

Patina: A Formalization of the Rust Programming Language

Eric Reed
University of Washington

February 2015

Abstract

Rust is a new systems language that uses some advanced type system features, specifically affine types and regions, to statically guarantee memory safety and eliminate the need for a garbage collector. While each individual addition to the type system is well understood in isolation and are known to be sound, the combined system is not known to be sound. Furthermore, Rust uses a novel checking scheme for its regions, known as the Borrow Checker, that is not known to be correct. Since Rust's goal is to be a safer alternative to C/C++, we should ensure that this safety scheme actually works. We present a formal semantics that captures the key features relevant to memory safety, unique pointers and borrowed references, specifies how they guarantee memory safety, and describes the operation of the Borrow Checker. We use this model to prove the soundness of some core operations and justify the conjecture that the model, as a whole, is sound. Additionally, our model provides a syntactic version of the Borrow Checker, which may be more understandable than the non-syntactic version in Rust.

Overview

The Rust Programming Language [1] makes strong claims about ensuring memory safety without garbage collection. We would like to prove that those claims are true. To that end, we use a small model of Rust, called Patina, that characterizes the key features of Rust under the most suspicion, namely its memory management.

Rust memory management has two goals: no memory should ever become unreachable without being deallocated and no uninitialized memory should ever be read. Rust tries to achieve this by maintaining two invariants: all memory has a unique owner responsible for deallocating it and no memory is ever simultaneously aliased and mutable. These invariants simplify the situation enough so that Rust only needs to track which L-values are uninitialized and which have been borrowed. Ownership tracking of heap memory is managed by unique pointers. Safe aliasing is managed by borrowed references.

The Patina model has three layers. The innermost layer deals with L-values and ensures no uninitialized data is ever used, which includes ensuring that uninitialized pointers and null pointers are never dereferenced. It also ensures that compile-time initialization information correctly models the runtime memory. The middle layer deals with R-values. It ensures that any L-values used do not violate initialization or borrow restrictions. It also ensures that borrowed references are not created unless they are safe. The outer layer deals with statements. It mostly just chains together the guarantees of the L-value and R-value layers. However, it is also responsible for ensuring no leaks would be created when deallocating memory.

Unique Pointers in Rust

In Rust, unique pointers are the owners of heap memory. Heap memory is allocated when a new unique pointer is created. Heap memory is freed when a unique pointer falls out of scope.

```
{
  let x: ~int; // A stack variable containing a unique pointer to an integer
  x = ~3; // Allocates a heap integer, initializes it with 3, and stores the pointer in x
  // The heap memory owned by x is freed when it falls out of scope
}
```

To avoid double frees, unique pointers must remain unique. Thus, when a unique pointer would be copied the original must be rendered unusable. That is, they are moved rather than copied.

```
{
  let x: ~int = ~3;
  let y: ~int = x; // x is moved into y. x is no longer usable.
  let z: ~int = x; // ERROR!
}
```

However, these deinitialized paths can be reinitialized.

```
{
  let x: ~int = ~3;
  let y: ~int = x; // x is now deinitialized
  x = ~1; // x is initialized again
}
```

Unique pointers can also be freed if they would be leaked by assignment.

```
{
  let x: ~~int = ~~3;
  *x = ~1; // The ~3 that *x points too would be leaked here. It is freed instead.
}
```

This is not a dynamic check. The compiler detects when heap memory should be freed and inserts the necessary code.

Unique Pointers in Patina

When unique pointers fall out of scope in Rust, the compiler inserts code at the end of the block that will actually free the allocated heap memory. This code will traverse the entire owned data structure and free memory without leaking anything. In the case of recursive data structures, this code will be a recursive function. Rather than model such complex behavior as a primitive in Patina, we will use only a shallow free statement. Aside from making the model simpler, this also provides us with a means of checking that the code Rust inserts to free memory is correct, i.e. by encoding it in terms of Patina's shallow frees. The only way to deallocate a unique pointer in Patina is with these explicit frees.

```
{
  let x: ~~int = ~~3;
  //free(x); // ERROR! would leak *x
  // The proper way
  free(*x);
  free(x);
}
```

However, the free statement is only valid for initialized pointers, which prevents double frees.

```
{
  let x: ~int = ~3;
  free(x); // *x is now deallocated. x is uninitialized
  free(x); // ERROR! x is not initialized
}
```

Finally, the free statement requires that the pointer is unaliased, preventing dangling pointers.

```
{
  let x: ~int = ~3;
  // Create an immutable borrowed reference to *x with &
  // Immutable borrowed references to some type A are denoted by &A
  let y: &int = &*x;
  free(x); // ERROR! would make y a dangling pointer
}
```

Borrowed References

Often, particularly in functions, we want to use some memory, but do not intend to free it. That is, we want to borrow access to memory owned by something else. This is fine for copyable values, like integers, but is tedious for noncopyable values such as unique pointers. To avoid freeing a unique pointer argument at the end of the function body, it must be returned from the function. Effectively, the programmer must thread noncopyable values through their function calls.

```
fn foo(x: ~int) -> (int, ~int) {
  return (*x + 1, x)
}

{
  let x1: ~int = ~1;
  let (y,x2): (int, ~int) = foo(x);
  // y = 2, x1 is no longer valid, x2 = ~1
}
```

Borrowed references allow programmers to have temporary access to a value, but they do not confer ownership so this tedium is not necessary.

```
fn foo(x: &int) -> int {
  return (*x + 1)
}

{
  let x: ~int = ~1;
  let y: int = foo(&*x); // Rust lets us write foo(x) here and will insert the &* for us
  // x = ~1, y = 2
}
```

However, unrestricted borrowed references would open a hole in our memory safety by enabling aliased, mutable memory.

```
{
  let x: ~int = ~1;
  let y: &~int = &x; // x and *y are the same memory location
  let z: ~int = x; // x is no longer usable - effectively x is uninitialized memory
  let w: int = **y; // ERROR! *y is uninitialized now, so **y dereferences a null pointer!
}

{
  let x: &int;
  {
```

```

    let y: int = 3;
    y = &x;
} // x would be freed here
let z: int = *x // ERROR *y would point to unallocated memory
}

{
let x: Option<int> = Some(3);
let y: &int = match x {
    Some(ref z) => z,
    None => fail!(), // Impossible
}
x = None // payload is now uninitialized
let z: int = *y; // ERROR! *y points to uninitialized memory
}

```

Aliased, mutable memory can be avoided in two ways: forbidding mutability or forbidding aliasing. Correspondingly, Rust has two kinds of borrowed references. Immutable borrows, denoted `&`, are aliasable, but require all memory reachable through them to be immutable for the duration of the borrow.

```

{
let x: ~int = ~1;
let y: &~int = &x; // this immutable borrow prevents x from being changed
let z: ~int = x; // Rust will not let this typecheck because it would change x
}

```

Conversely, mutable borrows, denoted `&mut`, are mutable, but require all memory reachable through them to be *only* accessible through the borrow for the duration. That is, they guarantee unique access.

```

{
let x: ~int = ~1;
let y: &mut ~int = &mut x; // this mutable borrow prevents others from using x
let z: ~int = x; // Rust will not let this typecheck because it would access x
}

```

Finally, both kinds of borrowed reference require that the referent outlive the reference, which prevents dangling references.

```

fn foo() -> &int {
    let x: int = 1;
    return &x; // Rust will not let this typecheck because it would be a dangling pointer
}

```

Types

Values in Patina can be described by the following types:

Lifetime	ℓ
Type Variable	X
Qualifier	$q ::= \text{imm} \mid \text{mut}$
Type	$\tau ::= \text{int} \mid \sim \tau \mid \& \ell q \tau \mid \langle \tau_i \rangle_{i \in \{1 \dots n\}} \mid [\tau_i]_{i \in \{1 \dots n\}} \mid \mu X. \tau \mid X$

$\sim \tau$ is a unique pointer to a τ . $\& \ell q \tau$ is a borrowed reference to a τ providing mutability guarantee q for lifetime ℓ , which is Rust's name for regions associated with stack frames. $\langle \tau_i \rangle_{i \in \{1 \dots n\}}$ is a discriminated union or a variant type. $[\tau_i]_{i \in \{1 \dots n\}}$ is a tuple type. Finally, $\mu X. \tau$ is a recursive type.

L-values

L-values in Patina are a combination of a variable and a path. Paths are relative and specify subsections of memory reachable from a L-value. Projection ($p \cdot i$) deconstructs tuples. Unrolling ($\text{unroll } [\tau] p$) deconstructs recursive types. Dereference ($*p$) deconstructs pointers and references. The base path (\bullet) does not deconstruct anything.

Variable x
 Path $p ::= \bullet \mid *p \mid p \cdot i \mid \text{unroll } [\tau] p$

The L-value layer of Patina has two core operations: evaluating a path to a memory location and reading the value at that location. We need to characterize the circumstances under which these operations will progress, and we need to show that these operations preserve certain information, specifically type and initialization data. Higher layers of Patina will build on these properties. We shall start with typing paths and then move into defining evaluation for paths.

The typing judgment for paths does not present any surprises. We use a partial map for the typing context and the type substitution operation is the standard capture-avoiding substitution.

$\Gamma : \text{Variable} \rightarrow \text{Type}$

$\boxed{\Gamma \vdash x@p : \tau}$

PT-BASE
 $\frac{\Gamma(x) = \tau}{\Gamma \vdash x@\bullet : \tau}$

PT-DEOWN
 $\frac{\Gamma \vdash x@p : \sim \tau}{\Gamma \vdash x@*p : \tau}$

PT-DEREF
 $\frac{\Gamma \vdash x@p : \& \ell q \tau}{\Gamma \vdash x@*p : \tau}$

PT-PROJ
 $\frac{\Gamma \vdash x@p : [\tau_i]_{i \in \{1 \dots n\}}}{\Gamma \vdash x@p \cdot i : \tau_i}$

PT-UNROLL
 $\frac{\Gamma \vdash x@p : \mu X. \tau}{\Gamma \vdash x@\text{unroll } [\mu X. \tau] p : [X \mapsto \mu X. \tau] \tau}$

Runtime Memory

Representation and Addressing

To accurately model Rust's memory usage, Patina restricts the contents of a memory cell to void data, an integer, or a pointer to another cell. Tuples and recursive types have no physical memory presence (beyond contiguity for tuples). The memory representation of a variant is a pair of a cell for the discriminant and whatever memory is necessary to store the payload.

We could model this by a map from addresses to memory cell values, but two issues make this inconvenient: the need for address arithmetic and non-unique typing. However, we can add a little extra structure to our model and eliminate these two issues. We wrap the cells inside layouts describing the type structure overlaying memory.

We also separate plain pointer cells into owned cells and reference cells. This is mostly useful for providing just enough type information about the pointer for memory operation purposes. We will explain how we address memory in a moment, so we use a placeholder for now.

Integer z
 Address ρ a placeholder
 Cell $c ::= \text{void} \mid z \mid \sim \rho \mid \& q \rho$
 Layout $l ::= c \mid \langle c : l \rangle \mid [l_i]_{i \in \{1 \dots n\}} \mid \text{roll } [\tau] l$

Due to the extra structure from layouts, addressing memory now requires more than simple labels. Instead we use a runtime analogue of variables and paths. Allocations are chunks of memory allocated and

freed atomically. They correspond to either variables or heap allocations. Routes are like paths except for two key differences.

Allocation α
Route $r ::= \bullet \mid r \cdot i \mid \text{pay } r \mid \text{unroll } [\tau] r$

The \bullet route, projection, and unrolling are effectively identical to their path equivalents. The $\text{pay } r$ route refers to the payload l in a variant $\langle c : l \rangle$. This is primarily used for match-by-reference. Unlike paths, there is no dereference route. This forces any pointer following into path evaluation (into routes) rather than doing so while reading memory at a route. The address placeholder from before is simply a pair of an allocation label and a route. The exception is that unique pointers need only an allocation since a unique pointer always points to a complete heap allocation. That is:

Integer z
Allocation α
Route $r ::= \bullet \mid r \cdot i \mid \text{pay } r \mid \text{unroll } [\tau] r$
Cell $c ::= \text{void} \mid z \mid \sim \alpha \mid \& q \alpha @ r$
Layout $l ::= c \mid \langle c : l \rangle \mid [l_i]_{i \in \{1 \dots n\}} \mid \text{roll } [\tau] l$

Reading

Reading a layout from a route in memory is a straightforward operation. We model the heap as a partial map from allocations to layouts. Routes and layouts interact as you would expect.

$H : \text{Allocation} \rightarrow \text{Layout}$

$\boxed{\text{READ}(H, \alpha @ r, l)}$

RD-BASE $\frac{H(\alpha) = l}{\text{READ}(H, \alpha @ \bullet, l)}$	RD-PROJ $\frac{\text{READ}(H, \alpha @ r, [l_i]_{i \in \{1 \dots n\}})}{\text{READ}(H, \alpha @ r \cdot i, l_i)}$	RD-PAY $\frac{\text{READ}(H, \alpha @ r, \langle c : l \rangle)}{\text{READ}(H, \alpha @ \text{pay } r, l)}$	RD-UNROLL $\frac{\text{READ}(H, \alpha @ r, \text{roll } [\tau'] l)}{\text{READ}(H, \alpha @ \text{unroll } [\tau] r, l)}$
--	--	---	---

As one would expect of a read operation, it's result is unique.

Lemma (Read Uniqueness). *If $\text{READ}(H, \alpha @ r, l)$ and $\text{READ}(H, \alpha @ r, l')$, then $l = l'$.*

Read Uniqueness. We will induct on the derivation of $\text{READ}(H, \alpha @ r, l)$.

RD-BASE Then $r = \bullet$.

By inversion, $H(\alpha) = l$.

There is only one rule (RD-BASE) for deriving $\text{READ}(H, \alpha @ \bullet, l')$.

Thus, by inversion, $H(\alpha) = l'$.

Ergo, $l = l'$.

RD-PROJ Then $r = r' \cdot i$ and $l = l_i$.

By inversion, $\text{READ}(H, \alpha @ r', [l_i]_{i \in \{1 \dots n\}})$.

There is only one rule (RD-PROJ) for deriving $\text{READ}(H, \alpha @ r' \cdot i, l')$.

Thus, by inversion, $\text{READ}(H, \alpha @ r', [l'_i]_{i \in \{1 \dots n'\}})$.

By induction, $[l_i]_{i \in \{1 \dots n\}} = [l'_i]_{i \in \{1 \dots n'\}}$.

Ergo, $n = n'$ and $l_i = l'_i$.

RD-PAY Then $r = \text{pay } r'$.

By inversion, $\text{READ}(H, \alpha @ r', \langle c : l \rangle)$.

There is only one rule (RD-PAY) for deriving $\text{READ}(H, \alpha @ \text{pay } r', l')$.

Thus, by inversion, $\text{READ}(H, \alpha @ r', \langle c' : l' \rangle)$.

By induction, $\langle c : l \rangle = \langle c' : l' \rangle$.
 Ergo, $c = c'$ and $l = l'$.

RD-UNROLL Then $r = \text{unroll } [\tau] r'$.
 By inversion, $\text{READ}(H, \alpha@r', \text{roll } [\tau'] l)$.
 There is only one rule (RD-UNROLL) for deriving $\text{READ}(H, \alpha@\text{unroll } [\tau] r', l')$.
 Thus, by inversion, $\text{READ}(H, \alpha@r', \text{roll } [\tau'] l')$.
 By induction, $\text{roll } [\tau'] l = \text{roll } [\tau''] l'$.
 Ergo, $\tau' = \tau''$ and $l = l'$.

□

Path Evaluation

Now that we can read memory, we can define path evaluation. The only actual work of path evaluation is following the dereferences of pointers to produce a route. In order to connect variables to runtime allocations, we use a partial map tracking the allocation labels of variables.

