

Simple Reference Immutability for System $F_{<}$

Keeps objects fresh for up to 5X longer!

EDWARD LEE, University of Waterloo, Canada

ONDŘEJ LHOTÁK, University of Waterloo, Canada

Reference immutability is a type based technique for taming mutation that has long been studied in the context of object-oriented languages, like Java. Recently, though, languages like Scala have blurred the lines between functional programming languages and object oriented programming languages. We explore how reference immutability interacts with features commonly found in these hybrid languages, in particular with higher-order functions – polymorphism – and subtyping. We construct a calculus System $F_{<,M}$ which encodes a reference immutability system as a simple extension of System $F_{<}$ and prove that it satisfies the standard soundness and immutability safety properties.

Additional Key Words and Phrases: System $F_{<}$, Reference Immutability, Type Systems

1 INTRODUCTION

Code written in a pure, functional language is *referentially transparent* – it has no side effects and hence can be run multiple times to produce the same result. Reasoning about referentially transparent code is easier for both humans and computers. However, purely functional code can be hard to write and inefficient, so many functional languages contain impure language features.

One important side effect that is difficult to reason about is *mutation* of state. Mutation arises naturally, but can cause bugs which can be hard to untangle; for example, two modules which at first glance are completely unrelated may interact through some shared mutable variable. Taming – or controlling – where and how mutation can occur can reduce these issues.

One method of taming mutation is *reference immutability* [Huang et al. 2012; Tschantz and Ernst 2005]. In this setting, the type of each reference to a value can be either mutable or immutable. An immutable reference cannot be used to mutate the value or any other values transitively reached from it.

Mutable and immutable references can coexist for the same value, so an immutable reference does not guarantee that the value will not change through some other, mutable reference. This is in contrast to the stronger guarantee of *object immutability*, which applies to values, and ensures that a particular value does not change through any of the references to it.

Reference immutability has long been studied in existing object-oriented programming languages such as Java [Huang et al. 2012; Tschantz and Ernst 2005; Zibin et al. 2007] and C# [Gordon et al. 2012]. However, reference immutability is largely unexplored in the context of functional languages with impure fragments – languages like Scala or OCaml, for example. Many programs in Scala are mostly immutable [Haller and Axelsson 2017]. A system that formally enforces specified patterns of immutability would help programmers and compilers better reason about immutability in such programs.

One feature that is important in all languages but especially essential in functional programs is polymorphism. The interaction of polymorphism and reference immutability raises interesting questions. Should type variables abstract over annotated types including their immutability annotations (such as `@readonly String`), or only over the base types without immutability annotations (such as `String`)? Should uses of type variables admit an immutability annotation like other types do? For

Authors' addresses: Edward Lee, Computer Science, University of Waterloo, 200 University Ave W., Waterloo, ON, N2L 3G1, Canada; Ondřej Lhoták, Computer Science, University of Waterloo, 200 University Ave W., Waterloo, ON, N2L 3G1, Canada.

2023. 2475-1421/2023/10-ART1 \$15.00

<https://doi.org/>

example, should `@readonly X` be allowed, where X is a type variable rather than a concrete type? If yes, then how should one interpret an annotated variable itself instantiated with an annotated type? For example, what should the type `@readonly X` mean if the variable X is instantiated with `@mutable String`?

Our contribution to this area is a *simple* and *sound* treatment of reference immutability in System $F_{<}$: [Cardelli et al. 1991]. Specifically, we formulate a simple extension System $F_{<:M}$ of System $F_{<}$: with the following properties:

- **Immutability safety:** When dealing with reference immutability, one important property to show is *immutability safety*: showing that when a reference is given a read-only type, then the underlying value is not modified through that reference. In System $F_{<:M}$ we introduce a dynamic form of immutability, a term-level `seal` construct, which makes precise the runtime guarantees that we expect from a reference that is statically designated as immutable by the type system. We do this by formalizing System λ_M , an untyped calculus with references and seals. Dynamic seals are transitive in that they seal any new references that are read from a field of an object through a sealed reference.
- **System $F_{<}$ -style polymorphism:** System $F_{<:M}$ preserves the same bounded-quantification structure of System $F_{<}$. At the same time, it allows type variables to be further modified by immutability modifiers.
- **Immutable types are types:** To allow for System $F_{<}$ -style polymorphism, we need to treat immutable types as types themselves. To do so, instead of type qualifiers, we introduce a type *operator* `readonly` that can be freely applied to existing types (including type variables). The `readonly` operator turns a type into an immutable version of the same type. While this complicates the definition of subtyping and proofs of canonical forms lemmas, we resolve these issues by reducing types to a normal form.

