference is represented as syntactic logic variables in CHCs, which can later be specialized in the deduction algorithm. Their proof depends heavily on syntactic structures, which requires careful design of the calculus COR, the reduction from COR to CHCs, and the deduction on CHCs. Because of the syntactic nature, it is unclear how we can extend their proof to other features and libraries of Rust, although they conjectured that their prophecy-based approach is applicable to a wide class of features and libraries.

Unlike RustHorn, the functional-correctness proof of our thesis is highly extensible, thanks to the power of the higher-order separation logic Iris and the lifetime logic. We support function types, concurrency and vector types, which are not supported by the formalization of RustHorn. Also, we can flexibly add new typing rules and libraries, whose correctness can be checked separately.

Still, for the logic model we can make using the refined type system we give the proof of *soundness* (with the adequacy theorems [Theorem 6.1](#) and [Theorem 4.1](#)) but not the proof of *completeness*, whereas RustHorn proves both soundness and completeness of the reduction. We believe that we can in theory prove completeness of our logic model by appropriately restricting the range of programs, in a way that enforces resolution of the prophecy variable of every unique reference. Still, in order to attain an extensible proof of completeness of the logic model, we probably need to design a new program logic.

**RustBelt** Our *semantic* approach to program verification comes from RustBelt ([Jung et al., 2018a](#)). RustBelt verified *safety* of the core type system and some basic libraries of Rust, but it did not verify functional correctness. In particular, RustBelt models unique reference as a kind of invariant saying that *some* object of the type is stored at the target address and does not track precise information about the object.

We extended their approach for verification of *functional correctness* of basic operations and libraries in Rust using the *prophecy*-based reduction of RustHorn. We refined semantic types with pure values parametrized over  $\pi$ , the assignment on prophecy variables, in order to handle prophecy information.

Still, the verification platform of this thesis does not support libraries with *interior mutability*. RustBelt verified *safety* of various libraries with interior mutability, including `Cell`, `RefCell`, `Mutex` and `RwLock`, using the power of invariants provided by Iris. We can rather easily incorporate their safety proofs into RustHornBelt. However, in order to verify more precise information about *functional correctness*, we need more machinery. This point is discussed more in depth in [Chapter 7](#).

**Semantic Soundness Proof** There are some existing studies on *semantic* soundness proof (i.e., soundness proof based on logical relations) for a substructural type system, including those by [Morrisett et al. \(2005\)](#); [Ahmed et al. \(2005\)](#); [Krishnaswami et al. \(2012\)](#); [Jung et al. \(2018a\)](#); [Dang et al. \(2020\)](#), although semantic approaches remain rather rare compared to *syntactic* approaches. However, we are not aware of any existing study that presents semantic soundness proof for a substructural type system that takes into account *functional correctness*, not only memory and thread safety.

**Other Studies on Verification of Rust Programs** There are a number of studies on automated low-level verification of Rust programs. [Toman et al. \(2015\)](#) developed a bounded model checker CRUST for automatically detecting violation of Rust’s ownership invariants by a Rust library with *unsafe code* within some bounds. The tool exhaustively generates tests on client use of a Rust library, imposing some upper bound on the number of library method calls, and passes the tests to SMBC, an existing SMT-based bounded model checker for the C programming language. They successfully re-discovered some real-world bugs of the standard Rust library that had been fixed at that time but been missed for a long time. [Lindner et al. \(2018\)](#) performed automated verification on Rust using symbolic execution in KLEE on the LLVM bitcode extracted from a Rust program. They particularly focused on the property that the program does not cause abortion, for example by out-of-bounds access or division by zero. [Baranowski et al. \(2018\)](#) performed automated verification on Rust programs using the SMACK verifier. In order to reuse SMACK’s existing verification tools for C, they modeled common Rust libraries as special C code tailored to SMACK.

There are also a number of existing studies on verification of Rust programs that exploit the guarantees by Rust’s ownership principle. [Ullrich \(2016\)](#) developed a tool for automatically translating Rust programs of some subset into the purely functional language of the Lean Theorem Prover and verified functional correctness of some Rust programs in Lean using that translation. The basic idea of his translation is to model a Rust function that inputs and consumes a unique reference (e.g., `Vec::push`) as a function that inputs the target value of the unique reference and outputs the new target value of the unique reference. The translation also supports one-to-one subdivision of unique references like `Vec::index_mut` through functional lenses. However, the translation still cannot handle various common usages of unique references, including *split* of a unique reference (e.g., `&mut (T, U)` into `&mut T` and `&mut U`), unlike the prophecy-based approach of RustHorn and RustHornBelt. [Hahn \(2016\)](#) and [Astrauskas et al. \(2019\)](#) performed semi-automated verification of Rust programs with annotations for verification on top of Viper ([Müller et al., 2016](#)), a verification platform that employs a separation logic with fractional permission. In particular, [Astrauskas et al. \(2019\)](#) supported one-to-one subdivision of unique references like `Vec::index_mut` using an assertion called a pledge, which is modeled in Viper’s separation logic using a magic wand. Also, Viper’s separation logic is much simpler than Iris and can possibly be more suited for automation. Still, their encoding does not support general use cases of unique references, including *split* of a unique reference, unlike RustHorn and RustHornBelt.

## Chapter 7

### Conclusion and Future Work

We proposed a novel extensible logical foundation to specify and verify functional correctness of Rust programs, using a prophecy-based clean logic model in the style of RustHorn and taking a semantic approach in Iris in the style of RustBelt, under the project name RustHornBelt. We designed a new type system with a logic model for Rust, which is highly flexible and supports various features that were not supported by RustHorn. As a basis for that, we also presented a new flexible platform of prophecy in Iris, which manages information about prophecy only in the ghost state and allows an operation that we call dependent resolution. We also presented a new technique for spending many logical steps at a time in Iris.

We believe that this work serves as a significant step for extensible and scalable verification of Rust programs, but we still have a lot of work to do in this direction.

**Mechanization in Coq** We plan to *mechanize* the results of this thesis in the Coq Proof Assistant in the near future. We already have a large size of Coq code that we can utilize. The separation logic Iris is mechanized in Coq with great support of tactics and proof modes. The lifetime logic is mechanized on top of the Coq library of Iris. Also, we can reuse some parts of the implementation of RustBelt. Still, since we handle information about *functional correctness* involving *prophecy* in RustHornBelt, we need to develop useful lemmas and tactics for verifying typing rules and libraries efficiently. In particular, it is great if we can verify involved libraries (e.g., `Vec` and `HashMap`) in a *scalable* way. We can possibly adopt some methods from a recently emerging study by Sammler et al. (2020b), which works on automated and verification of functional correctness of C programs with an extensible and reliable foundation built in Coq and Iris.