$V : \text{Variable} \rightarrow \text{Allocation}$

$V; H \vdash x@p \rightarrow \alpha@r$

$$\frac{\text{PE-BASE} \quad V(x) = \alpha}{V; H \vdash x@p \rightarrow \alpha@p}$$

$$\frac{\text{PE-DEOWN} \quad V; H \vdash x@p \rightarrow \alpha@r \quad \text{READ}(H, \alpha@r, \sim \alpha')}{V; H \vdash x@*p \rightarrow \alpha'@p}$$

$$\frac{\text{PE-DEREF} \quad V; H \vdash x@p \rightarrow \alpha@r \quad \text{READ}(H, \alpha@r, \& q \alpha'@r')}{V; H \vdash x@*p \rightarrow \alpha'@r'}$$

$$\frac{\text{PE-PROJ} \quad V; H \vdash x@p \rightarrow \alpha@r}{V; H \vdash x@p \cdot i \rightarrow \alpha@r \cdot i}$$

$$\frac{\text{PE-UNROLL} \quad V; H \vdash x@p \rightarrow \alpha@r}{V; H \vdash x@\text{unroll } [\tau] p \rightarrow \alpha@\text{unroll } [\tau] r}$$

Our core operations are now defined, but we still need to describe the typing of our runtime data so that we can properly state and prove our type preservation lemmas.

Runtime Typing

Route Typing

Typing routes is much like typing paths; however, the `pay r` route requires special attention. Since this route points to the payload of a variant, its type depends upon the value of the discriminant of the variant. This is used for handling match by reference. The match arms have a pointer to the payload in scope, but the type of the payload depends upon the value of the discriminant. For example,

```
// x: <[], int> is a Patina encoding of Option<int>
match x by imm y // y will have route 'pay r' were 'r' is the route of x
// [] (None) case
{
  // y: & imm [] is in scope here
}
```

```

// int (Some) case
{
  // y: & imm int is in scope here
}

```

We handle this with an additional context that records the discriminant value of a variant stored at a route.

$$\Sigma : \text{Allocation} \rightarrow \text{Type}$$

$$\Psi : \text{Allocation} \times \text{Route} \rightarrow \text{Integer}$$

$$\boxed{\Sigma; \Psi \vdash \alpha@r : \tau}$$

$$\frac{\text{RT-BASE} \quad \Sigma(\alpha) = \tau}{\Sigma; \Psi \vdash \alpha@{\bullet} : \tau}$$

$$\frac{\text{RT-PROJ} \quad \Sigma; \Psi \vdash \alpha@r : [\tau_i]_{i \in \{1 \dots n\}}}{\Sigma; \Psi \vdash \alpha@r \cdot i : \tau_i}$$

$$\frac{\text{RT-PAY} \quad \Psi(\alpha, r) = i \quad \Sigma; \Psi \vdash \alpha@r : \langle \tau_i \rangle_{i \in \{1 \dots n\}}}{\Sigma; \Psi \vdash \alpha@{\text{pay}} r : \tau_i}$$

$$\frac{\text{RT-UNROLL} \quad \Sigma; \Psi \vdash \alpha@r : \mu X. \tau}{\Sigma; \Psi \vdash \alpha@{\text{unroll}} [\mu X. \tau] r : [X \mapsto \mu X. \tau] \tau}$$

Layout and Cell Typing

Typing for layouts and cells is what you would expect. However, we restrict the type of `void` to the possible types for other kinds of cells, i.e. `int`, unique pointers, and borrowed references. This means that multi-cell layout types (variants, tuples, and recursive types) imply a non-void value (though they can still contain voids).

$$\boxed{\Sigma; \Psi \vdash l : \tau}$$

$$\frac{\text{LT-VOIDINT}}{\Sigma; \Psi \vdash \text{void} : \text{int}}$$

$$\frac{\text{LT-VOIDDOWN}}{\Sigma; \Psi \vdash \text{void} : \sim \tau}$$

$$\frac{\text{LT-VOIDREF}}{\Sigma; \Psi \vdash \text{void} : \& \ell q \tau}$$

$$\frac{\text{LT-INT}}{\Sigma; \Psi \vdash z : \text{int}}$$

$$\frac{\text{LT-OWN} \quad \Sigma; \Psi \vdash \alpha@{\bullet} : \tau}{\Sigma; \Psi \vdash \sim \alpha : \sim \tau}$$

$$\frac{\text{LT-REF} \quad \Sigma; \Psi \vdash \alpha@r : \tau}{\Sigma; \Psi \vdash \& q \alpha@r : \& \ell q \tau}$$

$$\frac{\text{LT-VOIDVAR}}{\Sigma; \Psi \vdash \langle \text{void} : \text{void} \rangle : \langle \tau_i \rangle_{i \in \{1 \dots n\}}}$$

$$\frac{\text{LT-VAR} \quad \Sigma; \Psi \vdash l : \tau_i}{\Sigma; \Psi \vdash \langle i : l \rangle : \langle \tau_i \rangle_{i \in \{1 \dots n\}}}$$

$$\frac{\text{LT-REC} \quad \forall i. \Sigma; \Psi \vdash l_i : \tau_i}{\Sigma; \Psi \vdash [l_i]_{i \in \{1 \dots n\}} : [\tau_i]_{i \in \{1 \dots n\}}}$$

$$\frac{\text{LT-ROLL} \quad \Sigma; \Psi \vdash l : [X \mapsto \mu X. \tau] \tau}{\Sigma; \Psi \vdash \text{roll} [\mu X. \tau] l : \mu X. \tau}$$

Heap Well-Formedness

With runtime typing now in hand, we can say what it means for the heap to be well-formed. We say a heap H is described by a heap type Σ and a discriminant context Ψ if every allocation in H has the type assigned to it by Σ and Ψ correctly records the discriminants of all the variants in H . The last hypothesis enforces the uniqueness of unique pointers by requiring that if a unique pointer is at any address, then that address is unique. This strengthens our induction hypothesis enough to show that no memory has multiple owners, which we need to prove that no double frees occur. Ensuring every piece of memory has at least one owner is handled elsewhere. Formally,

$$\begin{array}{c}
\text{dom } H = \text{dom } \Sigma \\
\forall \alpha \in \text{dom}(\Sigma). \Sigma(\alpha) \text{ closed} \\
\forall \alpha \in \text{dom } H. \Sigma; \Psi \vdash H(\alpha) : \Sigma(\alpha) \\
\forall (\alpha, r) \in \text{dom } \Psi. \text{READ}(H, \alpha @ r, \langle \Psi(\alpha, r) : l \rangle) \\
\forall \alpha, \alpha', \alpha'' \in \text{dom } H. \forall r, r'. \text{READ}(H, \alpha @ r, \sim \alpha'') \wedge \text{READ}(H, \alpha' @ r', \sim \alpha'') \implies \alpha @ r = \alpha' @ r' \\
\hline
\vdash H : \Sigma; \Psi
\end{array}$$

Read Safety

We are now in position to prove that our read operation is safe. That is, if the heap is well formed and a route has some type, then we can read some layout at that route and that layout has that same type. This is effectively a merger of the type preservation and progress lemmas for READ.

Lemma (Read Safety). *If $\vdash H : \Sigma; \Psi$ and $\Sigma; \Psi \vdash \alpha @ r : \tau$, then there is a layout l such that $\text{READ}(H, \alpha @ r, l)$ and $\Sigma; \Psi \vdash l : \tau$.*

Read Safety. We will use induction on the derivation of $\Sigma; \Psi \vdash \alpha @ r : \tau$.

RT-BASE Then $r = \bullet$.

By inversion, $\Sigma(\alpha) = \tau$.

From $\vdash H : \Sigma; \Psi$, we know that $\text{dom } H = \text{dom } \Sigma$. Thus, $\exists l. H(\alpha) = l$.

Then by RD-BASE, we have $\text{READ}(H, \alpha @ \bullet, l)$.

Again from $\vdash H : \Sigma; \Psi$, we know that $\Sigma; \Psi \vdash H(\alpha) : \Sigma(\alpha)$, which is just $\Sigma; \Psi \vdash l : \tau$ as required.

RT-PROJ Then $r = r' \cdot i$ and $\tau = \tau_i$.

By inversion, $\Sigma; \Psi \vdash \alpha @ r' : [\tau_i]_{i \in \{1 \dots n\}}$.

By induction, $\text{READ}(H, \alpha @ r', l)$ and $\Sigma; \Psi \vdash l : [\tau_i]_{i \in \{1 \dots n\}}$ for some l .

There is only one rule (LT-REC) for deriving $\Sigma; \Psi \vdash l : [\tau_i]_{i \in \{1 \dots n\}}$.

Thus, by inversion, $l = [l'_j]_{j \in \{1 \dots n\}}$ and $\forall j. \Sigma; \Psi \vdash l'_j : \tau_j$. Specifically, $\Sigma; \Psi \vdash l'_i : \tau_i$.

Then by RD-PROJ, we have $\text{READ}(H, \alpha @ r' \cdot i, l'_i)$.

RT-PAY Then $r = \text{pay } r'$ and $\tau = \tau_i$.

By inversion, $\Psi(\alpha, r') = i$ and $\Sigma; \Psi \vdash \alpha @ r' : \langle \tau_i \rangle_{i \in \{1 \dots n\}}$.

By induction, $\text{READ}(H, \alpha @ r', l)$ and $\Sigma; \Psi \vdash l : \langle \tau_i \rangle_{i \in \{1 \dots n\}}$.

From $\vdash H : \Sigma; \Psi$, we know that because $\Psi(\alpha, r') = i$ it follows that $\text{READ}(H, \alpha @ r', \langle i : l' \rangle)$ for some l' .

By Read Uniqueness, we know that $l = \langle i : l' \rangle$.

There is only one rule (LT-VAR) for deriving $\Sigma; \Psi \vdash \langle i : l' \rangle : \langle \tau_i \rangle_{i \in \{1 \dots n\}}$.

Thus, by inversion, $\Sigma; \Psi \vdash l' : \tau_i$.

Then by RD-PAY, we have $\text{READ}(H, \alpha @ \text{pay } r', l')$.

RT-UNROLL Then $r = \text{unroll } [\mu X. \tau'] r'$ and $\tau = [X \mapsto \mu X. \tau'] \tau'$.

By inversion, $\Sigma; \Psi \vdash \alpha @ r' : \mu X. \tau'$.

By induction, $\text{READ}(H, \alpha @ r', l)$ and $\Sigma; \Psi \vdash l : \mu X. \tau'$ for some l .

There is only one rule (LT-ROLL) for deriving $\Sigma; \Psi \vdash l : \mu X. \tau'$.

Thus, by inversion, $l = \text{roll } [\mu X. \tau'] l'$ and $\Sigma; \Psi \vdash l' : [X \mapsto \mu X. \tau'] \tau'$.

Then by RD-UNROLL, we have $\text{READ}(H, \alpha @ \text{unroll } [\mu X. \tau'] r', l')$.

□

Path Type Preservation

The final piece we need for preservation is how to relate the typing context Γ , the variable map V , and the heap type Σ . This is simply the requirement that the map V preserves typing.

$$\frac{\forall x \in \text{dom } \Gamma. \Gamma(x) = \Sigma(V(x))}{\Gamma; V \vdash \Sigma}$$

Finally, we can state and prove path type preservation. Under a well-formed heap and contexts, typed paths evaluate to routes of the same type.

Lemma (Path Type Preservation). *If $\vdash H : \Sigma; \Psi, \Gamma; V \vdash \Sigma, \Gamma \vdash x@p : \tau$, and $V; H \vdash x@p \rightarrow \alpha@r$, then $\Sigma; \Psi \vdash \alpha@r : \tau$.*

Path Type Preservation. We will use induction on the derivation of $V; H \vdash x@p \rightarrow \alpha@r$.

PE-BASE Then $p = \bullet$ and $r = \bullet$.

By inversion, $V(x) = \alpha$.

There is only one rule (PT-BASE) for deriving $\Gamma \vdash x@\bullet : \tau$.

Thus, by inversion, we have $\Gamma(x) = \tau$.

From $\Gamma; V \vdash \Sigma$, we know $\Gamma(x) = \Sigma(V(x))$. Thus, $\tau = \Sigma(\alpha)$.

Then by RT-BASE, we have $\Sigma; \Psi \vdash \alpha@\bullet : \tau$.

PE-DEOWN Then $p = *p'$ and $r = \bullet$.

By inversion, $V; H \vdash x@p' \rightarrow \alpha'@r'$ and $\text{READ}(H, \alpha'@r', \sim \alpha)$.

There are two possible rules for deriving $\Gamma \vdash x@*p' : \tau$.

PT-DEOWN By inversion, $\Gamma \vdash x@p' : \sim \tau$.

By induction, $\Sigma; \Psi \vdash \alpha'@r' : \sim \tau$.

By Read Safety, we have $\text{READ}(H, \alpha'@r', l)$ and $\Sigma; \Psi \vdash l : \sim \tau$ for some l .

By Read Uniqueness, we have $l = \sim \alpha$.

There is only one rule (LT-OWN) for deriving $\Sigma; \Psi \vdash \sim \alpha : \sim \tau$.

Thus, by inversion, $\Sigma; \Psi \vdash \alpha@\bullet : \tau$.

PT-DEREF By inversion, $\Gamma \vdash x@p' : \& \ell q \tau$.

By induction, $\Sigma; \Psi \vdash \alpha'@r' : \& \ell q \tau$.

By Read Safety, we have $\text{READ}(H, \alpha'@r', l)$ and $\Sigma; \Psi \vdash l : \& \ell q \tau$ for some l .

By Read Uniqueness, we have $l = \sim \alpha$.

There are no rules for deriving $\Sigma; \Psi \vdash \sim \alpha : \& \ell q \tau$.

Thus, by inversion, this case is impossible.

PE-DEREF Similar to the PE-DEOWN case, but switching the roles of PT-DEOWN and PT-DEREF.

PE-PROJ Then $p = p' \cdot i$ and $r = r' \cdot i$.

By inversion, $V; H \vdash x@p' \rightarrow \alpha@r'$.

There is only one rule (PT-PROJ) for deriving $\Gamma \vdash x@p' \cdot i : \tau$.

Thus, by inversion, we have $\tau = \tau_i$ and $\Gamma \vdash x@p' : [\tau_i]_{i \in \{1 \dots n\}}$.

By induction, we have $\Sigma; \Psi \vdash \alpha@r' : [\tau_i]_{i \in \{1 \dots n\}}$.

Then by RT-PROJ, we have $\Sigma; \Psi \vdash \alpha@r' \cdot i : \tau_i$.

PE-UNROLL Then $p = \text{unroll } [\tau'] p'$ and $r = \text{unroll } [\tau'] r'$.

By inversion, $V; H \vdash x@p' \rightarrow \alpha@r'$.

There is only one rule (PT-UNROLL) for deriving $\Gamma \vdash x@\text{unroll } [\tau'] p' : \tau$.

Thus, by inversion, we have $\tau' = \mu X. \tau''$, $\tau = [X \mapsto \mu X. \tau''] \tau''$, and $\Gamma \vdash x@p' : \mu X. \tau''$.

By induction, we have $\Sigma; \Psi \vdash \alpha@r' : \mu X. \tau''$.

Then by RT-UNROLL, we have $\Sigma; \Psi \vdash \alpha@\text{unroll } [\mu X. \tau''] r' : [X \mapsto \mu X. \tau''] \tau''$.