Our hope is to enable reference immutability systems in functional languages by giving simple, sound foundations in System $F_{<}$, a calculus that underpins many practical functional programming languages.

The rest of this paper is organized as follows. In Section 2 we give an overview of reference immutability. In Section 3 we introduce an un-typed core calculus, System λ_M , to describe sealing and how it relates to reference immutability safety at run time. In Section 4 we present System $F_{<:M}$, which enriches System λ_M with types, and show that it satisfies the standard soundness theorems. In Section 5 we use the soundness results from System $F_{<:M}$ and the dynamic safety results from System λ_M to show that our desired immutability safety properties hold in System $F_{<:M}$. We survey related and possible future work in Section 7 and we conclude in Section 8.

Our development is mechanized in the Coq artifact that we will submit to the OOPSLA artifact evaluation process.

2 REFERENCE IMMUTABILITY

Reference immutability at its core is concerned with two key ideas:

- **Immutable references:** References to values can be made immutable, so that the underlying value **cannot** be modified through that reference.
- **Transitive immutability:** An immutable reference to a *compound value* that contains other references cannot be used to obtain a mutable reference to another value. For example, if x is a read-only reference to a pair, the result of evaluating `x.first` should be *viewpoint adapted* [Dietl et al. 2007] to be a read-only reference, even if the pair contains references that are otherwise mutable.

For example, consider the following snippet of Scala-like code that deals with polymorphic mutable pairs.

```
case class Pair[X](var first: X, var second: X)

def good(x : Pair[Int]) = { x.first = 5 }
def bad1(y : @readonly Pair[Int]) = { y.first = 7 }
def bad2(y : @readonly Pair[Pair[Int]]) = { y.first.first = 5 }
def access(z: @readonly Pair[Pair[Int]]): @readonly Pair[Int] = { z.first }
```

A reference immutability system would deem the function `good` to be well-typed because it mutates the pair through a mutable reference `x`. However, it would disallow `bad1` because it mutates the pair through a read-only reference `y`. Moreover, it would also disallow `bad2` because it mutates the pair referenced indirectly through the read-only reference `y`. This can also be seen by looking at the `access` function, which returns a read-only reference of type `@readonly Pair[Int]` to the first component of the pair referenced by `z`.

2.1 Why though?

Immutable values are crucial even in impure functional programming languages because pure code is often easier to reason about. This benefits both the programmer writing the code, making debugging easier, and the compiler when applying optimizations.

Although most values, even in impure languages, are immutable by default [Haller and Axelsson 2017], mutable values are sometimes necessary for various reasons. For example, consider a compiler for a pure, functional, language. Such a compiler might be split into multiple passes, one which first builds and generates a *symbol table of procedures* during semantic analysis, and one which then uses that symbol table during *code generation*. For efficiency, we may wish to build both the table and the procedures in that table with an impure loop.

```
object analysis {
  class Procedure(name : String) {
    val locals : mutable.Map[String, Procedure] = mutable.Map.empty
    def addLocalProcedure(name: String, proc: Procedure) = {
      local += (name -> proc)
    }
  }

  val table : mutable.Map[String, Procedure] = mutable.Map.empty

  val analyze(ast: AST) = {
    ast.forEach(() => { table.add(new Procedure(...)) })
  }
}
```

The symbol table and the properties of the procedure should not be mutable everywhere, though; during code generation, our compiler should be able to use the information in the table to generate code but shouldn't be able to change the table nor the information in it! How do we enforce this though?

One solution is to create an immutable copy of the symbol table for the code generator, but this can be fragile. A naive solution which merely clones the table itself will not suffice, for example:

```
object analysis {
  private val table[analysis] = ...
```

```

def symbolTable : Map[String, Procedure] = table.toMap // create immutable
  copy of table.
}

object codegen {
  def go() = {
    analysis.symbolTable["main"].locals += ("bad" -> ...) // whoops...
  }
}

```

While this does create an immutable copy of the symbol table for the code generator, it does not create immutable copies of the procedures held in the table itself! We would need to recursively rebuild a new, immutable symbol table with new, immutable procedures to guarantee immutability, which can be an expensive proposition, both in terms of code and in terms of runtime costs.