**Semi-Automated Verification of Rust Programs** Although this thesis presents a logical foundation to give clean logical models to basic operations and libraries of Rust by the technique of prophecy, we have yet to design and implement a *semi-automated* verification platform for Rust programs that can perform scalable verification of Rust programs using the clean models given by our foundation. We expect that we can achieve such a scalable verification platform for Rust on top of existing semi-automated verification platforms for functional programming languages such as F\* (Swamy et al., 2016) and Why3 (Filliâtre and Paskevich, 2013).

**Interior Mutability** In the verification platform of this thesis, we omitted support of libraries with *interior mutability* like `Cell` and `RefCell`. We can prove memory and thread *safety* of such libraries using atomic and non-atomic *invariants* asserting we have *some* object of the expect type inside the cell-like object, which is the proof

technique used by RustBelt.<sup>1</sup> However, if we further want to prove *functional correctness*, we need to track more precise information about the *value* of the inner object of the cell-like object.

With interior mutability, we need to handle almost general imperative programming, although we have solid memory and thread safety. The situation is similar for example to verification of OCaml programs with reference cells, where memory safety is ensured by the type system and the garbage collection.

In principle, we can precisely encode a Rust program using interior mutability by carrying around a *global* array which maps the id of each cell-like object to its inner value in the store-passing style, but this is not very *scalable* for involved programs. One idea is to use *region types* (Fluet et al., 2006) to separate out some groups of cell-like objects that we manipulate, but we have yet to tailor this technique to Rust.

**Concurrency** In the verification platform of this thesis, we supported only the structured concurrent execution  $e \parallel e'$  for concurrency.

To support more features of concurrency, we can use the proof techniques of RustBelt, which verified *safety* of various features for concurrency in Rust including threads, channels, mutexes `Mutex<T>` and read-write locks `RwLock<T>`. Still, in order to verify *functional correctness* of concurrent programs, we need other techniques.

One idea is to use *prophecy* also for concurrency. For example, when we create a channel to get a sender and a receiver, we can prophesy the list of values of the objects that will pass through the channel.<sup>2</sup> When the sender sends an object, it can resolve the head of the prophecy list and subdivide the prophecy list into its tail. When the receiver receives an object, we know the value of the object from the prophesied list. When we clone the sender, we split the prophecy list in a non-deterministic way. We believe that we can formalize this prophecy-based encoding of a channel using our formulation of prophecy (Chapter 3). Still, this encoding is not very precise. For example, in this encoding, we cannot detect dead locks or discuss interaction between multiple senders.

Another idea is to incorporate the *rely-guarantee* reasoning (Vafeiadis and Parkinson, 2007; Feng et al., 2007) for more fine-grained verification on concurrency, but we have yet to tailor this to RustHornBelt.

**Termination-Sensitive Verification** The verification platform of this thesis does not support *termination-sensitive* verification. Using time credits (Atkey, 2011; Mével et al., 2019), we can verify a specific upper bound on the execution time of a program, which can probably be incorporated into RustHornBelt. However, when a function takes an unbounded but finite number of physical steps (e.g., repeat a loop  $n$  times for a random natural number  $n$ ), it is hard to verify its termination in Iris, particularly because Iris uses *finite* step indexing. We may be able to overcome this situation by using recently emerging Transfinite Iris (Spies et al., 2020) instead of Iris for the underlying separation logic. However, as discussed in §4.3, whether we can use Transfinite Iris for RustHornBelt remains unclear because a number of good properties of Iris are lost in Transfinite Iris.

---

<sup>1</sup> To support types like `Cell` in the presence of concurrency, we need to introduce Rust's `Send` and `Sync` traits, which manage what kind of object can be passed between thread. RustBelt models these traits respectively as the condition that the ownership or sharing predicate of the type does not depend on the thread id.

<sup>2</sup> This idea is similar to *sequence prophecies* used by Jung et al. (2020b).

## References

- Martín Abadi and Leslie Lamport. 1988. The Existence of Refinement Mappings. In *Proceedings of the Third Annual Symposium on Logic in Computer Science (LICS '88), Edinburgh, Scotland, UK, July 5-8, 1988*. IEEE Computer Society, 165–175. <https://doi.org/10.1109/LICS.1988.5115>
- Martín Abadi and Leslie Lamport. 1991. The Existence of Refinement Mappings. *Theor. Comput. Sci.* 82, 2 (1991), 253–284. [https://doi.org/10.1016/0304-3975\(91\)90224-P](https://doi.org/10.1016/0304-3975(91)90224-P)
- Amal J. Ahmed, Matthew Fluet, and Greg Morrisett. 2005. A Step-Indexed Model of Substructural State. In *Proceedings of the 10th ACM SIGPLAN International Conference on Functional Programming, ICFP 2005, Tallinn, Estonia, September 26-28, 2005*, Olivier Danvy and Benjamin C. Pierce (Eds.). ACM, 78–91. <https://doi.org/10.1145/1086365.1086376>
- Hussain M. J. Almhori and David Evans. 2018. Fidelius Charm: Isolating Unsafe Rust Code. In *Proceedings of the Eighth ACM Conference on Data and Application Security and Privacy, CODASPY 2018, Tempe, AZ, USA, March 19-21, 2018*, Ziming Zhao, Gail-Joon Ahn, Ram Krishnan, and Gabriel Ghinita (Eds.). ACM, 248–255. <https://doi.org/10.1145/3176258.3176330>
- Brian Anderson, Lars Bergstrom, Manish Goregaokar, Josh Matthews, Keegan McAllister, Jack Moffitt, and Simon Sapin. 2016. Engineering the Servo Web Browser Engine Using Rust. In *Proceedings of the 38th International Conference on Software Engineering, ICSE 2016, Austin, TX, USA, May 14-22, 2016 - Companion Volume*, Laura K. Dillon, Willem Visser, and Laurie Williams (Eds.). ACM, 81–89. <https://doi.org/10.1145/2889160.2889229>
- Andrew W. Appel and David A. McAllester. 2001. An Indexed Model of Recursive Types for Foundational Proof-Carrying Code. *ACM Trans. Program. Lang. Syst.* 23, 5 (2001), 657–683. <https://doi.org/10.1145/504709.504712>
- 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 (2020), 136:1–136:27. <https://doi.org/10.1145/3428204>
- 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 (2019), 147:1–147:30. <https://doi.org/10.1145/3360573>
- Robert Atkey. 2011. Amortised Resource Analysis with Separation Logic. *Log. Methods Comput. Sci.* 7, 2 (2011). [https://doi.org/10.2168/LMCS-7\(2:17\)2011](https://doi.org/10.2168/LMCS-7(2:17)2011)
- Abhiram Balasubramanian, Marek S. Baranowski, Anton Burtsev, Aurojit Panda, Zvonimir Rakamaric, and Leonid Ryzhyk. 2017. System Programming in Rust: Beyond