□

Tracking Initialization

Types are only half of the information we need to preserve. To properly type check Patina, we need to know which paths point to initialized memory and which do not. This is important for ensuring we never use uninitialized memory and that we correctly free heap memory. We will accomplish this by creating a shadow heap. Instead of tracking actual values, we will simply track whether a cell is initialized. Paths will serve the same function for this shadow heap as routes do for the real heap.

$$\begin{aligned} \text{Hole } h &::= \text{uninit} \mid \text{int} \mid \sim \sigma \mid \& q \tau \mid \langle \tau_i \rangle_{i \in \{1 \dots n\}} \\ \text{Shadow } \sigma &::= h \mid [\sigma_i]_{i \in \{1 \dots n\}} \mid \text{roll } [\tau] \sigma \end{aligned}$$

Since shadows and holes are supposed to be a shadow of the heap, it is unsurprising that they are very similar to layouts and cells. However, there are a few key differences. First, variants are considered an atomic value (and thus a hole) rather than a structure. This is because the discriminant and payload of a variant always share the same initialization state. Second, while the pointer cells contain routes, the pointer holes do not contain paths (the analogue of routes). For references, there is no need because references always point to fully initialized things. Therefore, the reference hole simply records the type of the referent. For unique pointers, there is no path that could possibly be used (otherwise the pointer would not be unique). Therefore, the unique pointer hole records the shadow of the pointed-to heap memory.

Two examples are useful for understanding shadows. First, consider nested pointers:

```
{
  let x: ~~int = ~~3;
  // Here the shadow of x is '~ ~ int'
  // x, *x, and **x are all initialized

  let y: ~int = *x; // *x is now deinitialized
  // Here the shadow of x is '~ uninit' and the shadow of y is '~ int'
  // x is initialized, but *x and **x are not
  // both y and *y are initialized
}
```

Second, consider tuples:

```
{
  let x: [~int, ~int] = [~1, ~2];
  // Here the shadow of x is '[~int, ~int]'
  // x, x.1, *(x.1), x.2, and *(x.2) are all initialized

  free(x.1);
  // Here the shadow of x is '[uninit, ~int]'
  // x, x.1, x.2, and *(x.2) are all initialized, but *(x.1) is not
}
```

The shadow heap gives the typechecker enough information to ensure we do not use or dereference uninitialized memory. For example, here is how it can catch dereferencing a null pointer:

```
{
  let x: ~~int = ~~3;
  free(*x); // deallocates **x, leaving *x uninitialized
  // at this point the shadow of x is '~ uninit'
  let y: int = **x; // ERROR! dereferences a null pointer *x
  // the type checker know the shadow of *x is 'uninit'
  // if a further dereference was allowed (i.e. **x), a null pointer derefernce would occur
}
```

```
// the type checker can avoid null pointer dereferences by forbidding derefs of 'uninit' data
}
```

It can also catch a missing free that would leak memory.

```
{
let x: ~int = ~3;
// the shadow of x is '~ int'

// the stack variable x will be popped off the stack, leaking the heap memory at *x
// the type checker can identify this error by seeing that the shadow of x contains '~' data
// by requiring that shadows cannot contain '~' when freed, the type checker prevents leaks
}
```

Checking Shallow Initialization

We will need an analogue of READ for our shadow heap for the simple purpose of extracting the shadow at a path in the same way we need to extract the layout at a route from the heap. However, this analogue also serves as a very useful property checker: any dereferences in the path are dereferences of initialized memory. This property, called *shallow initialization* is exactly what we need for path progress. A shallowly initialized path can successfully dereference all the pointers necessary to evaluate to a route. For this reason, we call this operation SHALLOW. To support dereferencing borrowed references, we need to construct a shape from the referenced type. We do this with the INIT function.

INIT : Type \rightarrow Shadow

$$\begin{aligned} \text{INIT}(\text{int}) &= \text{int} \\ \text{INIT}(\sim \tau) &= \sim \text{INIT}(\tau) \\ \text{INIT}(\& \ell q \tau) &= \& q \tau \\ \text{INIT}([\tau_i]_{i \in \{1 \dots n\}}) &= [\text{INIT}(\tau_i)]_{i \in \{1 \dots n\}} \\ \text{INIT}(\langle \tau_i \rangle_{i \in \{1 \dots n\}}) &= \langle \tau_i \rangle_{i \in \{1 \dots n\}} \\ \text{INIT}(\mu X. \tau) &= \text{roll } [\mu X. \tau] \text{ INIT}([X \mapsto \mu X. \tau] \tau) \end{aligned}$$

We also need a new context recording the shadows of stack variables.

$$\Upsilon : \text{Variable} \rightarrow \text{Shadow}$$

SHALLOW($\Upsilon, x@p, \sigma$)

$$\begin{array}{c} \text{SI-BASE} \\ \frac{\Upsilon(x) = \sigma}{\text{SHALLOW}(\Upsilon, x@p, \sigma)} \end{array} \quad \begin{array}{c} \text{SI-DEOWN} \\ \frac{\text{SHALLOW}(\Upsilon, x@p, \sim \sigma)}{\text{SHALLOW}(\Upsilon, x@*p, \sigma)} \end{array} \quad \begin{array}{c} \text{SI-DEREF} \\ \frac{\text{SHALLOW}(\Upsilon, x@p, \& q \tau) \quad \text{INIT}(\tau) = \sigma}{\text{SHALLOW}(\Upsilon, x@*p, \sigma)} \end{array}$$

$$\begin{array}{c} \text{SI-PROJ} \\ \frac{\text{SHALLOW}(\Upsilon, x@p, [\sigma_i]_{i \in \{1 \dots n\}})}{\text{SHALLOW}(\Upsilon, x@p \cdot i, \sigma_i)} \end{array} \quad \begin{array}{c} \text{SI-UNROLL} \\ \frac{\text{SHALLOW}(\Upsilon, x@p, \text{roll } [\tau'] \sigma)}{\text{SHALLOW}(\Upsilon, x@\text{unroll } [\tau] p, \sigma)} \end{array}$$

Shadow Typing

We need to ensure that the shadows assigned to variables are consistent with the types assigned to variables, i.e. we want to avoid situations like $\Gamma(x) = \text{int}$ and $\Upsilon(x) = \& \text{imm int}$. We can easily do this by defining

a typing judgement for shadows and defining a well-formedness condition for Υ that ensures it is consistent with Γ . Similar to typing for layouts, `uninit` can only have types that other holes can have.

$\vdash \sigma : \tau$				
ST-UNINT $\frac{}{\vdash \text{uninit} : \text{int}}$	ST-UNOWN $\frac{}{\vdash \text{uninit} : \sim \tau}$	ST-UNREF $\frac{}{\vdash \text{uninit} : \& \ell q \tau}$	ST-UNVAR $\frac{}{\vdash \text{uninit} : \langle \tau_i \rangle_{i \in \{1 \dots n\}}}$	ST-INT $\frac{}{\vdash \text{int} : \text{int}}$
ST-OWN $\frac{\vdash \sigma : \tau}{\vdash \sim \sigma : \sim \tau}$	ST-REF $\frac{}{\vdash \& q \tau : \& \ell q \tau}$	ST-VAR $\frac{}{\vdash \langle \tau_i \rangle_{i \in \{1 \dots n\}} : \langle \tau_i \rangle_{i \in \{1 \dots n\}}}$	ST-REC $\frac{\forall i. \vdash \sigma_i : \tau_i}{\vdash [\sigma_i]_{i \in \{1 \dots n\}} : [\tau_i]_{i \in \{1 \dots n\}}}$	
ST-ROLL $\frac{\vdash \sigma : [X \mapsto \mu X. \tau] \tau}{\vdash \text{roll } [\mu X. \tau] \sigma : \mu X. \tau}$				
$\text{dom } \Gamma = \text{dom } \Upsilon$ $\forall x \in \text{dom } \Gamma. \Gamma(x) \text{ closed}$ $\forall x \in \text{dom } \Gamma. \vdash \Upsilon(x) : \Gamma(x)$ $\frac{}{\vdash \Upsilon : \Gamma}$				

The INIT Function Preserves Types

Since the `INIT` function is supposed to be the fully initialized shadow of the argument type, it should be the case that the resulting shadow has that type. Since `INIT` is only defined on closed types, we must restrict ourselves to them.

Lemma (INIT Type Preservation). *If τ is closed, then $\vdash \text{INIT}(\tau) : \tau$.*

INIT Type Preservation. Induction on the form of τ .

INT If $\tau = \text{int}$, then $\text{INIT}(\text{int}) = \text{int}$. From ST-INT, we have $\vdash \text{int} : \text{int}$.

OWN If $\tau = \sim \tau'$, then $\text{INIT}(\sim \tau') = \sim \text{INIT}(\tau')$.

Since τ is closed, τ' is also closed.

Thus, by induction, $\vdash \text{INIT}(\tau') : \tau'$.

From ST-OWN, we have $\vdash \sim \text{INIT}(\tau') : \sim \tau'$.

REF If $\tau = \& \ell q \tau'$, then $\text{INIT}(\& \ell q \tau') = \& q \tau'$.

$\vdash \& q \tau' : \& \ell q \tau'$ follows immediately by ST-REF.

REC If $\tau = [\tau'_i]_{i \in \{1 \dots n\}}$, then $\text{INIT}([\tau'_i]_{i \in \{1 \dots n\}}) = [\text{INIT}(\tau'_i)]_{i \in \{1 \dots n\}}$.

Since τ is closed, τ'_i is closed for all i .

Thus, by induction, for all i we have $\vdash \text{INIT}(\tau'_i) : \tau'_i$.

Then by ST-REC, we have $\vdash [\text{INIT}(\tau'_i)]_{i \in \{1 \dots n\}} : [\tau'_i]_{i \in \{1 \dots n\}}$.

VAR If $\tau = \langle \tau'_i \rangle_{i \in \{1 \dots n\}}$, then $\text{INIT}(\langle \tau'_i \rangle_{i \in \{1 \dots n\}}) = \langle \tau'_i \rangle_{i \in \{1 \dots n\}}$.

$\vdash \langle \tau'_i \rangle_{i \in \{1 \dots n\}} : \langle \tau'_i \rangle_{i \in \{1 \dots n\}}$ follows immediately from ST-VAR.

FIX If $\tau = \mu X. \tau'$, then $\text{INIT}(\mu X. \tau') = \text{roll } [\mu X. \tau'] \text{INIT}([X \mapsto \mu X. \tau'] \tau')$.

Since τ is closed, τ' might not be closed; however, the only possible free variable is X .

Thus, the unrolling $[X \mapsto \mu X. \tau'] \tau'$ is closed.

Hence, by induction, $\vdash \text{INIT}([X \mapsto \mu X. \tau'] \tau') : [X \mapsto \mu X. \tau'] \tau'$.

Then by ST-ROLL, we have $\vdash \text{roll } [\mu X. \tau'] \text{INIT}([X \mapsto \mu X. \tau'] \tau') : \mu X. \tau'$.

□

Shallow Preserves Types

Just as with READ, SHALLOW preserves types. This lemma is used primarily to demonstrate the impossibility of certain subcases of dereferences in later proofs.

Lemma (SHALLOW Type Preservation). *If $\vdash \Upsilon : \Gamma$, $\Gamma \vdash x@p : \tau$, and $\text{SHALLOW}(\Upsilon, x@p, \sigma)$ then $\vdash \sigma : \tau$.*

SHALLOW *Type Preservation*. Induction on the derivation of $\text{SHALLOW}(\Upsilon, x@p, \sigma)$.

SI-BASE Then $p = \bullet$.

By inversion, $\Upsilon(x) = \sigma$.

There is only one rule (PT-BASE) for deriving $\Gamma \vdash x@\bullet : \tau$.

Thus, by inversion, we have $\Gamma(x) = \tau$.

From $\vdash \Upsilon : \Gamma$, we have $\vdash \Upsilon(x) : \Gamma(x)$.

Ergo, we have $\vdash \sigma : \tau$.

SI-DEOWN Then $p = *p'$.

By inversion, $\text{SHALLOW}(\Upsilon, x@p', \sim \sigma)$.

There are two possible rules for deriving $\Gamma \vdash x@*p' : \tau$.

PT-DEOWN By inversion, $\Gamma \vdash x@p' : \sim \tau$.

By induction, $\vdash \sim \sigma : \sim \tau$.

There is only one rule (ST-OWN) for deriving $\vdash \sim \sigma : \sim \tau$.

Thus, by inversion, $\vdash \sigma : \tau$.

PT-DEREF By inversion, $\Gamma \vdash x@p' : \& \ell q \tau$.

By induction, $\vdash \sim \sigma : \& \ell q \tau$.

There are no rules for deriving $\vdash \sim \sigma : \& \ell q \tau$.

Thus, by inversion, this case is impossible.

SI-DEREF Then $p = *p'$.

By inversion, $\text{SHALLOW}(\Upsilon, x@p', \& q \tau)$ and $\text{INIT}(\tau) = \sigma$.

There are two possible rules for deriving $\Gamma \vdash x@*p' : \tau$.

PT-DEOWN By inversion, $\Gamma \vdash x@p' : \sim \tau$.

By induction, $\vdash \& q \tau : \sim \tau$.

There are no rules for deriving $\vdash \& q \tau : \sim \tau$.

Thus, by inversion, this case is impossible.

PT-DEREF By inversion, $\Gamma \vdash x@p' : \& \ell q \tau$.

Since $\vdash \Upsilon : \Gamma$ guarantees that $\Gamma(x)$ is closed and a simple inspection of the path typing rules shows that this implies all well typed paths have closed types, we know that τ is closed.

By INIT Type Preservation, $\text{INIT}(\tau) = \sigma$ implies $\vdash \sigma : \tau$.

SI-PROJ Then $p = p' \cdot i$ and $\sigma = \sigma_i$.

By inversion, $\text{SHALLOW}(\Upsilon, x@p', [\sigma_i]_{i \in \{1 \dots n\}})$.

There is only one rule (PT-PROJ) for deriving $\Gamma \vdash x@p' \cdot i : \tau$.

Thus, by inversion, we have $\tau = \tau_i$ and $\Gamma \vdash x@p' : [\tau_i]_{i \in \{1 \dots n\}}$.

By induction, we have $\vdash [\sigma_i]_{i \in \{1 \dots n\}} : [\tau_i]_{i \in \{1 \dots n\}}$.

There is only one rule (ST-REC) for deriving $\vdash [\sigma_i]_{i \in \{1 \dots n\}} : [\tau_i]_{i \in \{1 \dots n\}}$.

Thus, by inversion, we have $\vdash \sigma_i : \tau_i$.

SI-UNROLL Then $p = \text{unroll } [\tau'] p'$.

By inversion, $\text{SHALLOW}(\Upsilon, x@p', \text{roll } [\tau''] \sigma)$.

There is only one rule (PT-UNROLL) for deriving $\Gamma \vdash x@\text{unroll } [\tau'] p' : \tau$.

By inversion, $\tau' = \mu X. \tau'''$, $\tau = [X \mapsto \mu X. \tau'''] \tau'''$, and $\Gamma \vdash x@p' : \mu X. \tau'''$.

By induction, $\vdash \text{roll } [\tau''] \sigma : \mu X. \tau'''$.

There is only one rule (ST-ROLL) for deriving this.
Thus, by inversion, we have $\tau'' = \mu X. \tau'''$ and $\vdash \sigma : [X \mapsto \mu X. \tau'''] \tau'''$.

□

Shadow Checking the Heap

Of course, we also need to check that our shadow heap correctly models the actual heap. For the most part, it is exactly the relation you would expect. The use of INIT in LS-REF and LS-VAR ensures that the pointees of borrowed references and the payload of initialized variants are fully initialized. The DROPPABLE(l) judgment checks that l owns no initialized heap memory, i.e. $\sim \alpha$ is not reachable from l except through references. These layouts are safe to deallocate because we cannot leak heap memory by doing so. The $V \vdash H : \Upsilon$ judgment ensures that the shadow heap corresponds to the real heap. It also ensures that all heap memory is reachable from the stack and that every heap allocation has at least one owner: either a stack variable or a unique pointer.