- Safety. In *Proceedings of the 16th Workshop on Hot Topics in Operating Systems, HotOS 2017, Whistler, BC, Canada, May 8-10, 2017*, Alexandra Fedorova, Andrew Warfield, Ivan Beschastnikh, and Rachit Agarwal (Eds.). ACM, 156–161. <https://doi.org/10.1145/3102980.3103006>
- Marek S. Baranowski, Shaobo He, and Zvonimir Rakamaric. 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 (Lecture Notes in Computer Science, Vol. 11138)*, Shuvendu K. Lahiri and Chao Wang (Eds.). Springer, 528–535. [https://doi.org/10.1007/978-3-030-01090-4\\_32](https://doi.org/10.1007/978-3-030-01090-4_32)
- Ariel Ben-Yehuda. 2015. `std::thread::JoinGuard` (and `scoped`) are Unsound Because of Reference Cycles. <https://github.com/rust-lang/rust/issues/24292>
- Jean-Philippe Bernardy, Mathieu Boespflug, Ryan R. Newton, Simon Peyton Jones, and Arnaud Spiwack. 2018. Linear Haskell: Practical Linearity in a Higher-Order Polymorphic Language. *PACMPL* 2, POPL (2018), 5:1–5:29. <https://doi.org/10.1145/3158093>
- Lars Birkedal, Rasmus Ejlers Møgelberg, Jan Schwinghammer, and Kristian Støvring. 2011. First Steps in Synthetic Guarded Domain Theory: Step-Indexing in the Topos of Trees. In *Proceedings of the 26th Annual IEEE Symposium on Logic in Computer Science, LICS 2011, June 21-24, 2011, Toronto, Ontario, Canada*. IEEE Computer Society, 55–64. <https://doi.org/10.1109/LICS.2011.16>
- Nikolaj Bjørner, Arie Gurfinkel, Kenneth L. McMillan, and Andrey Rybalchenko. 2015. Horn Clause Solvers for Program Verification. In *Fields of Logic and Computation II - Essays Dedicated to Yuri Gurevich on the Occasion of His 75th Birthday (Lecture Notes in Computer Science, Vol. 9300)*, Lev D. Beklemishev, Andreas Blass, Nachum Dershowitz, Bernd Finkbeiner, and Wolfram Schulte (Eds.). Springer, 24–51. [https://doi.org/10.1007/978-3-319-23534-9\\_2](https://doi.org/10.1007/978-3-319-23534-9_2)
- Konrad Borowski. 2019. `PartialEq` Implementation for `RangeInclusive` is Unsound. <https://github.com/rust-lang/rust/issues/67194>
- Adrien Champion, Naoki Kobayashi, and Ryosuke Sato. 2018. HoIce: An ICE-Based Non-linear Horn Clause Solver. In *Programming Languages and Systems - 16th Asian Symposium, APLAS 2018, Wellington, New Zealand, December 2-6, 2018, Proceedings (Lecture Notes in Computer Science, Vol. 11275)*, Sukyoung Ryu (Ed.). Springer, 146–156. [https://doi.org/10.1007/978-3-030-02768-1\\_8](https://doi.org/10.1007/978-3-030-02768-1_8)
- David G. Clarke, John Potter, and James Noble. 1998. Ownership Types for Flexible Alias Protection. In *Proceedings of the 1998 ACM SIGPLAN Conference on Object-Oriented Programming Systems, Languages & Applications (OOPSLA '98), Vancouver, British Columbia, Canada, October 18-22, 1998*, Bjørn N. Freeman-Benson and Craig Chambers (Eds.). ACM, 48–64. <https://doi.org/10.1145/286936.286947>
- Byron Cook and Eric Koskinen. 2011. Making Prophecies with Decision Predicates. In *Proceedings of the 38th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL 2011, Austin, TX, USA, January 26-28, 2011*, Thomas Ball and Mooly Sagiv (Eds.). ACM, 399–410. <https://doi.org/10.1145/1926385.1926431>

- Coq Community. 2021. *The Coq Proof Assistant*. <https://coq.inria.fr/>
- Hoang-Hai Dang, Jacques-Henri Jourdan, Jan-Oliver Kaiser, and Derek Dreyer. 2020. RustBelt Meets Relaxed Memory. *Proc. ACM Program. Lang.* 4, POPL (2020), 34:1–34:29. <https://doi.org/10.1145/3371102>
- Paulo Emílio de Vilhena and François Pottier. 2021. A Separation Logic for Effect Handlers. *Proc. ACM Program. Lang.* 5, POPL (2021), 1–28. <https://doi.org/10.1145/3434314>
- Paulo Emílio de Vilhena, François Pottier, and Jacques-Henri Jourdan. 2020. Spy Game: Verifying a Local Generic Solver in Iris. *Proc. ACM Program. Lang.* 4, POPL (2020), 33:1–33:28. <https://doi.org/10.1145/3371101>
- Robert DeLine and Manuel Fähndrich. 2001. Enforcing High-Level Protocols in Low-Level Software. In *Proceedings of the 2001 ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI), Snowbird, Utah, USA, June 20-22, 2001*, Michael Burke and Mary Lou Soffa (Eds.). ACM, 59–69. <https://doi.org/10.1145/378795.378811>
- Yu Ding, Ran Duan, Long Li, Yueqiang Cheng, Yulong Zhang, Tanghui Chen, Tao Wei, and Huibo Wang. 2017. Rust SGX SDK: Towards Memory Safety in Intel SGX Enclave. In *Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security, CCS 2017, Dallas, TX, USA, October 30 - November 03, 2017*, Bhavani M. Thuraisingham, David Evans, Tal Malkin, and Dongyan Xu (Eds.). ACM, 2491–2493. <https://doi.org/10.1145/3133956.3138824>
- Dropbox. 2020. *Rewriting the Heart of Our Sync Engine - Dropbox*. <https://dropbox.tech/infrastructure/rewriting-the-heart-of-our-sync-engine>
- Paul Emmerich, Sebastian Voit, Georg Carle, Simon Ellmann, Fabian Bonk, Alex Egger, Esaú García Sánchez-Torija, Thomas Günzel, Sebastian Di Luzio, Alexandru Obada, and Maximilian Stadlmeier. 2019. The Case for Writing Network Drivers in High-Level Programming Languages. In *2019 ACM/IEEE Symposium on Architectures for Networking and Communications Systems, ANCS 2019, Cambridge, United Kingdom, September 24-25, 2019*. IEEE, 1–13. <https://doi.org/10.1109/ANCS.2019.8901892>
- Manuel Fähndrich and Robert DeLine. 2002. Adoption and Focus: Practical Linear Types for Imperative Programming. In *Proceedings of the 2002 ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI), Berlin, Germany, June 17-19, 2002*, Jens Knoop and Laurie J. Hendren (Eds.). ACM, 13–24. <https://doi.org/10.1145/512529.512532>
- Chris Fallin. 2020. Safe, Flexible Aliasing with Deferred Borrows. In *34th European Conference on Object-Oriented Programming, ECOOP 2020, November 15-17, 2020, Berlin, Germany (Virtual Conference) (LIPIcs, Vol. 166)*, Robert Hirschfeld and Tobias Pape (Eds.). Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 30:1–30:26. <https://doi.org/10.4230/LIPIcs.ECOOP.2020.30>
- Grigory Fedyukovich, Samuel J. Kaufman, and Rastislav Bodik. 2017. Sampling Invariants from Frequency Distributions. In *2017 Formal Methods in Computer Aided Design, FMCAD 2017, Vienna, Austria, October 2-6, 2017*, Daryl Stewart and Georg Weissenbacher (Eds.). IEEE, 100–107. <https://doi.org/10.23919/FMCAD.2017.8102247>

- Xinyu Feng, Rodrigo Ferreira, and Zhong Shao. 2007. On the Relationship Between Concurrent Separation Logic and Assume-Guarantee Reasoning. In *Programming Languages and Systems, 16th European Symposium on Programming, ESOP 2007, Held as Part of the Joint European Conferences on Theory and Practice of Software, ETAPS 2007, Braga, Portugal, March 24 - April 1, 2007, Proceedings (Lecture Notes in Computer Science, Vol. 4421)*, Rocco De Nicola (Ed.). Springer, 173–188. [https://doi.org/10.1007/978-3-540-71316-6\\_13](https://doi.org/10.1007/978-3-540-71316-6_13)
- Jean-Christophe Filliâtre and Andrei Paskevich. 2013. Why3 - Where Programs Meet Provers. In *Programming Languages and Systems - 22nd European Symposium on Programming, ESOP 2013, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2013, Rome, Italy, March 16-24, 2013. Proceedings (Lecture Notes in Computer Science, Vol. 7792)*, Matthias Felleisen and Philippa Gardner (Eds.). Springer, 125–128. [https://doi.org/10.1007/978-3-642-37036-6\\_8](https://doi.org/10.1007/978-3-642-37036-6_8)
- Matthew Fluet, Greg Morrisett, and Amal J. Ahmed. 2006. Linear Regions Are All You Need. In *Programming Languages and Systems, 15th European Symposium on Programming, ESOP 2006, Held as Part of the Joint European Conferences on Theory and Practice of Software, ETAPS 2006, Vienna, Austria, March 27-28, 2006, Proceedings (Lecture Notes in Computer Science, Vol. 3924)*, Peter Sestoft (Ed.). Springer, 7–21. [https://doi.org/10.1007/11693024\\_2](https://doi.org/10.1007/11693024_2)
- J. Nathan Foster, Michael B. Greenwald, Jonathan T. Moore, Benjamin C. Pierce, and Alan Schmitt. 2007. Combinators for Bidirectional Tree Transformations: A Linguistic Approach to the View-Update Problem. *ACM Trans. Program. Lang. Syst.* 29, 3 (2007), 17. <https://doi.org/10.1145/1232420.1232424>
- Dan Frumin, Robbert Krebbers, and Lars Birkedal. 2018. ReLoC: A Mechanised Relational Logic for Fine-Grained Concurrency. In *Proceedings of the 33rd Annual ACM/IEEE Symposium on Logic in Computer Science, LICS 2018, Oxford, UK, July 09-12, 2018*, Anuj Dawar and Erich Grädel (Eds.). ACM, 442–451. <https://doi.org/10.1145/3209108.3209174>
- Dan R. Ghica and Alex I. Smith. 2014. Bounded Linear Types in a Resource Semiring. In *Programming Languages and Systems - 23rd European Symposium on Programming, ESOP 2014, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2014, Grenoble, France, April 5-13, 2014, Proceedings (Lecture Notes in Computer Science, Vol. 8410)*, Zhong Shao (Ed.). Springer, 331–350. [https://doi.org/10.1007/978-3-642-54833-8\\_18](https://doi.org/10.1007/978-3-642-54833-8_18)
- Léon Gondelman, Simon Oddershede Gregersen, Abel Nieto, Amin Timany, and Lars Birkedal. 2021. Distributed Causal Memory: Modular Specification and Verification in Higher-Order Distributed Separation logic. *Proc. ACM Program. Lang.* 5, POPL (2021), 1–29. <https://doi.org/10.1145/3434323>
- Sergey Grebenshchikov, Nuno P. Lopes, Corneliu Popeea, and Andrey Rybalchenko. 2012. Synthesizing Software Verifiers from Proof Rules. In *ACM SIGPLAN Conference on Programming Language Design and Implementation, PLDI '12, Beijing, China - June 11 - 16, 2012*, Jan Vitek, Haibo Lin, and Frank Tip (Eds.). ACM, 405–416. <https://doi.org/10.1145/2254064.2254112>
- Arie Gurfinkel, Temesghen Kahsai, Anvesh Komuravelli, and Jorge A. Navas. 2015. The SeaHorn Verification Framework. In *Computer Aided Verification - 27th International Conference, CAV 2015, San Francisco, CA, USA, July 18-24, 2015, Proceedings, Part I*