DROPPABLE(l)

$\frac{\text{D-VOID}}{\text{DROPPABLE}(\text{void})}$	$\frac{\text{D-INT}}{\text{DROPPABLE}(z)}$	$\frac{\text{D-REF}}{\text{DROPPABLE}(\& q \alpha @ r)}$	$\frac{\text{D-VAR} \quad \text{DROPPABLE}(c) \quad \text{DROPPABLE}(l)}{\text{DROPPABLE}(\langle c : l \rangle)}$
	$\frac{\text{D-REC} \quad \forall i. \text{DROPPABLE}(l_i)}{\text{DROPPABLE}([l_i]_{i \in \{1 \dots n\}})}$		$\frac{\text{D-ROLL} \quad \text{DROPPABLE}(l)}{\text{DROPPABLE}(\text{roll } [\tau] l)}$

$H \vdash l : \sigma$

$\frac{\text{LS-VOID}}{H \vdash \text{void} : \text{uninit}}$	$\frac{\text{LS-INT}}{H \vdash z : \text{int}}$	$\frac{\text{LS-OWN} \quad \text{READ}(H, \alpha @ \bullet, l) \quad H \vdash l : \sigma}{H \vdash \sim \alpha : \sim \sigma}$	$\frac{\text{LS-REF} \quad \text{READ}(H, \alpha @ r, l) \quad H \vdash l : \text{INIT}(\tau)}{H \vdash \& q \alpha @ r : \& q \tau}$
$\frac{\text{LS-VOIDVAR}}{H \vdash \langle \text{void} : \text{void} \rangle : \text{uninit}}$	$\frac{\text{LS-DROPVAR} \quad \text{DROPPABLE}(l)}{H \vdash \langle i : l \rangle : \text{uninit}}$	$\frac{\text{LS-VAR} \quad H \vdash l : \text{INIT}(\tau_i)}{H \vdash \langle i : l \rangle : \langle \tau_i \rangle_{i \in \{1 \dots n\}}}$	$\frac{\text{LS-REC} \quad \forall i. H \vdash l_i : \sigma_i}{H \vdash [l_i]_{i \in \{1 \dots n\}} : [\sigma_i]_{i \in \{1 \dots n\}}}$
	$\frac{\text{LS-ROLL} \quad H \vdash l : \sigma}{H \vdash \text{roll } [\mu X. \tau] l : \text{roll } [\mu X. \tau] \sigma}$		

$$\begin{array}{l}
\text{dom } V = \text{dom } \Upsilon \\
\text{dom } H = \text{REACHABLE}(V, H) \\
\forall \alpha \in \text{dom } H. \exists x. V(x) = \alpha \vee \exists \alpha' @ r. \text{READ}(H, \alpha' @ r, \sim \alpha) \\
\forall x \in \text{dom } V. H \vdash H(V(x)) : \Upsilon(x) \\
\hline
V \vdash H : \Upsilon
\end{array}$$

Shadow Preservation

The initialization data, in the form of a shadow, is preserved by path evaluation and READ so that the resulting layout is correctly described by the shadow.

Lemma (Shadow Preservation). *If $V \vdash H : \Upsilon$, $\text{SHALLOW}(\Upsilon, x@p, \sigma)$, $V; H \vdash x@p \rightarrow \alpha@r$, and $\text{READ}(H, \alpha@r, l)$ then $H \vdash l : \sigma$.*

Shadow Preservation. Induction on the derivation fo $V; H \vdash x@p \rightarrow \alpha@r$.

PE-BASE Then $p = \bullet$ and $r = \bullet$.

By inversion, $V(x) = \alpha$.

There is only one rule (SI-BASE) for deriving $\text{SHALLOW}(\Upsilon, x@\bullet, \sigma)$.

Thus, by inversion, $\Upsilon(x) = \sigma$.

There is only one rule (RD-BASE) for deriving $\text{READ}(H, \alpha@\bullet, l)$.

Thus, by inversion, $H(\alpha) = l$.

By $V \vdash H : \Upsilon$, we have $H \vdash H(V(x)) : \Upsilon(x)$.

However, this is just $H \vdash l : \sigma$ as required.

PE-DEOWN Then $p = *p'$ and $r = \bullet$.

By inversion, $V; H \vdash x@p' \rightarrow \alpha'@r'$ and $\text{READ}(H, \alpha'@r', \sim \alpha)$.

There are two possible rules for deriving $\text{SHALLOW}(\Upsilon, x@*p', \sigma)$.

SI-DEOWN By inversion, $\text{SHALLOW}(\Upsilon, x@p', \sim \sigma)$.

By induction, $H \vdash \sim \alpha : \sim \sigma$.

There is only one rule (LS-OWN) for deriving this.

Thus, by inversion, we have $\text{READ}(H, \alpha@\bullet, l')$ and $H \vdash l' : \sigma$.

By Read Uniqueness, we have $l' = l$.

Ergo, $H \vdash l : \sigma$ as required.

SI-DEREF By inversion, $\text{SHALLOW}(\Upsilon, x@p', \& q \tau)$ and $\text{INIT}(\tau) = \sigma$.

By induction, $H \vdash \sim \alpha : \& q \tau$.

There are no rules for deriving this.

Thus, by inversion, this case is impossible.

PE-DEREF Then $p = *p'$.

By inversion, $V; H \vdash x@p' \rightarrow \alpha'@r'$ and $\text{READ}(H, \alpha'@r', \& q \alpha@r)$.

There are two possible rules for deriving $\text{SHALLOW}(\Upsilon, x@*p', \sigma)$.

SI-DEOWN By inversion, $\text{SHALLOW}(\Upsilon, x@p', \sim \sigma)$.

By induction, $H \vdash \& q \alpha@r : \sim \sigma$.

There are no rules for deriving this.

Thus, by inversion, this case is impossible.

SI-DEREF By inversion, $\text{SHALLOW}(\Upsilon, x@p', \& q \tau)$ and $\text{INIT}(\tau) = \sigma$.

By induction, $H \vdash \& q \alpha@r : \& q \tau$.

There is only one rule (LS-REF) for deriving this.

Thus, by induction, we have $\text{READ}(H, \alpha@r, l')$ and $H \vdash l' : \text{INIT}(\tau)$.

By Read Uniqueness, we have $l' = l$.

Ergo, $H \vdash l : \sigma$ as required.

PE-PROJ Then $p = p' \cdot i$ and $r = r' \cdot i$.

By inversion, $V; H \vdash x@p' \rightarrow \alpha@r'$.

There is only one rule (SI-PROJ) for deriving $\text{SHALLOW}(\Upsilon, x@p' \cdot i, \sigma)$.

Thus, by inversion, we have $\sigma = \sigma_i$ and $\text{SHALLOW}(\Upsilon, x@p', [\sigma_i]_{i \in \{1..n\}})$.

There is only one rule (RD-PROJ) for deriving $\text{READ}(H, \alpha@r' \cdot i, l)$.

Thus, by inversion, we have $l = l_i$ and $\text{READ}(H, \alpha@r', [l_i]_{i \in \{1..n\}})$.

By induction, we have $H \vdash [l_i]_{i \in \{1..n\}} : [\sigma_i]_{i \in \{1..n\}}$.

There is only one rule (LS-REC) for deriving this.

Thus, by inversion, we have $H \vdash l_i : \sigma_i$ as required.

PE-UNROLL Then $p = \text{unroll } [\tau] p'$ and $r = \text{unroll } [\tau] r'$.
 By inversion, $V; H \vdash x@p' \rightarrow \alpha@r'$.
 There is only one rule (SI-UNROLL) for deriving $\text{SHALLOW}(\Upsilon, x@\text{unroll } [\tau] p', \sigma)$.
 Thus, by inversion, we have $\text{SHALLOW}(\Upsilon, x@p', \text{roll } [\tau'] \sigma)$.
 There is only one rule (RD-UNROLL) for deriving $\text{READ}(H, \alpha@\text{unroll } [\tau] r', l)$.
 Thus, by inversion, we have $\text{READ}(H, \alpha@r', \text{roll } [\tau''] l)$.
 By induction, $H \vdash \text{roll } [\tau''] l : \text{roll } [\tau'] \sigma$.
 There is only one rule (LS-ROLL) for deriving this.
 Thus, by inversion, we have $\tau' = \tau'' = \mu X. \tau'''$ and $H \vdash l : \sigma$ required.

□

Path Progress

Finally, we have all we need to proof path evaluation progress under the right circumstances. Assuming all our contexts and the heap are well-formed, then a shallowly initialized, well-typed path can evaluate to some route. We can then rely on path preservation to show that route has the same type as the path. Read safety then lets us show that we can read a layout at that route and that it has the same type. Shadow preservation guarantees us that our model of the layout reflects reality. All of the key pieces we need for the higher layers of Patina.

Lemma (Path Progress). *If $\vdash H : \Sigma; \Psi, \Gamma; V \vdash \Sigma, \vdash \Upsilon : \Gamma, V \vdash H : \Upsilon, \Gamma \vdash x@p : \tau$, and $\text{SHALLOW}(\Upsilon, x@p, \sigma)$ then there is some $\alpha@r$ such that $V; H \vdash x@p \rightarrow \alpha@r$.*

Path Progress. Induction on the derivation of $\Gamma \vdash x@p : \tau$.

PT-BASE Then $p = \bullet$.
 By inversion, $\Gamma(x) = \tau$.
 From $\Gamma; V \vdash \Sigma$, we have $\Gamma(x) = \Sigma(V(x))$.
 Ergo, $\exists \alpha. V(x) = \alpha$.
 Then by PE-BASE, we have $V; H \vdash x@\bullet \rightarrow \alpha@\bullet$.

PT-DEOWN Then $p = *p'$.
 By inversion, $\Gamma \vdash x@p' : \sim \tau$.
 There are two possible rules for deriving $\text{SHALLOW}(\Upsilon, x@*p', \sigma)$.

SI-DEOWN By inversion, $\text{SHALLOW}(\Upsilon, x@p', \sim \sigma)$.
 By induction, $\exists \alpha@r. V; H \vdash x@p' \rightarrow \alpha@r$.
 By Path Type Preservation, $\Sigma; \Psi \vdash \alpha@r : \sim \tau$.
 By Read Safety, $\text{READ}(H, \alpha@r, l)$ and $\Sigma; \Psi \vdash l : \sim \tau$ for some l .
 By Shadow Preservation, $H \vdash l : \sim \sigma$.
 There is only one rule (LS-OWN) for deriving this.
 Thus, by inversion, we have $l = \sim \alpha', \text{READ}(H, \alpha'@\bullet, l')$, and $H \vdash l' : \sigma$.
 Using PE-DEOWN with $V; H \vdash x@p' \rightarrow \alpha@r$ and $\text{READ}(H, \alpha@r, \sim \alpha')$,
 we get $V; H \vdash x@*p' \rightarrow \alpha'@\bullet$, which is what was required.

SI-DEREF By inversion, $\text{SHALLOW}(\Upsilon, x@p', \& q \tau)$.
 By SHALLOW Type Preservation, $\vdash \& q \tau : \sim \tau$.
 There are no rules for deriving this.
 Thus, by inversion, this case is impossible.

PT-DEREF Then $p = *p'$.
 By inversion, $\Gamma \vdash x@p' : \& \ell q \tau$.
 There are two possible rules for deriving $\text{SHALLOW}(\Upsilon, x@*p', \sigma)$.

SI-DEOWN By inversion, $\text{SHALLOW}(\Upsilon, x@p', \sim \sigma)$.
 By SHALLOW Type Preservation, $\vdash \sim \sigma : \& \ell q \tau$.
 There are no rules for deriving this.
 Thus, by inversion, this case is impossible.

SI-DEREF By inversion, $\text{SHALLOW}(\Upsilon, x@p', \& q \tau)$.
 By induction, $\exists \alpha@r. V; H \vdash x@p' \rightarrow \alpha@r$.
 By Path Type Preservation, $\Sigma; \Psi \vdash \alpha@r : \& \ell q \tau$.
 By Read Safety, $\text{READ}(H, \alpha@r, l)$ and $\Sigma; \Psi \vdash l : \& \ell q \tau$ for some l .
 By Shadow Preservation, $H \vdash l : \& q \tau$.
 There is only one rule (LS-REF) for deriving this.
 Thus, by inversion, we have $l = \& q \alpha'@r'$ and $\text{READ}(H, \alpha'@r', l')$.
 Using PE-DEREF with $V; H \vdash x@p' \rightarrow \alpha@r$ and $\text{READ}(H, \alpha@r, \& q \alpha'@r')$,
 we get $V; H \vdash x@*p' \rightarrow \alpha'@r'$, which is what was required.

PT-PROJ Then $p = p' \cdot i$ and $\tau = \tau_i$.
 By inversion, $\Gamma \vdash x@p' : [\tau_i]_{i \in \{1 \dots n\}}$.
 There is only one rule (SI-PROJ) for deriving $\text{SHALLOW}(\Upsilon, x@*p', \sigma)$.
 Thus, by inversion, we have $\sigma = \sigma_i$ and $\text{SHALLOW}(\Upsilon, x@p', [\sigma_i]_{i \in \{1 \dots n\}})$.
 By induction, $\exists \alpha@r. V; H \vdash x@p' \rightarrow \alpha@r$.
 Then by PE-PROJ, we have $V; H \vdash x@p' \cdot i \rightarrow \alpha@r \cdot i$.

PT-UNROLL Then $p = \text{unroll } [\mu X. \tau'] p'$ and $\tau = [X \mapsto \mu X. \tau'] \tau'$.
 By inversion, $\Gamma \vdash x@p' : \mu X. \tau'$.
 There is only one rule (SI-UNROLL) for deriving $\text{SHALLOW}(\Upsilon, x@unroll [\mu X. \tau'] p', \sigma)$.
 Thus, by inversion, $\text{SHALLOW}(\Upsilon, x@p', \text{roll } [\tau'] \sigma)$.
 By induction, $\exists \alpha@r. V; H \vdash x@p' \rightarrow \alpha@r$.
 Then by PE-UNROLL, we have $V; H \vdash x@unroll [\mu X. \tau'] p' \rightarrow \alpha@unroll [\mu X. \tau'] r'$.

□

R-Values

R-values in Patina are expressions that can evaluate to a layout. This layer of Patina involves the borrow checker, which ensures the borrowed reference guarantees. In order to focus on the details of the borrow checker, we will pare down our types. Variants, tuples, and recursive types have straightforward behavior here and would unnecessarily clutter the discussion. Going forward, we will use this subset of our earlier language:

Context	Γ	:	Variable \rightarrow Type
Shadow Heap	Υ	:	Variable \rightarrow Shadow
Map	V	:	Variable \rightarrow Allocation
Heap Type	Σ	:	Allocation \rightarrow Type
Heap	H	:	Allocation \rightarrow Layout
Lifetime	ℓ		
Variable	x		
Integer	z		
Allocation	α		
Qualifier	q	::=	imm mut
Type	τ	::=	int $\sim \tau$ $\& \ell q \tau$
Path	p	::=	\bullet $*p$
Cell	c	::=	void z $\sim \alpha$ $\& q \alpha$
Layout	l	::=	c
Hole	h	::=	uninit int $\sim \sigma$ $\& q \tau$
Shadow	σ	::=	h

Routes would only consist of \bullet , so they are now elided. Without variants, the Ψ context is irrelevant, so it is elided as well.