- (*Lecture Notes in Computer Science, Vol. 9206*), Daniel Kroening and Corina S. Pasareanu (Eds.). Springer, 343–361. [https://doi.org/10.1007/978-3-319-21690-4\\_20](https://doi.org/10.1007/978-3-319-21690-4_20)
- Florian Hahn. 2016. Rust2Viper: Building a Static Verifier for Rust. Master Thesis. ETH Zürich. <https://www.research-collection.ethz.ch/handle/20.500.11850/155723>
- Philipp Haller and Martin Odersky. 2010. Capabilities for Uniqueness and Borrowing. In *ECOOP 2010 - Object-Oriented Programming, 24th European Conference, Maribor, Slovenia, June 21-25, 2010. Proceedings (Lecture Notes in Computer Science, Vol. 6183)*, Theo D’Hondt (Ed.). Springer, 354–378. [https://doi.org/10.1007/978-3-642-14107-2\\_17](https://doi.org/10.1007/978-3-642-14107-2_17)
- Timothy L. Harris, Keir Fraser, and Ian A. Pratt. 2002. A Practical Multi-word Compare-and-Swap Operation. In *Distributed Computing, 16th International Conference, DISC 2002, Toulouse, France, October 28-30, 2002 Proceedings (Lecture Notes in Computer Science, Vol. 2508)*, Dahlia Malkhi (Ed.). Springer, 265–279. [https://doi.org/10.1007/3-540-36108-1\\_18](https://doi.org/10.1007/3-540-36108-1_18)
- Jonas Kastberg Hinrichsen, Jesper Bengtson, and Robbert Krebbers. 2020. Actris: Session-Type Based Reasoning in Separation Logic. *Proc. ACM Program. Lang.* 4, POPL (2020), 6:1–6:30. <https://doi.org/10.1145/3371074>
- Hossein Hojjat and Philipp Rümmer. 2018. The ELARICA Horn Solver. In *2018 Formal Methods in Computer Aided Design, FMCAD 2018, Austin, TX, USA, October 30 - November 2, 2018*, Nikolaj Bjørner and Arie Gurfinkel (Eds.). IEEE, 1–7. <https://doi.org/10.23919/FMCAD.2018.8603013>
- Koen Jacobs, Amin Timany, and Dominique Devriese. 2021. Fully Abstract from Static to Gradual. *Proc. ACM Program. Lang.* 5, POPL (2021), 1–30. <https://doi.org/10.1145/3434288>
- Thomas Bracht Laumann Jespersen, Philip Munksgaard, and Ken Friis Larsen. 2015. Session Types for Rust. In *Proceedings of the 11th ACM SIGPLAN Workshop on Generic Programming, WGP@ICFP 2015, Vancouver, BC, Canada, August 30, 2015*, Patrick Bahr and Sebastian Erdweg (Eds.). ACM, 13–22. <https://doi.org/10.1145/2808098.2808100>
- Trevor Jim, J. Gregory Morrisett, Dan Grossman, Michael W. Hicks, James Cheney, and Yanling Wang. 2002. Cyclone: A Safe Dialect of C. In *Proceedings of the General Track: 2002 USENIX Annual Technical Conference, June 10-15, 2002, Monterey, California, USA*, Carla Schlatter Ellis (Ed.). USENIX, 275–288. <http://www.usenix.org/publications/library/proceedings/usenix02/jim.html>
- Ralf Jung. 2017. *MutexGuard<Cell<i32>> Must not be Sync*. <https://github.com/rust-lang/rust/issues/41622>
- Ralf Jung. 2020. *Understanding and Evolving the Rust Programming Language*. Ph. D. Dissertation. Saarland University, Saarbrücken, Germany. <https://publikationen.sulb.uni-saarland.de/handle/20.500.11880/29647>
- Ralf Jung, Hoang-Hai Dang, Jeehoon Kang, and Derek Dreyer. 2020a. Stacked Borrows: an Aliasing Model for Rust. *Proc. ACM Program. Lang.* 4, POPL (2020), 41:1–41:32. <https://doi.org/10.1145/3371109>

- Ralf Jung, Jacques-Henri Jourdan, Robbert Krebbers, and Derek Dreyer. 2018a. RustBelt: Securing the Foundations of the Rust Programming Language. *Proc. ACM Program. Lang.* 2, POPL (2018), 66:1–66:34. <https://doi.org/10.1145/3158154>
- Ralf Jung, Jacques-Henri Jourdan, Robbert Krebbers, and Derek Dreyer. 2018b. RustBelt: Securing the Foundations of the Rust Programming Language – Technical Appendix. <https://plv.mpi-sws.org/rustbelt/pop18/appendix.pdf>
- Ralf Jung, Robbert Krebbers, Jacques-Henri Jourdan, Ales Bizjak, Lars Birkedal, and Derek Dreyer. 2018c. Iris from the Ground Up: A Modular Foundation for Higher-Order Concurrent Separation Logic. *J. Funct. Program.* 28 (2018), e20. <https://doi.org/10.1017/S0956796818000151>
- Ralf Jung, Rodolphe Lepigre, Gaurav Parthasarathy, Marianna Rapoport, Amin Timany, Derek Dreyer, and Bart Jacobs. 2020b. The Future is Ours: Prophecy Variables in Separation Logic. *Proc. ACM Program. Lang.* 4, POPL (2020), 45:1–45:32. <https://doi.org/10.1145/3371113>
- Ralf Jung, David Swasey, Filip Sieczkowski, Kasper Svendsen, Aaron Turon, Lars Birkedal, and Derek Dreyer. 2015. Iris: Monoids and Invariants as an Orthogonal Basis for Concurrent Reasoning. In *Proceedings of the 42nd Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL 2015, Mumbai, India, January 15-17, 2015*, Sriram K. Rajamani and David Walker (Eds.). ACM, 637–650. <https://doi.org/10.1145/2676726.2676980>
- Jan-Oliver Kaiser, Hoang-Hai Dang, Derek Dreyer, Ori Lahav, and Viktor Vafeiadis. 2017. Strong Logic for Weak Memory: Reasoning About Release-Acquire Consistency in Iris. In *31st European Conference on Object-Oriented Programming, ECOOP 2017, June 19-23, 2017, Barcelona, Spain (LIPIcs, Vol. 74)*, Peter Müller (Ed.). Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 17:1–17:29. <https://doi.org/10.4230/LIPIcs.ECOOP.2017.17>
- Steve Klabnik, Carol Nichols, and Rust Community. 2018. *The Rust Programming Language*. <https://doc.rust-lang.org/book/>
- Anvesh Komuravelli, Arie Gurfinkel, and Sagar Chaki. 2014. SMT-based Model Checking for Recursive Programs. In *Computer Aided Verification - 26th International Conference, CAV 2014, Held as Part of the Vienna Summer of Logic, VSL 2014, Vienna, Austria, July 18-22, 2014. Proceedings (Lecture Notes in Computer Science, Vol. 8559)*, Armin Biere and Roderick Bloem (Eds.). Springer, 17–34. [https://doi.org/10.1007/978-3-319-08867-9\\_2](https://doi.org/10.1007/978-3-319-08867-9_2)
- Robert A. Kowalski. 1974. Predicate Logic as Programming Language. In *Information Processing, Proceedings of the 6th IFIP Congress 1974, Stockholm, Sweden, August 5-10, 1974*, Jack L. Rosenfeld (Ed.). North-Holland, 569–574.
- Robbert Krebbers, Jacques-Henri Jourdan, Ralf Jung, Joseph Tassarotti, Jan-Oliver Kaiser, Amin Timany, Arthur Charguéraud, and Derek Dreyer. 2018. MoSeL: A General, Extensible Modal Framework for Interactive Proofs in Separation Logic. *Proc. ACM Program. Lang.* 2, ICFP (2018), 77:1–77:30. <https://doi.org/10.1145/3236772>
- Robbert Krebbers, Amin Timany, and Lars Birkedal. 2017. Interactive Proofs in Higher-Order Concurrent Separation Logic. In *Proceedings of the 44th ACM SIGPLAN Symposium on Principles of Programming Languages, POPL 2017, Paris, France, Jan-*

- uary 18-20, 2017, Giuseppe Castagna and Andrew D. Gordon (Eds.). ACM, 205–217. <https://doi.org/10.1145/3009837>
- Neelakantan R. Krishnaswami, Aaron Turon, Derek Dreyer, and Deepak Garg. 2012. Superficially Substructural Types. In *ACM SIGPLAN International Conference on Functional Programming, ICFP'12, Copenhagen, Denmark, September 9-15, 2012*, Peter Thiemann and Robby Bruce Findler (Eds.). ACM, 41–54. <https://doi.org/10.1145/2364527.2364536>
- Benjamin Lamowski, Carsten Weinhold, Adam Lackorzynski, and Hermann Härtig. 2017. Sandcrust: Automatic Sandboxing of Unsafe Components in Rust. In *Proceedings of the 9th Workshop on Programming Languages and Operating Systems, Shanghai, China, October 28, 2017*, Julia Lawall (Ed.). ACM, 51–57. <https://doi.org/10.1145/3144555.3144562>
- Leslie Lamport and Stephan Merz. 2017. Auxiliary Variables in TLA+. *CoRR* abs/1703.05121 (2017). arXiv:1703.05121 <http://arxiv.org/abs/1703.05121>
- Amit Levy, Bradford Campbell, Branden Ghena, Pat Pannuto, Prabal Dutta, and Philip Levis. 2017. The Case for Writing a Kernel in Rust. In *Proceedings of the 8th Asia-Pacific Workshop on Systems, Mumbai, India, September 2, 2017*. ACM, 1:1–1:7. <https://doi.org/10.1145/3124680.3124717>
- Amit A. Levy, Michael P. Andersen, Bradford Campbell, David E. Culler, Prabal Dutta, Branden Ghena, Philip Levis, and Pat Pannuto. 2015. Ownership is theft: experiences building an embedded OS in Rust. In *Proceedings of the 8th Workshop on Programming Languages and Operating Systems, PLOS 2015, Monterey, California, USA, October 4, 2015*, Shan Lu (Ed.). ACM, 21–26. <https://doi.org/10.1145/2818302.2818306>
- Hongjin Liang and Xinyu Feng. 2013. Modular Verification of Linearizability with Non-Fixed Linearization Points. In *ACM SIGPLAN Conference on Programming Language Design and Implementation, PLDI '13, Seattle, WA, USA, June 16-19, 2013*, Hans-Juergen Boehm and Cormac Flanagan (Eds.). ACM, 459–470. <https://doi.org/10.1145/2491956.2462189>
- Marcus Lindner, Jorge Aparicius, and Per Lindgren. 2018. No Panic! Verification of Rust Programs by Symbolic Execution. In *16th IEEE International Conference on Industrial Informatics, INDIN 2018, Porto, Portugal, July 18-20, 2018*. IEEE, 108–114. <https://doi.org/10.1109/INDIN.2018.8471992>
- Nicholas D. Matsakis. 2018. *An Alias-based Formulation of the Borrow Checker*. <http://smallcultfollowing.com/babysteps/blog/2018/04/27/an-alias-based-formulation-of-the-borrow-checker/>
- Nicholas D. Matsakis. 2020. *Polonius: Either Borrower or Lender Be, but Responsibly*. [https://www.youtube.com/watch?v=\\_agDeiWek8w](https://www.youtube.com/watch?v=_agDeiWek8w)
- Nicholas D. Matsakis and Felix S. Klock, II. 2014. The Rust Language. In *Proceedings of the 2014 ACM SIGAda annual conference on High integrity language technology, HILT 2014, Portland, Oregon, USA, October 18-21, 2014*, Michael Feldman and S. Tucker Taft (Eds.). ACM, 103–104. <https://doi.org/10.1145/2663171.2663188>
- Yusuke Matsushita. 2019. CHC-based Program Verification Exploiting Ownership Types (所有権型を利用したCHCベースのプログラム検証). Senior Thesis. University of Tokyo. <http://www.kb.is.s.u-tokyo.ac.jp/~yskm24t/papers/senior-thesis.pdf>