We define three kinds of expressions: integer constants, using the value at an L-value, and creating a borrowed reference to the value at an L-value.

$$\text{Expression } e ::= z \mid x@p \mid \& \ell q x@p$$

Tracking Loans¹

When we use an L-value in an expression, we must ensure the operation we are performing on it does not violate the promises of existing loans, and we must ensure that the value it refers to is fully initialized. When we create a borrowed reference, we must ensure that the promises of the borrow can actually be upheld for the duration of the borrow. In order to make these checks, we need to know what is borrowed, how it is borrowed, and for how long it is borrowed. We will accomplish this by adding a stack of loans to every level of the shadow heap:

$$\begin{aligned} \text{Hole } h &::= \text{uninit} \mid \text{int} \mid \sim \sigma \mid \& q \sigma \\ \text{Bank } \$ &::= \emptyset \mid \$, (\ell, q) \\ \text{Shadow } \sigma &::= \$: h \end{aligned}$$

Each pair (ℓ, q) describes the duration and manner of a loan, with more recent loans at the front of the stack. Since every path can be borrowed, every path must have a loan stack (or bank) associated with it. The easiest way to do that is to tag every shadow with a bank. Note that the borrowed reference hole has changed. When we only cared about initialization, we used the hole $\& q \tau$. Since everything reachable via the borrow was initialized, we did not need to record the initialization state. Now, however, we do need to record the loan state of things reachable via the borrow. Thus, we make the reference hole recursive: $\& q \sigma$.

¹Rust itself and earlier models of Patina used a dramatically different form of loan tracking and borrow checking. A universal set of loans is kept and threaded through the type checker. Each loan generates a set of paths that cannot be accessed without violating a loan guarantee. The borrow checker would then make sure that no path used was in that set of forbidden paths.

This structure does not lend itself to proofs, so we devised this alternate system. The system we present below is syntax driven and, we think, easier to understand. Since proving this system matches the original would run into the problem we tried to avoid in the first place, we satisfied ourselves with testing the two implementations. The original model and this model agree on all the tests used in the original model.

Shadow Typing with Banks

Since our shadows now include loan information, we need to modify our shadow typing. While we have included lifetimes in our model from the beginning, this is the first time we will actually utilize them. Before we can discuss how to check banks, we must first discuss what lifetimes are, where they come from, and what they give us.

The Lifetime Relation

Lifetimes, as the name implies, encompass a notion of duration. Specifically, they describe the duration a value is allocated and the duration of a loan. Each new block of scope defines a new lifetime. So there is a clear ordering on lifetimes, $\ell \leq \ell'$ if the duration of ℓ is not longer than the duration of ℓ' . More formally:

$$\mathcal{L} \subseteq \text{Lifetime} \times \text{Lifetime}$$

$$\frac{\mathcal{L} \text{ is a poset}}{\vdash \mathcal{L}}$$

where $(\ell, \ell') \in \mathcal{L}$ means $\ell \leq \ell'$. Rust allows some other sources of lifetimes not relevant to the current model: lifetime parameters for functions and *static*, which is the oldest lifetime and typically used for string constants. Though the lifetimes in our model form a total order, the general Rust lifetime relation is only a partial order. Since we intend this model to be a potential basis for all of Rust, we will ignore the totality of our lifetimes.

Bank Coherence

As mentioned before, lifetimes come from scope blocks. This means that every variable has a lifetime, which comes from the block it was defined in. Since variables are deallocated at the end of the block in which they are defined, that lifetime is an accurate model. We record the lifetimes of variables in a new context:

$$L : \text{Variable} \rightarrow \text{Lifetime}$$

Furthermore, if variables only exist for a particular lifetime, then it makes sense that loans of memory owned by the variable cannot be longer than the lifetime of the variable. Otherwise the deallocation of the variable could create dangling pointers elsewhere.

Finally, we must ensure that more recent loans are not inconsistent with older loans of the same path. If a path is already loaned out for some lifetime ℓ and we loan it out again for some lifetime ℓ' , then $(\ell', \ell) \in \mathcal{L}$. A loan only controls a value during the loan's lifetime, so it cannot make promises for lifetimes longer than that. Additionally, given a mutable loan we can reloan it immutably, but not the other way. A mutable loan knows it has unique access to a value, so it can guarantee the value can be immutable if it wishes. On the other hand, an immutable loan promises that the value will not be changed for its duration, which violates the requirement that a mutable loan should be able to mutate its referent. We characterize this relationship with an ordering on qualifiers: $\text{imm} < \text{mut}$.

$$\boxed{\mathcal{L} \vdash \$ \leq \ell}$$

$$\begin{array}{ccc} \text{BWF-NONE} & \text{BWF-ONE} & \text{BWF-TWO} \\ \frac{}{\mathcal{L} \vdash \emptyset \leq \ell} & \frac{(\ell', \ell) \in \mathcal{L}}{\mathcal{L} \vdash \emptyset, (\ell', q) \leq \ell} & \frac{(\ell_2, \ell_1) \in \mathcal{L} \quad q_2 \leq q_1 \quad \mathcal{L} \vdash \$, (\ell_1, q_1) \leq \ell}{\mathcal{L} \vdash \$, (\ell_1, q_1), (\ell_2, q_2) \leq \ell} \end{array}$$

Bounded Shadow Typing

We need to update our shadow typing judgment to support checking bank coherence, but we need to check two new things as well. First, we will require that variables can only reference memory that they will not outlast. This means we require borrowed references to have lifetimes not shorter than the bounding lifetime, which is the reverse of the requirement on loans. This should make sense, we can only use memory that will exist at least as long as us, but we can only loan out our memory to others that will not last longer than us.

Second, we need to make sure that more specific loans do not contradict more general loans. For example, if a variable x is loaned immutably, then a mutable loan of $*x$ would not make sense. The immutable loan of x promises that no memory reachable via x will change, but the mutable loan of $*x$ promises that the memory reachable via $*x$ can be changed through the borrowed reference. Similarly, if x was loaned mutably, then $*x$ could not be loaned at all. The loan of x promises that only the borrowed reference has access to memory reachable from x , but the loan of $*x$ creates an alias to part of that memory. In general, a wider immutable loan requires more specific loans to not be mutable loans, and a wider mutable loan prevents any more specific loans.

Maybe Qualifier $mq ::= \text{none} \mid q$

$\boxed{\text{NOT-MUT}(\$)}$

$\frac{\text{NM-EMPTY}}{\text{NOT-MUT}(\emptyset)}$

$\frac{\text{NM-IMM}}{\text{NOT-MUT}(\$, (\ell, \text{imm}))}$

$\boxed{\mathcal{L} \vdash \sigma : \tau \text{ controlled by } mq \leq \ell}$

$\frac{\text{ST-NONE} \quad \mathcal{L} \vdash \emptyset \leq \ell \quad \mathcal{L} \vdash h : \tau \text{ controlled by } \text{none} \leq \ell}{\mathcal{L} \vdash (\emptyset : h) : \tau \text{ controlled by } \text{none} \leq \ell}$

$\frac{\text{ST-SOME} \quad \mathcal{L} \vdash \$, (\ell', q) \leq \ell \quad \mathcal{L} \vdash h : \tau \text{ controlled by } q \leq \ell}{\mathcal{L} \vdash (\$, (\ell', q) : h) : \tau \text{ controlled by } \text{none} \leq \ell}$

$\frac{\text{ST-IMM} \quad \mathcal{L} \vdash \$ \leq \ell \quad \text{NOT-MUT}(\$) \quad \mathcal{L} \vdash h : \tau \text{ controlled by } \text{imm} \leq \ell}{\mathcal{L} \vdash (\$: h) : \tau \text{ controlled by } \text{imm} \leq \ell}$

$\frac{\text{ST-MUT} \quad \mathcal{L} \vdash \emptyset \leq \ell \quad \mathcal{L} \vdash h : \tau \text{ controlled by } \text{mut} \leq \ell}{\mathcal{L} \vdash (\emptyset : h) : \tau \text{ controlled by } \text{mut} \leq \ell}$

$\boxed{\mathcal{L} \vdash h : \tau \text{ controlled by } mq \leq \ell}$

$$\begin{array}{c}
\text{HT-UNINT} \\
\hline
\mathcal{L} \vdash \text{uninit} : \text{int controlled by } mq \leq \ell
\end{array}
\qquad
\begin{array}{c}
\text{HT-UNOWN} \\
\hline
\mathcal{L} \vdash \text{uninit} : \sim \tau \text{ controlled by } mq \leq \ell
\end{array}$$

$$\begin{array}{c}
\text{HT-UNREF} \\
\hline
(l, \ell') \in \mathcal{L} \\
\hline
\mathcal{L} \vdash \text{uninit} : \& \ell' q \tau \text{ controlled by } mq \leq \ell
\end{array}
\qquad
\begin{array}{c}
\text{HT-INT} \\
\hline
\mathcal{L} \vdash \text{int} : \text{int controlled by } mq \leq \ell
\end{array}$$

$$\begin{array}{c}
\text{HT-OWN} \\
\hline
\mathcal{L} \vdash \sigma : \tau \text{ controlled by } mq \leq \ell \\
\hline
\mathcal{L} \vdash \sim \sigma : \sim \tau \text{ controlled by } mq \leq \ell
\end{array}
\qquad
\begin{array}{c}
\text{HT-REFMUT} \\
\hline
(l, \ell') \in \mathcal{L} \quad \mathcal{L} \vdash \sigma : \tau \text{ controlled by } mq \leq \ell \quad \sigma \text{ init} \\
\hline
\mathcal{L} \vdash \& \text{mut } \sigma : \& \ell' \text{ mut } \tau \text{ controlled by } mq \leq \ell
\end{array}$$

$$\begin{array}{c}
\text{HT-REFIMM} \\
\hline
(l, \ell') \in \mathcal{L} \quad \mathcal{L} \vdash \sigma : \tau \text{ controlled by } \text{imm} \leq \ell \quad \sigma \text{ init} \\
\hline
\mathcal{L} \vdash \& \text{imm } \sigma : \& \ell' \text{ imm } \tau \text{ controlled by } mq \leq \ell
\end{array}$$

where “ σ init” in HT-REFMUT and HT-REFIMM is a judgment that ensures a shadow is fully initialized, i.e. it cannot reach a uninit. Note that the interior of immutable borrowed references is always controlled by an immutable loan. This prevents mutable loans of the interior of immutable borrows, which are nonsensical.

Bounded Shadow Heap Well-Formedness

With our shadow typing updated to check banks, we can update what we mean by a well-formed shadow heap. The lifetime relation and the lifetime map are now included. We ensure the lifetime map is well-formed by requiring its domain to be the same as the domain of the type context and the shadow heap, and by requiring that every lifetime in its range to be in the lifetime relation.

$$\begin{array}{c}
\text{dom } \Gamma = \text{dom } \Upsilon = \text{dom } L \\
\forall x \in \text{dom } \Gamma. \Gamma(x) \text{ closed} \\
\forall x \in \text{dom } L. (L(x), L(x)) \in L \\
\forall x \in \text{dom } \Upsilon. \mathcal{L} \vdash \Upsilon(x) : \Gamma(x) \text{ controlled by } \text{none} \leq L(x) \\
\hline
\mathcal{L} \vdash \Upsilon : \Gamma \leq L
\end{array}$$

Ensuring Operations Do Not Violate Loans

The most important duty of the borrow checker is making sure that performing some operation does not violate the promises of existing loans. We cannot read a path that has promised unique access to another. We cannot write to a path that has promised to remain unchanged or that we cannot read. We cannot move from a path that we do not own or that we cannot write.

Additionally, we can only use a path that is fully initialized, or our program might try to touch an uninitialized value and get stuck. This means that $\text{SHALLOW}(\Upsilon, x@p, \sigma)$ holds and that σ init holds. That is, enough memory is initialized to evaluate the L-value to a memory location, and everything reachable from that location is initialized. We use these dual judgments as a base to define our validity checks for reading, writing, and moving. Reading will be more restricted than simple full initialization. Writing will be more restricted than reading. Moving will be more restricted than writing.

Reading

The restriction we add to full initialization to verify that reading a path is valid is that no part of the path or anything reachable from the path has been mutable loaned. A mutable loan promises unique access via the corresponding mutable borrowed reference, which precludes reading the loaned path directly. We can

also short-circuit our check in a two cases. If the path we want to read dereferences an immutable borrow, then we do not need to check the interior of the borrow. By its nature, the immutable borrow guarantees readability, and any mutable loaning of the interior is already forbidden. For similar reasons, when checking the memory reachable from the path we do not need to examine the interior of immutable borrows.

CONTAINS-NO-MUT-LOANS(σ)

$$\frac{\text{CNML-INT} \quad \text{NOT-MUT}(\$)}{\text{CONTAINS-NO-MUT-LOANS}(\$: \text{int})} \qquad \frac{\text{CNML-OWN} \quad \text{NOT-MUT}(\$) \quad \text{CONTAINS-NO-MUT-LOANS}(\sigma)}{\text{CONTAINS-NO-MUT-LOANS}(\$: \sim \sigma)}$$

$$\frac{\text{CNML-REFIMM} \quad \text{NOT-MUT}(\$)}{\text{CONTAINS-NO-MUT-LOANS}(\$: \& \text{imm } \sigma)} \qquad \frac{\text{CNML-REFMUT} \quad \text{NOT-MUT}(\$) \quad \text{CONTAINS-NO-MUT-LOANS}(\sigma)}{\text{CONTAINS-NO-MUT-LOANS}(\$: \& \text{mut } \sigma)}$$

CAN-READ(p, σ)

$$\frac{\text{CR-BASE} \quad \text{CONTAINS-NO-MUT-LOANS}(\sigma)}{\text{CAN-READ}(\bullet, \sigma)} \qquad \frac{\text{CR-DEOWN} \quad \text{NOT-MUT}(\$) \quad \text{CAN-READ}(p, \sigma)}{\text{CAN-READ}(*p, \$: \sim \sigma)} \qquad \frac{\text{CR-DEREFIMM} \quad \text{NOT-MUT}(\$)}{\text{CAN-READ}(*p, \$: \& \text{imm } \sigma)}$$

$$\frac{\text{CR-DEREFMUT} \quad \text{NOT-MUT}(\$) \quad \text{CAN-READ}(p, \sigma)}{\text{CAN-READ}(*p, \$: \& \text{mut } \sigma)}$$

where one would check that $x@p$ is readable by $\text{CAN-READ}(p, \Upsilon(x))$.

Writing

In order to safely write a path, no part of the path or anything reachable from the path can be loaned at all. An immutable loan implies the data is frozen, which prevents us from writing. A mutable loan implies something else has unique access to the data, which also prevents us from writing. Additionally, we cannot write to the interior of an immutable borrow, even if we can write to the borrowed reference itself. Memory reachable from that immutable borrow is frozen.

CONTAINS-NO-LOANS(σ)

$$\frac{\text{CNL-INT}}{\text{CONTAINS-NO-LOANS}(\emptyset : \text{int})} \qquad \frac{\text{CNL-OWN} \quad \text{CONTAINS-NO-LOANS}(\sigma)}{\text{CONTAINS-NO-LOANS}(\emptyset : \sim \sigma)} \qquad \frac{\text{CNL-REF} \quad \text{CONTAINS-NO-LOANS}(\sigma)}{\text{CONTAINS-NO-LOANS}(\emptyset : \& q \sigma)}$$

Note that the borrowed reference cases of have merged again after being separate for reading. We cannot guarantee that the interior of an immutable borrow will not be loaned out without actually checking. However, note that $\text{CONTAINS-NO-LOANS}(\sigma)$ implies $\text{CONTAINS-NO-MUT-LOANS}(\sigma)$.

CAN-WRITE(p, σ)

$$\frac{\text{CW-BASE} \quad \text{CONTAINS-NO-LOANS}(\sigma)}{\text{CAN-WRITE}(\bullet, \sigma)} \qquad \frac{\text{CW-DEOWN} \quad \text{CAN-WRITE}(p, \sigma)}{\text{CAN-WRITE}(*p, \emptyset : \sim \sigma)} \qquad \frac{\text{CW-DEREFMUT} \quad \text{CAN-WRITE}(p, \sigma)}{\text{CAN-WRITE}(*p, \emptyset : \& \text{mut } \sigma)}$$

Again, note that $\text{CAN-WRITE}(p, \sigma)$ implies $\text{CAN-READ}(p, \sigma)$. Rust requires that the ability to write implies the ability to read, which is reflected in our judgments.

Moving

The final operation, moving out of a path, requires ownership over the path. Moving is used when the value being read has an affine type, which prevents it from being merely copied. The old location of the value must be purged. This requires the ability to write, but since it leaves the original path uninitialized it cannot be used via mutable borrows. Doing so would violate the guarantee that borrowed references always point to initialized memory.

$\boxed{\text{CAN-MOVE}(p, \sigma)}$

$$\frac{\text{CM-BASE} \quad \text{CONTAINS-NO-LOANS}(\sigma)}{\text{CAN-MOVE}(\bullet, \sigma)}$$

$$\frac{\text{CM-DEOWN} \quad \text{CAN-MOVE}(p, \sigma)}{\text{CAN-MOVE}(*p, \emptyset : \sim \sigma)}$$

Note that $\text{CAN-MOVE}(p, \sigma)$ implies $\text{CAN-WRITE}(p, \sigma)$.

Affine Types and Copyability

Values of an affine type cannot be used more than once, e.g. unique pointers. We model this in Patina by deinitializing an affine value on use. However, we do not want to do this unnecessarily so we need to distinguish between affine and non-affine types. Affine types must be moved, but non-affine types can be merely copied. In Patina, unique pointers and borrowed mutable references are the sources of affineness. Unique pointers must be affine to preserve their uniqueness. Borrowed mutable references must be affine to preserve the unique access they promise.

$\boxed{\tau \text{ copy}}$

$$\frac{\text{C-INT}}{\text{int copy}}$$

$$\frac{\text{C-REFIMM}}{\& \ell \text{ imm } \tau \text{ copy}}$$

$\boxed{\tau \text{ move}}$

$$\frac{\text{M-OWN}}{\sim \tau \text{ move}}$$

$$\frac{\text{M-REFMUT}}{\& \ell \text{ mut } \tau \text{ move}}$$

The two judgments here are not particularly interesting without compound types, but they still create a necessary distinction. Note that these two judgments partition the set of types.

Using an L-value

Now that we can check that an operation is safe and we can determine from the type which operation to perform, we can specify how to check that using the value at an L-value is safe and track how doing so changes the shadow heap.

$\boxed{\Gamma; \Upsilon \vdash x@p : \tau \Rightarrow \Upsilon'}$

$$\frac{\text{UP-COPY} \quad \Gamma \vdash x@p : \tau \quad \tau \text{ copy} \quad \text{CAN-READ}(p, \Upsilon(x))}{\Gamma; \Upsilon \vdash x@p : \tau \Rightarrow \Upsilon}$$

$$\frac{\text{UP-MOVE} \quad \Gamma \vdash x@p : \tau \quad \tau \text{ move} \quad \text{CAN-MOVE}(p, \Upsilon(x)) \quad \text{USE}(\Upsilon, x@p, \Upsilon')}{\Gamma; \Upsilon \vdash x@p : \tau \Rightarrow \Upsilon'}$$

where $\text{USE}(\Upsilon, x@p, \Upsilon')$ modifies the shadow of $x@p$ to deinitialize the affine values.

Conjecture (L-Value Use Shadow Preservation). *If $\vdash \mathcal{L}$, $\mathcal{L} \vdash \Upsilon : \Gamma \leq L$, and $\Gamma; \Upsilon \vdash x@p : \tau \Rightarrow \Upsilon'$ then $\mathcal{L} \vdash \Upsilon' : \Gamma \leq L$.*

Justification. Using a path either leaves the shadow heap unchanged or deinitializes only a part of it, neither of which should change the types of shadows or the lifetimes of loans. Deinitializing part of the shadow heap could invalidate existing loans, but the CAN-READ and CAN-MOVE checks should rule out those cases. \square

Ensuring Loan Promises Can Be Upheld

When creating a borrowed reference, the borrow checker must ensure that is it actually possible to make the promises that the loan demands. For example, we cannot promise the interior of a borrowed reference will be valid for longer than that reference's lifetime, nor can we promise that the interior of a immutable borrowed reference will have unique access.

Validity

Both kinds of borrowed references require that their referent be initialized for the lifetime of the reference, which here we call *validity*. Variables are valid for their entire lifetime and can be loaned out for any duration not longer than that. Unique pointers inherit the lifetime of their owner. Borrowed references specify how long they are valid by their lifetime, and they can be loaned out for any duration not longer than that. Since we already know the interior of a borrowed reference is valid for the lifetime of that reference, we do not need to examine further.

$$\boxed{\Gamma; \mathcal{L}; L \vdash x@p \text{ valid for } \ell}$$

$$\begin{array}{c} \text{VF-BASE} \\ \frac{(\ell, L(x)) \in \mathcal{L}}{\Gamma; \mathcal{L}; L \vdash x@{\bullet} \text{ valid for } \ell} \end{array} \qquad \begin{array}{c} \text{VF-DEOWN} \\ \frac{\Gamma \vdash x@p : \sim \tau \quad \Gamma; \mathcal{L}; L \vdash x@p \text{ valid for } \ell}{\Gamma; \mathcal{L}; L \vdash x@*p \text{ valid for } \ell} \end{array}$$

$$\begin{array}{c} \text{VF-DEREF} \\ \frac{\Gamma \vdash x@p : \& \ell' q \tau \quad (\ell, \ell') \in \mathcal{L}}{\Gamma; \mathcal{L}; L \vdash x@*p \text{ valid for } \ell} \end{array}$$

Uniqueness

Mutable loans require the loaned memory have *unique* access. If the path being loaned has unique access, then it can give that access to the new mutable borrow. Variables, being top level, are inherently unique along with unique pointers. Mutable borrows provide unique access to their interiors for their duration. Immutable borrows are inherently not unique.

$$\boxed{\Gamma; \mathcal{L} \vdash x@p \text{ unique for } \ell}$$

$$\begin{array}{c} \text{UF-BASE} \\ \frac{}{\Gamma; \mathcal{L} \vdash x@{\bullet} \text{ unique for } \ell} \end{array} \qquad \begin{array}{c} \text{UF-DEOWN} \\ \frac{\Gamma \vdash x@p : \sim \tau \quad \Gamma; \mathcal{L} \vdash x@p \text{ unique for } \ell}{\Gamma; \mathcal{L} \vdash x@*p \text{ unique for } \ell} \end{array}$$

$$\begin{array}{c} \text{UF-DEREFMUT} \\ \frac{\Gamma \vdash x@p : \& \ell' \text{ mut } \tau \quad (\ell, \ell') \in \mathcal{L} \quad \Gamma; \mathcal{L} \vdash x@p \text{ unique for } \ell}{\Gamma; \mathcal{L} \vdash x@*p \text{ unique for } \ell} \end{array}$$

Freezability

Immutable loans require the loaned memory to be *frozen*, unchangable by anything, for the duration of the loan. When creating an immutable borrow, we need to check a slightly weaker property: *freezable*, unchangable by anything *else*. A freezable path can be mutated, but only by the path itself. If the path does not mutate itself, e.g. by loaning itself immutably, then we know it is frozen. Unique paths are inherently freezable since no other path exists. Dereferencing through an immutable borrow, while not unique, is freezable because the interior of the borrow is already guaranteed to be frozen for its lifetime.

$$\boxed{\Gamma; \mathcal{L} \vdash x@p \text{ freezable for } \ell}$$

$$\begin{array}{c} \text{FF-BASE} \\ \hline \Gamma; \mathcal{L} \vdash x@p \bullet \text{ freezable for } \ell \end{array} \qquad \begin{array}{c} \text{FF-DEOWN} \\ \Gamma \vdash x@p : \sim \tau \quad \Gamma; \mathcal{L} \vdash x@p \text{ freezable for } \ell \\ \hline \Gamma; \mathcal{L} \vdash x@*p \text{ freezable for } \ell \end{array}$$

$$\frac{\text{FF-DEREFMUT} \quad \Gamma \vdash x@p : \& \ell' \text{ mut } \tau \quad (\ell, \ell') \in \mathcal{L} \quad \Gamma; \mathcal{L} \vdash x@p \text{ freezable for } \ell}{\Gamma; \mathcal{L} \vdash x@*p \text{ freezable for } \ell}$$

$$\frac{\text{FF-DEREFIMM} \quad \Gamma \vdash x@p : \& \ell' \text{ imm } \tau \quad (\ell, \ell') \in \mathcal{L}}{\Gamma; \mathcal{L} \vdash x@*p \text{ freezable for } \ell}$$

Note how $\Gamma; \mathcal{L} \vdash x@p$ unique for ℓ implies $\Gamma; \mathcal{L} \vdash x@p$ freezable for ℓ . This implies that anything we can borrow mutably we can also borrow immutably, which accurately reflects Rust.

Expression Typing

With our model of the borrow checker in hand, we can finally specify the typing of expressions.

$$\boxed{\Gamma; \mathcal{L}; L; \Upsilon \vdash e : \tau; \Upsilon'}$$

Integer constants are trivial and do not change the shadow heap.

$$\frac{\text{ET-INT}}{\Gamma; \mathcal{L}; L; \Upsilon \vdash z : \text{int}; \Upsilon}$$

Using the value at an L-value simply relies on our earlier judgment.

$$\frac{\text{ET-USE} \quad \Gamma; \Upsilon \vdash x@p : \tau \Rightarrow \Upsilon'}{\Gamma; \mathcal{L}; L; \Upsilon \vdash x@p : \tau; \Upsilon'}$$

Creating an immutable borrowed reference requires checking that the immutable loan guarantees will be upheld for the entire duration of the loan. This involves checking that we can read the L-value now, and that we can guarantee its validity and freezability for the duration of the loan. We record the new loan in the output shadow heap.

$$\frac{\text{ET-REFIMM} \quad \Gamma \vdash x@p : \tau \quad \text{CAN-READ}(p, \Upsilon(x)) \quad \Gamma; \mathcal{L}; L \vdash x@p \text{ valid for } \ell \quad \Gamma; \mathcal{L} \vdash x@p \text{ freezable for } \ell \quad \text{RECORD}(\Upsilon, x@p, (\ell, \text{imm}), \Upsilon')}{\Gamma; \mathcal{L}; L; \Upsilon \vdash \& \ell \text{ imm } x@p : \& \ell \text{ imm } \tau; \Upsilon'}$$

Similarly, creating a mutable borrowed reference requires checking the mutable loan guarantees. Instead of reading, we must be able to write. Instead of freezability, we require uniqueness.

$$\text{ET-REFMUT} \quad \frac{\Gamma \vdash x@p : \tau \quad \text{CAN-WRITE}(p, \Upsilon(x)) \quad \Gamma; \mathcal{L}; L \vdash x@p \text{ valid for } \ell \quad \Gamma; \mathcal{L} \vdash x@p \text{ unique for } \ell \quad \text{RECORD}(\Upsilon, x@p, (\ell, \text{mut}), \Upsilon')}{\Gamma; \mathcal{L}; L; \Upsilon \vdash \& \ell \text{ mut } x@p : \& \ell \text{ mut } \tau; \Upsilon'}$$

Conjecture (Expression Shadow Preservation). *If $\vdash \mathcal{L}$, $\mathcal{L} \vdash \Upsilon : \Gamma \leq L$, and $\Gamma; \mathcal{L}; L; \Upsilon \vdash e : \tau; \Upsilon'$ then $\mathcal{L} \vdash \Upsilon' : \Gamma \leq L$.*

Justification. This is trivial in the integer case and follows from L-Value Use Shadow Preservation in the L-value use case, so we only need to discuss the borrow creation cases. In both borrow creation cases, the only change to the shadow heap is the addition of one new loan.

When creating an immutable borrow, we know from CAN-READ that no mutable loan exists on way to the path we are borrowing or in memory reachable from it. The shadow of the interior of the loaned path is now controlled by the immutable loan of the path, but since there were no mutable loans there before and it has not changed it should still be well-formed. Similarly, the shadows on the way to the path have not changed directly and they should not be affected by the addition of an immutable loan in their interior. The main concern is whether the addition of the new immutable loan to the bank will result in a well-formed bank. We know from the validity check that the lifetime of the loan is within the lifetime bound of the bank. If the bank is empty or the lifetime of the old loan is not shorter than the lifetime of the new loan, then the bank will still be well-formed. We could restrict RECORD to only succeed in this situation.

When creating a mutable borrow, we know from CAN-WRITE that no loan exists on the way to the path or in memory reachable from it. So similar to the immutable case, the shadow of the interior of the path will still be well-formed even though it is now controlled by a mutable loan. The shadows on the way to the path are still well-formed since they have no loans to begin with. As before, the only question is whether the bank of the path will remain well-formed. We know from CAN-WRITE that the bank is empty, so the new bank, consisting of just one mutable loan, will be well-formed. \square

Using a Layout

The runtime analogue of using an L-value is using a layout. Affine values are replaced by void and copyable values are left untouched. It is not defined for void, which prevents expression evaluation from utilizing uninitialized memory. A program that tries to do so will get stuck.

USE : Layout \rightarrow Layout

$$\begin{aligned} \text{USE}(z) &= z \\ \text{USE}(\sim \alpha) &= \text{void} \\ \text{USE}(\& \text{imm } \alpha) &= \& \text{imm } \alpha \\ \text{USE}(\& \text{mut } \alpha) &= \text{void} \end{aligned}$$

Note that $\text{USE}(l) = l'$ implies $\text{DROPPABLE}(l')$. Once we use a value, we know we can safely deallocate it.

Expression Evaluation

Evaluating expressions is straightforward. The only unusual aspect is the case of using the value at an L-value. To prevent copies of affine values, we must overwrite the original value with the used version. If it is a copyable value, then it will remain unchanged. If it is an affine value, then it will be replaced with void. Note that this means the heap resulting from expression evaluation is either unchanged or a more deinitialized version of the starting heap. Also note that the result heap may temporarily contain a leak

when considered alone, but not when considered together with the result layout.

$$\boxed{V; H \vdash e \rightarrow l; H'}$$

$$\begin{array}{c} \text{EE-INT} \\ \hline V; H \vdash z \rightarrow z; H \end{array} \qquad \begin{array}{c} \text{EE-REF} \\ \hline V; H \vdash x@p \rightarrow \alpha \\ \hline V; H \vdash \& \ell \ q \ x@p \rightarrow \& \ q \ \alpha; H \end{array}$$

$$\begin{array}{c} \text{EE-USE} \\ \hline V; H \vdash x@p \rightarrow \alpha \quad \text{READ}(H, \alpha, l) \quad \text{WRITE}(H, \alpha, \text{USE}(l), H') \\ \hline V; H \vdash x@p \rightarrow l; H' \end{array}$$

Expression Preservation

Having defined both evaluation and typing for expressions, we can finally state the preservation lemma for expressions. Assuming all input contexts are well-formed and given a well typed expression that evaluates to a layout, then that layout has the same type as the expression, the resulting shadow heap is well-formed, and the resulting heap is consistent with the resulting shadow heap.