- Yusuke Matsushita, Takeshi Tsukada, and Naoki Kobayashi. 2020a. RustHorn: CHC-based Verification for Rust Programs. In *Programming Languages and Systems - 29th European Symposium on Programming, ESOP 2020, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2020, Dublin, Ireland, April 25-30, 2020, Proceedings (Lecture Notes in Computer Science, Vol. 12075)*, Peter Müller (Ed.). Springer, 484–514. [https://doi.org/10.1007/978-3-030-44914-8\\_18](https://doi.org/10.1007/978-3-030-44914-8_18)
- Yusuke Matsushita, Takeshi Tsukada, and Naoki Kobayashi. 2020b. RustHorn: CHC-based Verification for Rust Programs (full version). *CoRR* abs/2002.09002 (2020). arXiv:2002.09002 <https://arxiv.org/abs/2002.09002>
- Karl Mazurak, Jianzhou Zhao, and Steve Zdancewic. 2010. Lightweight Linear Types in System F<sup>o</sup>. In *Proceedings of TLDI 2010: 2010 ACM SIGPLAN International Workshop on Types in Languages Design and Implementation, Madrid, Spain, January 23, 2010*, Andrew Kennedy and Nick Benton (Eds.). ACM, 77–88. <https://doi.org/10.1145/1708016.1708027>
- Glen Mével, Jacques-Henri Jourdan, and François Pottier. 2019. Time Credits and Time Receipts in Iris. In *Programming Languages and Systems - 28th European Symposium on Programming, ESOP 2019, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2019, Prague, Czech Republic, April 6-11, 2019, Proceedings (Lecture Notes in Computer Science, Vol. 11423)*, Luís Caires (Ed.). Springer, 3–29. [https://doi.org/10.1007/978-3-030-17184-1\\_1](https://doi.org/10.1007/978-3-030-17184-1_1)
- Matt Miller. 2019. *Trends, Challenge, and Shifts in Software Vulnerability Mitigation*. <https://www.youtube.com/watch?v=PjbGojJnBZQ>
- Kai Mindermann, Philipp Keck, and Stefan Wagner. 2018. How Usable Are Rust Cryptography APIs?. In *2018 IEEE International Conference on Software Quality, Reliability and Security, QRS 2018, Lisbon, Portugal, July 16-20, 2018*. IEEE, 143–154. <https://doi.org/10.1109/QRS.2018.00028>
- J. Garrett Morris. 2016. The Best of Both Worlds: Linear Functional Programming Without Compromise. In *Proceedings of the 21st ACM SIGPLAN International Conference on Functional Programming, ICFP 2016, Nara, Japan, September 18-22, 2016*, Jacques Garrigue, Gabriele Keller, and Eijiro Sumii (Eds.). ACM, 448–461. <https://doi.org/10.1145/2951913.2951925>
- Greg Morrisett, Amal J. Ahmed, and Matthew Fluet. 2005. L<sup>3</sup>: A Linear Language with Locations. In *Typed Lambda Calculi and Applications, 7th International Conference, TLCA 2005, Nara, Japan, April 21-23, 2005, Proceedings (Lecture Notes in Computer Science, Vol. 3461)*, Pawel Urzyczyn (Ed.). Springer, 293–307. [https://doi.org/10.1007/11417170\\_22](https://doi.org/10.1007/11417170_22)
- Mozilla. 2021. *Rust language — Mozilla Research*. <https://research.mozilla.org/rust/>
- Peter Müller, Malte Schwerhoff, and Alexander J. Summers. 2016. Viper: A Verification Infrastructure for Permission-Based Reasoning. In *Verification, Model Checking, and Abstract Interpretation - 17th International Conference, VMCAI 2016, St. Petersburg, FL, USA, January 17-19, 2016. Proceedings (Lecture Notes in Computer Science, Vol. 9583)*, Barbara Jobstmann and K. Rustan M. Leino (Eds.). Springer, 41–62. [https://doi.org/10.1007/978-3-662-49122-5\\_2](https://doi.org/10.1007/978-3-662-49122-5_2)
- npm. 2019. *Rust Case Study: Community Makes Rust an Easy Choice for npm*. <https://www.rust-lang.org/static/pdfs/Rust-npm-Whitepaper.pdf>

- Peter W. O’Hearn. 2007. Resources, Concurrency, and Local Reasoning. *Theor. Comput. Sci.* 375, 1-3 (2007), 271–307. <https://doi.org/10.1016/j.tcs.2006.12.035>
- Peter W. O’Hearn, John C. Reynolds, and Hongseok Yang. 2001. Local Reasoning about Programs that Alter Data Structures. In *Computer Science Logic, 15th International Workshop, CSL 2001. 10th Annual Conference of the EACSL, Paris, France, September 10-13, 2001, Proceedings (Lecture Notes in Computer Science, Vol. 2142)*, Laurent Fribourg (Ed.). Springer, 1–19. [https://doi.org/10.1007/3-540-44802-0\\_1](https://doi.org/10.1007/3-540-44802-0_1)
- Boqin Qin, Yilun Chen, Zeming Yu, Linhai Song, and Yiying Zhang. 2020. Understanding Memory and Thread Safety Practices and Issues in real-World Rust Programs. In *Proceedings of the 41st ACM SIGPLAN International Conference on Programming Language Design and Implementation, PLDI 2020, London, UK, June 15-20, 2020*, Alastair F. Donaldson and Emrina Torlak (Eds.). ACM, 763–779. <https://doi.org/10.1145/3385412.3386036>
- Azalea Raad, John Wickerson, Gil Neiger, and Viktor Vafeiadis. 2020. Persistency Semantics of the Intel-x86 Architecture. *Proc. ACM Program. Lang.* 4, POPL (2020), 11:1–11:31. <https://doi.org/10.1145/3371079>
- Eric Reed. 2015. Patina: A Formalization of the Rust Programming Language. Master Thesis. University of Washington. <https://dada.cs.washington.edu/research/tr/2015/03/UW-CSE-15-03-02.pdf>
- John C. Reynolds. 2002. Separation Logic: A Logic for Shared Mutable Data Structures. In *17th IEEE Symposium on Logic in Computer Science (LICS 2002), 22-25 July 2002, Copenhagen, Denmark, Proceedings*. IEEE Computer Society, 55–74. <https://doi.org/10.1109/LICS.2002.1029817>
- Rust Community. 2020. *2094-nll — The Rust RFC Book*. <https://rust-lang.github.io/rfcs/2094-nll.html>
- Rust Community. 2021a. *Rust Programming Language*. <https://www.rust-lang.org/>
- Rust Community. 2021b. *Sponsors — Rust Programming Language*. <https://www.rust-lang.org/sponsors>
- Rust Community. 2021c. *rust-lang/polonius: Defines the Rust Borrow Checker*. <https://github.com/rust-lang/polonius>
- Michael Sammler, Deepak Garg, Derek Dreyer, and Tadeusz Litak. 2020a. The High-Level Benefits of Low-Level Sandboxing. *Proc. ACM Program. Lang.* 4, POPL (2020), 32:1–32:32. <https://doi.org/10.1145/3371100>
- Michael Sammler, Rodolphe Lepigre, Robbert Krebbers, Kayvan Memarian, Derek Dreyer, and Deepak Garg. 2020b. RefinedC: A Foundational Refinement Type System for C Based on Separation Logic Programming. (2020). <https://plv.mpi-sws.org/refinedc/paper.pdf>
- Ali Sezgin, Serdar Tasiran, and Shaz Qadeer. 2010. Tressa: Claiming the Future. In *Verified Software: Theories, Tools, Experiments, Third International Conference, VSTTE 2010, Edinburgh, UK, August 16-19, 2010. Proceedings (Lecture Notes in Computer Science, Vol. 6217)*, Gary T. Leavens, Peter W. O’Hearn, and Sriram K. Rajamani (Eds.). Springer, 25–39. [https://doi.org/10.1007/978-3-642-15057-9\\_2](https://doi.org/10.1007/978-3-642-15057-9_2)