Conjecture (Expression Preservation). *If $\vdash H : \Sigma, \Gamma; V \vdash \Sigma, \vdash \mathcal{L}, \mathcal{L} \vdash \Upsilon : \Gamma \leq L, V \vdash H : \Upsilon, \Gamma; \mathcal{L}; L; \Upsilon \vdash e : \tau; \Upsilon'$, and $V; H \vdash e \rightarrow l; H'$ then $\Sigma \vdash l : \tau, \vdash H' : \Sigma, \mathcal{L} \vdash \Upsilon' : \Gamma \leq L$, and $V \vdash H' : \Upsilon'$.*

Justification. This is trivial in the integer case. In the reference case, the heap does not change. From Expression Shadow Preservation, we get the well-formedness of the new shadow heap, which has been extended with one new loan. Since the initialization information in the shadow heap has not changed, it should still model the heap correctly. That the type of the layout matches the type of the expression is simple to establish.

In the use case, we again get the well-formedness of the new shadow heap from Expression Shadow Preservation. The type of the layout we can get via Path Type Preservation and Read Safety. The write we perform on the heap does not change the type of anything, so the new heap will still typecheck under the heap type. In the new shadow heap, the used path was deinitialized (if affine). In the new heap, the used path was voided (if affine). The new heap should be consistent with the new shadow heap. \square

The Layout Resulting from Evaluation is Fully Initialized

We should also note that the resulting layout is fully initialized. This is trivial in the integer case. In the reference and the use case we rely on typing judgments that asserted full initialization and the correspondence of the shadow heap with the runtime heap.

Expression Progress

The progress lemma for expressions is surprisingly straightforward. Since we do not nest expressions, no induction is necessary. Each case constructs the resulting values in a straightforward manner.

Conjecture (Expression Progress). *If $\vdash H : \Sigma, \Gamma; V \vdash \Sigma, \vdash \mathcal{L}, \mathcal{L} \vdash \Upsilon : \Gamma \leq L, V \vdash H : \Upsilon$, and $\Gamma; \mathcal{L}; L; \Upsilon \vdash e : \tau; \Upsilon'$ then there exists a layout l and a heap H' such that $V; H \vdash e \rightarrow l; H'$.*

Justification. In the integer case this is trivial. In both reference cases, we can rely on Path Progress to evaluate the L-value and use the address to construct the resulting layout, and the heap does not change. In the use case, we can rely on Path Progress to evaluate the L-value to an address. Then we can rely on Read Safety to find the layout. Since we could read sufficiently, we can also write sufficiently so we find the heap. \square

Statements

We have finally reached the outermost shell of Patina! The statements layer is primarily about heap manipulation and chaining expressions. It mostly just relies on judgments we have already defined. `skip` and `s; s` are simply for chaining together statements. `x@p ← e` is for initializing a path with the result of an expression. `x@p ← new x@p` is similar, but allocates the result on the heap and initializes the path with a new unique pointer. Treating this as a statement rather than an expression simplifies the expression preservation and progress lemmas. Note that both versions are for *initialization* and not *assignment*. We only allow initialization of droppable paths, which we guarantee will not leak memory. `x@p ↔ x@p` will perform a shallow swap of two writable paths. This is the only way to actually mutate through a mutable borrow since we do not allow the interior of a borrowed reference to be uninitialized even temporarily. `free x@p` performs a shallow free and is the inverse of `new`. `push ℓ x τ s` will create a new stack variable. `pop x s` will pop that stack variable once the statement that requires it finishes. It is for evaluation purposes and should not appear in user code.

Statements $s ::= \text{skip} \mid s; s \mid x@p \leftarrow e \mid x@p \leftarrow \text{new } x@p \mid x@p \leftrightarrow x@p \mid \text{free } x@p \mid \text{push } \ell \ x \ \tau \ s \mid \text{pop } x \ s$

Droppable Shadows

Earlier we defined `DROPPABLE(l)`, which specified layouts that were safe to deallocate. Now we define an equivalent notion for shadows, which we will use to ensure only droppable layouts are deallocated. We also require a droppable shadow to be unborrowed, which will prevent us from creating dangling borrowed references.

σ DROPPABLE

DS-UNINIT
 $\frac{}{\emptyset : \text{uninit DROPPABLE}}$

DS-INT
 $\frac{}{\emptyset : \text{int DROPPABLE}}$

DS-REF
 CONTAINS-NO-LOANS(σ)
 $\frac{}{\emptyset : \& \ q \ \sigma \ \text{DROPPABLE}}$

Subtyping

The only subtyping in Patina is lifetime subtyping. Borrowed references are contravariant in their lifetimes. This means that a reference for a longer lifetime is a subtype of a reference for a shorter lifetime, which is the behavior we want. The reference of a longer lifetime is usable anywhere the reference of a shorter lifetime is. Since mutable references can both read and write, they are invariant in their type argument. Immutable references can only read, so they are covariant in their type argument.

$\mathcal{L} \vdash \tau_0 <: \tau_1$

SUB-INT
 $\frac{}{\mathcal{L} \vdash \text{int} <: \text{int}}$

SUB-OWN
 $\frac{\mathcal{L} \vdash \tau_1 <: \tau_2}{\mathcal{L} \vdash \sim \tau_1 <: \sim \tau_2}$

SUB-REFIMM
 $\frac{(\ell_2, \ell_1) \in \mathcal{L} \quad \mathcal{L} \vdash \tau_1 <: \tau_2}{\mathcal{L} \vdash \& \ \ell_1 \ \text{imm} \ \tau_1 <: \& \ \ell_2 \ \text{imm} \ \tau_2}$

SUB-REFMUT
 $\frac{(\ell_2, \ell_1) \in \mathcal{L}}{\mathcal{L} \vdash \& \ \ell_1 \ \text{mut} \ \tau <: \& \ \ell_2 \ \text{mut} \ \tau}$

Statement Typing

$$\boxed{\Gamma; \mathcal{L}; L; \Upsilon \vdash s : \Upsilon'}$$

The skip statement does nothing so it is easy to check.

$$\frac{\text{ST-SKIP}}{\Gamma; \mathcal{L}; L; \Upsilon \vdash \text{skip} : \Upsilon}$$

Sequencing two statement simply sequences typechecking by threadings the shadow heap. Note that since Γ is not changing, the statements can only change the initialization and loan information of the shadow heap.

$$\frac{\text{ST-SEQ} \quad \Gamma; \mathcal{L}; L; \Upsilon_0 \vdash s_1 : \Upsilon_1 \quad \Gamma; \mathcal{L}; L; \Upsilon_1 \vdash s_2 : \Upsilon_2}{\Gamma; \mathcal{L}; L; \Upsilon_0 \vdash s_1; s_2 : \Upsilon_2}$$

Swapping requires only that the two L-values have the same type and that they both be writable. Since the writability guarantees no part is loaned, the shadow heap is not changed.

$$\frac{\text{ST-SWAP} \quad \begin{array}{l} \Gamma \vdash x_1 @ p_1 : \tau \quad \text{CAN-WRITE}(p_1, \Upsilon(x_1)) \\ \Gamma \vdash x_2 @ p_2 : \tau \quad \text{CAN-WRITE}(p_2, \Upsilon(x_2)) \end{array}}{\Gamma; \mathcal{L}; L; \Upsilon \vdash x_1 @ p_1 \leftrightarrow x_2 @ p_2 : \Upsilon}$$

Initializing with a value requires that the path be droppable so that nothing will be leaked. It also requires the expression type to be a subtype of the path type. As discussed above, this refers to lifetimes: the subtype lifetime is greater than or equal to the lifetime of the supertype. This requirement guarantees that the value stored at the path will be valid for at least as long as the path is valid. Finally, we initialize the path in the output shadow heap.

$$\frac{\text{ST-SET} \quad \begin{array}{l} \Gamma \vdash x @ p : \tau_p \quad \Gamma; \mathcal{L}; L; \Upsilon_0 \vdash e : \tau_e; \Upsilon_1 \quad \mathcal{L} \vdash \tau_e <: \tau_p \\ \text{SHALLOW}(\Upsilon_1, x @ p, \sigma) \quad \sigma \text{ DROPPABLE} \quad \text{INIT}(\Upsilon_1, x @ p, \Upsilon_2) \end{array}}{\Gamma; \mathcal{L}; L; \Upsilon_0 \vdash x @ p \leftarrow e : \Upsilon_2}$$

Initializing with a pointer is similar to initializing with a value. Except instead of expression typing, we must check that we can use the path on the right hand side.

$$\frac{\text{ST-NEW} \quad \begin{array}{l} \Gamma \vdash x_l @ p_l : \sim \tau_l \quad \Gamma; \Upsilon_0 \vdash x_r @ p_r : \sim \tau_r \Rightarrow \Upsilon_1 \quad \mathcal{L} \vdash \sim \tau_r <: \sim \tau_l \\ \text{SHALLOW}(\Upsilon_1, x_l @ p_l, \sigma) \quad \sigma \text{ DROPPABLE} \quad \text{INIT}(\Upsilon_1, x_l @ p_l, \Upsilon_2) \end{array}}{\Gamma; \mathcal{L}; L; \Upsilon_0 \vdash x_l @ p_l \leftarrow \text{new } x_r @ p_r : \Upsilon_2}$$

Freeing the memory of a unique pointer requires checking that that memory is droppable. We then mark it used in the output shadow heap, which since this a unique pointer will replace the shadow with uninit.

$$\frac{\text{ST-FREE} \quad \Gamma \vdash x @ p : \sim \tau \quad \text{SHALLOW}(\Upsilon, x @ p, \sim \sigma) \quad \sigma \text{ DROPPABLE} \quad \text{USE}(\Upsilon, x @ p, \Upsilon')}{\Gamma; \mathcal{L}; L; \Upsilon \vdash \text{free } x @ p : \Upsilon'}$$

Pushing a new stack variable involves checking that the body is well typed under the appropriately extended contexts. Newly pushed variables start uninitialized. By the end of the body, the variable must be droppable so that popping it will not leak memory. When checking the body, we extend the lifetime relation with the new lifetime by adding the new lifetime as the smallest lifetime in the relation (it is related to everything else). Afterwards, we must filter everything involving this lifetime from the shadow heap. This should only consist of removing loans from banks. We also "deallocate" the stack variable from the shadow heap by restricting the domain to the domain of the input.

$$\begin{array}{c}
\text{ST-PUSH} \\
\text{NEW-SMALLEST-LIFETIME}(\mathcal{L}, \ell, \mathcal{L}') \\
\Gamma, x \mapsto \tau; \mathcal{L}'; L, x \mapsto \ell; \Upsilon_0, x \mapsto \text{uninit} \vdash s : \Upsilon_1 \\
\Upsilon_1(x) \text{ DROPPABLE} \quad \text{FILTER}(\Upsilon_1, \ell, \Upsilon_2) \\
\hline
\Gamma; \mathcal{L}; L; \Upsilon_0 \vdash \text{push } \ell \ x \ \tau \ s : \Upsilon_2|_{\text{dom } \Upsilon_0}
\end{array}$$

Popping a stack variable is similar, but we do not need to worry about extending contexts. Instead, we must make sure to remove the stack variable when we restrict the domain of the output shadow heap.

$$\begin{array}{c}
\text{ST-POP} \\
\Gamma; \mathcal{L}; L; \Upsilon_0 \vdash s : \Upsilon_1 \\
\Upsilon_1(x) \text{ DROPPABLE} \quad \text{FILTER}(\Upsilon_1, \ell, \Upsilon_2) \\
\hline
\Gamma; \mathcal{L}; L; \Upsilon_0 \vdash \text{pop } x \ s : \Upsilon_2|_{\text{dom } \Upsilon_0 - \{x\}}
\end{array}$$

One interesting thing to note here is that the domain of the input and output shadow heaps is always the same except for the `pop` case, where it decreases by one. Since `pop` does not appear in user code, this means that if a program starts with no memory allocated and that program type checks, then the program frees all the memory it allocates. Assuming of course that evaluation preserves the shadow heap.

Statement Evaluation

Statement evaluation goes as you would expect. We use a new helper judgment `FRESH(α)` that generates a unused allocation. As before, we utilize domain restriction to "deallocate". Other than that, we simply use path and expression evaluation combined with `READ` and `WRITE`.

$$\boxed{V; H \vdash s \rightarrow V'; H' \vdash s'}$$

$$\begin{array}{c}
\text{SE-SKIPSEQ} \\
\hline
V; H \vdash \text{skip}; s \rightarrow V; H \vdash s \\
\\
\text{SE-STEPSEQ} \\
\hline
V; H \vdash s_1 \rightarrow V'; H' \vdash s'_1 \\
\hline
V; H \vdash s_1; s_2 \rightarrow V'; H' \vdash s'_1; s'_2 \\
\\
\text{SE-SET} \\
\hline
V; H_0 \vdash e \rightarrow l; H_1 \quad V; H_1 \vdash x@p \rightarrow \alpha \quad \text{WRITE}(H_1, \alpha, l, H_2) \\
\hline
V; H_0 \vdash x@p \leftarrow e \rightarrow V; H_2 \vdash \text{skip} \\
\\
\text{SE-NEW} \\
\hline
V; H_0 \vdash x_r@p_r \rightarrow \alpha_r \quad \text{READ}(H_0, \alpha_r, l_r) \quad \text{WRITE}(H_0, \alpha_r, \text{USE}(l_r), H_1) \\
V; H_1 \vdash x_l@p_l \rightarrow \alpha_l \quad \text{FRESH}(\alpha_n) \quad \text{WRITE}(H_1, \alpha_l, \sim \alpha_n, H_2) \\
\hline
V; H_0 \vdash x_l@p_l \leftarrow \text{new } x_r@p_r \rightarrow V; H_2, \alpha_n \mapsto l_r \vdash \text{skip} \\
\\
\text{SE-SWAP} \\
\hline
V; H_0 \vdash x_l@p_l \rightarrow \alpha_l \quad V; H_0 \vdash x_r@p_r \rightarrow \alpha_r \\
\text{READ}(H_0, \alpha_l, l_l) \quad \text{READ}(H_0, \alpha_r, l_r) \\
\text{WRITE}(H_0, \alpha_l, l_r, H_1) \quad \text{WRITE}(H_1, \alpha_r, l_l, H_2) \\
\hline
V; H_0 \vdash x_l@p_l \leftrightarrow x_r@p_r \rightarrow V; H_2 \vdash \text{skip} \\
\\
\text{SE-FREE} \\
\hline
V; H_0 \vdash x@p \rightarrow \alpha \quad \text{READ}(H_0, \alpha, \sim \alpha') \quad \text{WRITE}(H_0, \alpha, \text{USE}(\sim \alpha'), H_1) \\
\hline
V; H_0 \vdash \text{free } x@p \rightarrow V; H_1|_{\text{dom } H_1 - \{\alpha'\}} \vdash \text{skip} \\
\\
\text{SE-PUSH} \\
\hline
\text{FRESH}(\alpha) \\
\hline
V; H \vdash \text{push } \ell \ x \ \tau \ s \rightarrow V, x \mapsto \alpha; H, \alpha \mapsto \text{void} \vdash \text{pop } x \ s \\
\\
\text{SE-POPSTEP} \\
\hline
V; H \vdash s \rightarrow V'; H' \vdash s' \\
\hline
V; H \vdash \text{pop } x \ s \rightarrow V'; H' \vdash \text{pop } x \ s' \\
\\
\text{SE-POPSKIP} \\
\hline
V; H \vdash \text{pop } x \ \text{skip} \rightarrow V|_{\text{dom } V - \{x\}}; H|_{\text{dom } H - \{V(x)\}} \vdash \text{skip}
\end{array}$$

Statement Preservation

Statement preservation is the assertion that given all our well-formed contexts, a well-typed statement, and the fact that that statement evaluates to another statement, then we can also type the result statement and both statements produce the same output shadow heap, i.e. they effectively end in the same way. Put another way, this lemma asserts that memory safe programs step to memory safe programs.

It is difficult to specify how the new versions of contexts relate to the old versions. They may be larger or smaller depending upon the statement in question. We do not specify this relationship currently. Instead we just rely on the new variable map, heap, and statement to enforce the relationship.

Conjecture (Statement Preservation). *If $\vdash H : \Sigma, \Gamma; V \vdash \Sigma, \vdash \mathcal{L}, \mathcal{L} \vdash \Upsilon : \Gamma \leq L, V \vdash H : \Upsilon, \Gamma; \mathcal{L}; L; \Upsilon \vdash s : \Upsilon_0$, and $V; H \vdash s \rightarrow V'; H' \vdash s'$ then there exist $\Sigma', \Gamma', \mathcal{L}', L'$, and Υ' such that $\vdash H' : \Sigma', \Gamma'; V' \vdash \Sigma', \vdash \mathcal{L}', \mathcal{L}' \vdash \Upsilon' : \Gamma' \leq L', V' \vdash H' : \Upsilon'$, and $\Gamma'; \mathcal{L}'; L'; \Upsilon' \vdash s' : \Upsilon_0$.*

Justification. The skip-seq case is trivial. The step-seq case follows from induction. In the value initialization (set) case, only the shadow heap and the heap are changing. From Expression Preservation, we know that the heap and shadow heap are still aligned after evaluating the expression. The WRITE will initialize the path, which is reflected in the shadow heap. The pointer initialization (new) case is similar. Since the components of a swap must have the same shadow, the modified heap will still be well-formed. We required that the memory deallocated by free was droppable, which should mean that the heap is still well-formed once we void the pointer. Push steps immediately to pop and only extends contexts, which should still be

well-formed. The pop-step case follows from induction. In the pop-skip case, similar to the free case, we know that the variable is droppable, which should mean the resulting contexts are well-formed. \square

Statement Progress

Statement progress is essentially the assertion that a well-typed statement will never attempt to read uninitialized memory, free borrowed memory, and so on. Essentially, that well-typed programs have sufficient memory safety to take a step. The restriction to statements that end with an empty shadow heap allows us to avoid the cases with leaked memory.

Conjecture (Statement Progress). *If $\vdash H : \Sigma, \Gamma; V \vdash \Sigma, \vdash \mathcal{L}, \mathcal{L} \vdash \Upsilon : \Gamma \leq L, V \vdash H : \Upsilon, \Gamma; \mathcal{L}; L; \Upsilon \vdash s : \emptyset$ then there exists V', H' , and s' such that $V; H \vdash s \rightarrow V'; H' \vdash s'$ or $s = \text{skip}$.*

Justification. The skip case is immediate. The sequence case is either immediate or follows from induction. The swap case and the new case follow from path progress and read safety. The set case is similar, but also requires expression progress. The push case is immediate. The pop case is either immediate or follows from induction. \square

Patina Soundness

If our conjectures hold, then we have established soundness for Patina. This means that a well-typed statement that ends with an empty shadow heap will not leak memory, not access uninitialized memory, not create aliased, mutable memory, and so on. In effect, well-typed programs are statically guaranteed to be memory safe.

Future Work

The first extension to the existing work should be to prove our conjectures. This is necessary to put us on a firm footing for other future work, and to ensure we have actually included everything necessary in our model.

The expression and statement layers should be extended to match the path layer. Product types are straightforward to add to the borrow checker and will likely not pose much trouble. Sum types may be more difficult. In particular, matching by reference allows creation of references to the payload of a sum type. The borrow checker must ensure that the sum is not mutated while such a reference exists. Work on adding existential types to imperative languages can likely be adapted for this purpose [7]. Extending the model with recursive types will require the addition of top level (non-closure) recursive functions. Without the capability those functions give us it is impossible to free recursive values. Those functions will necessitate other additions, e.g. lifetime parameters. Type parameters likely do not have significant soundness implications by themselves, but may in their interaction with later features.

Until recently, Rust's closures have been in flux. As of this writing, Rust supports both stack and heap closures and also supports different manners of access to the environment (owning, immutable reference, or mutable reference). In the past, closures have been a common source of borrow checker bugs, which likely means that proving soundness will be difficult. However, it also means that it is one of the most essential features to add to Patina. It would likely be profitable to look to Cyclone [8], which took great care to ensure references could not escape via closures or existential types.

Traits, which are Rust's version of typeclasses or interfaces, may have similar issues. Traits can be used to create what is effectively an object or an existential type. Since objects and closures are so similar, it is likely that proving trait objects sound will face similar difficulties. Rust's closures are actually implemented as syntactic sugar for certain trait implementations in the current version. It is unclear whether this means formalization of closures will be easier or whether formalization of traits will be harder.

Recently, Rust has added some more advanced type system features. In particular, associated types and higher rank lifetime polymorphism are of interest. Associated types are effectively functions from types to types. Neither of these features will likely impact soundness, but they should still be carefully modeled.

There are many potential features that we could explore using our model. Several additional versions of borrowed references has been purposed, e.g. write-only references or references that allow moving out as long as something is moved back in before the reference is released. Adding this to Patina would likely result in a more general borrow checking framework.

One oft-requested feature is the ability to use static values in type parameters. This effectively gives Rust a weak form of dependent typing. The consequences for soundness in a affine dependently typed language with regions are likely significant, and the difficulty in proving them even more so.

Our model made several changes to how systems worked in the Rust compiler. It would be interesting to compare how the model’s version compares with the original in terms of performance, ease of understanding, and maintainability.

Related Work

Safe dialects of C, unique pointers, borrowed references, regions, and other topics relevant to Rust each have long and full histories. Rather than attempt to catalogue them in detail and compare Rust with each of its ancestors, we will limit our discussion to the most closely related work we are aware of.

Both ML Kit [12, 13] and Cyclone [8] heavily influenced the design of Rust. However, both of those languages, along with most research in the area, uses regions for memory management. In contrast, Rust’s memory is managed by affine types (unique ownership), and uses regions (lifetimes in Rust) more for ensuring the safety of borrowed references than for managing memory. Cyclone’s region subtyping is the source of Rust’s lifetime subtyping and so they ensure references do not escape in a similar fashion. Like Cyclone, Rust prevents access to uninitialized memory, but Rust goes further than Cyclone by statically preventing dereferences of null pointers. Unlike Cyclone, Rust does not have first class regions, but Rust also lacks Cyclone’s dynamic regions, which lessens the need for first class regions.

Cyclone later added unique pointers and a notion of borrowing [11, 9]. Cyclone’s version does not prevent memory leaks, which are possible both from freeing a unique pointer that contains another live unique pointer and from overwriting a live unique pointer. Rust statically guarantees no memory leaks in both of these situations. Cyclone requires an explicit swap operation for using unique versions of pointers that do not admit null values. Rust will allow any of its pointers to be null, but guarantees that they will not be used while null. Rust also does not require explicit swaps; Patina uses them, but that is an artifact of the model. Cyclone considers unique pointers “consumed” while a borrow exists, whereas Rust will allow the borrowed unique pointer to be used in certain situations (namely, reading through an immutably borrowed unique pointer is allowed). Cyclone supports abstraction over the aliasability and kind (box versus region) of a pointer, but Rust does not currently support any such abstraction over pointer types.

Overall, Rust and Cyclone are each more expressive in different ways. Cyclone’s dynamic regions make it much easier to specify the lifetime of a particular heap data structure, and they allow for heap structures that are difficult or impossible to express in Rust, such as a doubly linked list. Rust supports far more detailed reasoning about when memory should be freed due to its flow-sensitivity, and Rust offers safer and more convenient usage of linear data structures, particularly with regard to freeing them. Cyclone still has much Rust can learn from, e.g. a mechanism similar to dynamic regions could be used to support custom allocators in Rust.

Walker and Watkins discuss how regions and linear types could be combined [14]. Their model supports first class regions, like Cyclone, but unlike Rust. It also introduces the idea of linear regions and provides a mechanism for temporarily aliasing them. This could be a fruitful ground for Rust to explore when looking to support custom allocators. However, their model does not support mutation, which is the source of most of the complexity of Rust’s borrow checker, and this may impact its applicability to Rust.

Gordon et al. [6] introduce a static method for ensuring data-race freedom by enforcing that at most one mutable reference to an object exists, which is a similar task to Rust’s effort to avoid aliased, mutable memory.

Their work supports read-only references, which cannot be used to mutate, but do not forbid mutation. Since these are by definition aliased, mutable memory, Rust has no equivalent. However, their immutable reference is similar to Rust's immutable borrowed references, and their isolated reference is somewhat similar to Rust's unique pointers and mutable borrowed references. Isolated references can be converted to writeable references and back, which allows for behavior similar to a mutable borrowed reference. Isolated references can also be converted into immutable references, which is analogous to immutably reborrowing a mutably borrowed reference. Their correspondence with unique pointers is similar. Gordon, Ernst, and Grossman [5] later propose a general notion of rely-guarantee references, which treat aliases similar to different threads of control and use methods from rely-guarantee program logics to ensure safety. Rust's references can likely be encoded in this general framework.

Boyland's alias burying [2] has a similar goal to Rust, i.e. using static checks rather than destructive reads, but differs in many details. Boyland defines a strong version of the uniqueness invariant (true uniqueness) and a weak version (allowing borrowed aliases), and then develops a system based on maintaining the strong invariant. Rust meets only the weak uniqueness invariant, but it also guarantees that aliased, mutable memory never exists, which allows for similar control over mutation as the strong uniqueness invariant. The general idea of alias burying is to allow for aliases, but to invalidate them when a unique field is read. Rust ensures that a unique value will not be read if (mutable) aliases exist. Rust's borrowed references are also more first-class than Boyland's, which cannot be stored in data structures or returned from functions.

Clarke and Wrigstad's external uniqueness [4] has a notion of ownership similar to Rust's. Both have recursive ownership; however, Rust uses it for memory management rather than for access control. Unlike Rust, unique pointers are nullified on use. The Patina model nullifies on use, but soundness ensures it does not matter in practice. There is also a fresh owner notion for borrowed references that accomplishes the same thing as Rust's lifetimes, i.e. preventing escape. Unlike Rust, Clarke and Wrigstad's borrowed references are both copyable and mutable, which means they cannot make the same safety guarantees as Rust, i.e. no access to uninitialized memory. They list several options about how borrowed unique pointers can be handled, but Rust does not cleanly fit into these categories. Rust discriminates on the manner of borrowing and the operation to be performed to determine whether it is safe to use the borrowed unique pointer, but Clarke and Wrigstad suggest that a system must either allow all uses or forbid all uses. Their model also has problems with immutable (final) variables and unique versions of non-nullable pointers because they denote consumption via nullification and the finality prevents that. Rust, however, has both immutable variables and all of its pointers are non-null, but it has no issues mixing these with uniqueness. This is because Rust's ban on null pointers is not a ban on pointers being null, but rather on pointers being used while null. Additionally, Rust's immutable variables' usability can change if a unique value is consumed, which allows Rust to render an immutable variable unusable thereafter. Clarke and Wrigstad speculate that alias burying would enable them to support this combination, which is somewhat true. Rust does manage this via static analysis, but not via alias burying.

Boyland, Noble, and Retert's capabilities for sharing [3] provides a general framework for describing uniqueness and read-only concepts related to pointers. However, Rust's pointers do not quite fit into their capability system. The issue involves immutably loaned unique pointers and immutably loaned mutable borrowed references. In Rust, both allow reading through the pointer, but do not allow reading the pointer itself, which would require a move. The capabilities for sharing system supports only whole object permissions, which prevents us from specifying this case of Rust's pointers in the system.

Potanin et al. [10] survey the various ways of adding immutability to object-oriented languages and how it impacts aliasing, including Javari, OIGJ, and Joe₃. These systems usually add immutable objects and/or read-only references, but they do not provide the no aliased, mutable memory guarantee of Rust. As discussed earlier, read-only references are antithetical to Rust's memory safety invariant. Immutable objects are similar to Rust's immutably borrowed references, but the OO systems often go further by allowing immutability at the field level. Rust does not allow programmers to specify the mutability of particular fields, but immutable borrows can provide a temporary guarantee that is similar. Several of the OO systems have a concept of ownership, but it is primarily about encapsulation and access control. Rust handles encapsulation via its module system rather than via the type system or static analysis. As usual with OO

languages, they rely on garbage collection for memory management, which is very unlike Rust.

Acknowledgements

Thanks to Dan Grossman, Nicholas Matsakis, and Zach Tatlock for their advice on this work and on this paper. Thanks to the Mozilla Corporation for financially supporting this work.

References

- [1] The Rust Programming Language. <http://www.rust-lang.org>.
- [2] John Boyland. Alias burying: Unique variables without destructive reads. *Software: Practice and Experience*, 31(6):533–553, 2001.
- [3] John Boyland, James Noble, and William Retert. Capabilities for sharing. In JørgenLindskov Knudsen, editor, *ECOOP 2001 – Object-Oriented Programming*, volume 2072 of *Lecture Notes in Computer Science*, pages 2–27. Springer Berlin Heidelberg, 2001.
- [4] Dave Clarke and Tobias Wrigstad. External uniqueness is unique enough. In Luca Cardelli, editor, *ECOOP 2003 Object-Oriented Programming*, volume 2743 of *Lecture Notes in Computer Science*, pages 176–200. Springer Berlin Heidelberg, 2003.
- [5] Colin S. Gordon, Michael D. Ernst, and Dan Grossman. Rely-guarantee references for refinement types over aliased mutable data. In *Proceedings of the 34th ACM SIGPLAN Conference on Programming Language Design and Implementation, PLDI '13*, pages 73–84, New York, NY, USA, 2013. ACM.
- [6] Colin S. Gordon, Matthew J. Parkinson, Jared Parsons, Aleks Bromfield, and Joe Duffy. Uniqueness and reference immutability for safe parallelism. *SIGPLAN Not.*, 47(10):21–40, October 2012.
- [7] Dan Grossman. Existential types for imperative languages. In Daniel Le Mtayer, editor, *Programming Languages and Systems*, volume 2305 of *Lecture Notes in Computer Science*, pages 21–35. Springer Berlin Heidelberg, 2002.
- [8] Dan Grossman, Greg Morrisett, Trevor Jim, Michael Hicks, Yanling Wang, and James Cheney. Region-based memory management in cyclone. *SIGPLAN Not.*, 37(5):282–293, May 2002.
- [9] Michael Hicks, Greg Morrisett, Dan Grossman, and Trevor Jim. Experience with safe manual memory-management in cyclone. In *Proceedings of the 4th International Symposium on Memory Management, ISMM '04*, pages 73–84, New York, NY, USA, 2004. ACM.
- [10] Alex Potanin, Johan Östlund, Yoav Zibin, and Michael D. Ernst. Immutability. In *Aliasing in Object-Oriented Programming*, volume 7850 of *LNCS*, pages 233–269. Springer-Verlag, April 2013.
- [11] Nikhil Swamy, Michael Hicks, Greg Morrisett, Dan Grossman, and Trevor Jim. Safe manual memory management in cyclone. *Science of Computer Programming*, 62(2):122 – 144, 2006. Special Issue: Five perspectives on modern memory management - Systems, hardware and theory.
- [12] Mads Tofte, Lars Birkedal, Martin Elsmann, Niels Hallenberg, Tommy Højfeldt Olesen, and Peter Sestoft. Programming with regions in the MLKit (revised for version 4.3.0). Technical report, IT University of Copenhagen, Denmark, January 2006.
- [13] Mads Tofte and Jean-Pierre Talpin. Region-based memory management. *Information and Computation*, 132(2):109 – 176, 1997.

- [14] David Walker and Kevin Watkins. On regions and linear types (extended abstract). In *Proceedings of the Sixth ACM SIGPLAN International Conference on Functional Programming, ICFP '01*, pages 181–192, New York, NY, USA, 2001. ACM.