- Simon Spies, Lennard Gäher, Daniel Gratzer, Joseph Tassarotti, Robbert Krebbers, Derek Dreyer, and Lars Birkedal. 2020. Transfinite Iris: Resolving an Existential Dilemma of Step-Indexed Separation Logic. (2020). <https://iris-project.org/pdfs/2020-transfinite-iris-submission.pdf>
- Nikhil Swamy, Catalin Hritcu, Chantal Keller, Aseem Rastogi, Antoine Delignat-Lavaud, Simon Forest, Karthikeyan Bhargavan, Cédric Fournet, Pierre-Yves Strub, Markulf Kohlweiss, Jean Karim Zinzindohoue, and Santiago Zanella Béguelin. 2016. Dependent Types and Multi-Monadic Effects in F\*. In *Proceedings of the 43rd Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL 2016, St. Petersburg, FL, USA, January 20 - 22, 2016*, Rastislav Bodík and Rupak Majumdar (Eds.). ACM, 256–270. <https://doi.org/10.1145/2837614.2837655>
- Amin Timany, Léo Stefanescu, Morten Krogh-Jespersen, and Lars Birkedal. 2018. A Logical Relation for Monadic Encapsulation of State: Proving Contextual Equivalences in the Presence of runST. *Proc. ACM Program. Lang.* 2, POPL (2018), 64:1–64:28. <https://doi.org/10.1145/3158152>
- John Toman, Stuart Pernsteiner, and Emina Torlak. 2015. CRUST: A Bounded Verifier for Rust. In *30th IEEE/ACM International Conference on Automated Software Engineering, ASE 2015, Lincoln, NE, USA, November 9-13, 2015*, Myra B. Cohen, Lars Grunske, and Michael Whalen (Eds.). IEEE Computer Society, 75–80. <https://doi.org/10.1109/ASE.2015.77>
- Jesse A. Tov and Riccardo Pucella. 2011. Practical Affine Types. In *Proceedings of the 38th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL 2011, Austin, TX, USA, January 26-28, 2011*, Thomas Ball and Mooly Sagiv (Eds.). ACM, 447–458. <https://doi.org/10.1145/1926385.1926436>
- Aaron Joseph Turon, Jacob Thamsborg, Amal Ahmed, Lars Birkedal, and Derek Dreyer. 2013. Logical Relations for Fine-Grained Concurrency. In *The 40th Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL '13, Rome, Italy - January 23 - 25, 2013*, Roberto Giacobazzi and Radhia Cousot (Eds.). ACM, 343–356. <https://doi.org/10.1145/2429069.2429111>
- Sebastian Ullrich. 2016. Simple Verification of Rust Programs via Functional Purification. Master Thesis. Karlsruhe Institute of Technology. <https://pp.ipd.kit.edu/uploads/publikationen/ullrich16masterarbeit.pdf>
- Viktor Vafeiadis. 2008. *Modular Fine-Grained Concurrency Verification*. Ph.D. Dissertation. University of Cambridge, UK. <http://ethos.bl.uk/OrderDetails.do?uin=uk.bl.ethos.612221>
- Viktor Vafeiadis and Matthew J. Parkinson. 2007. A Marriage of Rely/Guarantee and Separation Logic. In *CONCUR 2007 - Concurrency Theory, 18th International Conference, CONCUR 2007, Lisbon, Portugal, September 3-8, 2007, Proceedings (Lecture Notes in Computer Science, Vol. 4703)*, Luís Caires and Vasco Thudichum Vasconcelos (Eds.). Springer, 256–271. [https://doi.org/10.1007/978-3-540-74407-8\\_18](https://doi.org/10.1007/978-3-540-74407-8_18)
- Philip Wadler. 1990. Linear Types Can Change the World!. In *Programming concepts and methods: Proceedings of the IFIP Working Group 2.2, 2.3 Working Conference on Programming Concepts and Methods, Sea of Galilee, Israel, 2-5 April, 1990*, Manfred Broy (Ed.). North-Holland, 561.

- Aaron Weiss, Daniel Patterson, Nicholas D. Matsakis, and Amal Ahmed. 2019. Oxide: The Essence of Rust. *CoRR* abs/1903.00982 (2019). arXiv:1903.00982 <http://arxiv.org/abs/1903.00982>
- Andrew K. Wright and Matthias Felleisen. 1994. A Syntactic Approach to Type Soundness. *Inf. Comput.* 115, 1 (1994), 38–94. <https://doi.org/10.1006/inco.1994.1093>
- Zipeng Zhang, Xinyu Feng, Ming Fu, Zhong Shao, and Yong Li. 2012. A Structural Approach to Prophecy Variables. In *Theory and Applications of Models of Computation - 9th Annual Conference, TAMC 2012, Beijing, China, May 16-21, 2012. Proceedings (Lecture Notes in Computer Science, Vol. 7287)*, Manindra Agrawal, S. Barry Cooper, and Angsheng Li (Eds.). Springer, 61–71. [https://doi.org/10.1007/978-3-642-29952-0\\_12](https://doi.org/10.1007/978-3-642-29952-0_12)
- Yoav Zibin, Alex Potanin, Paley Li, Mahmood Ali, and Michael D. Ernst. 2010. Ownership and Immutability in Generic Java. In *Proceedings of the 25th Annual ACM SIGPLAN Conference on Object-Oriented Programming, Systems, Languages, and Applications, OOPSLA 2010, October 17-21, 2010, Reno/Tahoe, Nevada, USA*, William R. Cook, Siobhán Clarke, and Martin C. Rinard (Eds.). ACM, 598–617. <https://doi.org/10.1145/1869459.1869509>