Moreover, creating an immutable copy might not even work in all cases. Consider an interpreter for a pure, functional language with support for `letrec x := e in f`. The environment in which e is interpreted contains a cyclic reference to x , which necessitates mutation in the interpreter. Without special tricks like laziness this sort of structure cannot be constructed, let alone copied, without mutation.

```

abstract class Value { }
type Env = Map[String, Value]
case class Closure(var env: Env, params: List[String], body: Exp) extends Value

def interpret_letrec(env: Env, x: String, e: Exp, f: Exp) : Value = {
  val v = interpret(env + (x -> Nothing), e)
  case v of {
    Closure(env, params, body) => v.env = v.env + (x -> v) // Update binding
  }
  interpret (env + (x -> v), f)
}

```

Here, the closure that v refers to needs to be mutable while it is being constructed, but since the underlying language is pure, it should be immutable afterwards. In particular, we should not be able to mutate the closure through the self-referential reference $v.env = env + (x \rightarrow v)$, nor should we be able to mutate the closure while interpreting f .

We would like a system that prevents writes to v from the self-referential binding in its environment and from the reference we pass to `interpret (env + (x -> v), f)`. This is what *reference immutability* provides.

```

abstract class Value { }
type Env = Map[String, @readonly Value]
case class Closure(env: var Env, params: List[String], body: Exp)

def interpret_letrec(env: Env, x: String, e: Exp, f: Exp) : Value = {
  val v = interpret(env + (x -> Nothing), e)
  case v of {
    Closure(env, params, body) => v.env = env + (x -> @readonly v) // update
      binding
  }
  interpret (env + (x -> @readonly v), f)
}

```

3 DYNAMIC IMMUTABILITY SAFETY

Now, to formalize reference immutability, we need to formalize exactly when references are used to update the values they refer to. For example, from above, how do we check that access does what it *claims* to do?

```
def access(z: @readonly Pair[Pair[Int]]): @readonly Pair[Int] = { z.first }
```

How do we check that access returns a reference to `z.first` that, at runtime, is never used to write to `z.first` or any other values transitively reachable from it through other references? How do we even express this guarantee precisely?

If we consider a reference as a collection of getter and setter methods for the fields of the object it refers to, we could ensure that a reference is immutable by dropping all the setter methods. To ensure that immutability is transitive, we would also need to ensure that the result of applying a getter method is also immutable, i.e. by also dropping its setter methods and recursively applying the same modification to *its* getter methods. We will make this precise by introducing the System λ_M calculus with a notion of *sealed* references.

3.1 System λ_M

To answer this question we introduce System λ_M , the untyped lambda calculus with collections of mutable references – namely, records – extended with a mechanism for *sealing* references. System λ_M is adapted from the CS-machine of Felleisen and Friedman [1987] and extended with rules for dealing with sealed references.

Sealed references: To address the question about dynamic, runtime safety – can we ensure that read-only references are never used to mutate values – references can be explicitly *sealed* so that any operation that will mutate the cell referenced will fail to evaluate; see Figure 1.

The `seal` form protects its result from writes. A term under a `seal` form reduces until it becomes a value. At that point, values that are not records, like functions and type abstractions, are just transparently passed through the `seal` construct. However, values that *are* – records – remain protected by the `seal` form, and do not reduce further. For example:

$$\text{seal } (\{y : 0x0001\})$$

is an irreducible value – a sealed record where the first field is stored at location 1 in the store. Intuitively, this can be viewed as removing the setter methods from an object reference. A sealed reference `seal v` behaves *exactly* like its unsealed variant `v` except that writes to `seal v` are forbidden and reads from `seal v` return sealed results.

Rules that mutate the cells corresponding to a record explicitly require an unsealed open record; see (**WRITE-FIELD**). This ensures that any ill-behaved program that mutates a store cell through a sealed record will get stuck, while an unsealed record can have its fields updated:

$$\begin{aligned} \langle \{x : 10\}.x = 5, [] \rangle &\longrightarrow \langle \{x : 0x0001\}.x = 5, [0x0001 : 10] \rangle \\ &\longrightarrow \langle 10, [0x0001 : 5] \rangle \end{aligned}$$

A sealed record cannot have its fields written to. Unlike record field reads, for which there is a sealed (**SEALED-FIELD**) counterpart to the standard record read rule (**FIELD**), there is no corresponding rule for writing to a sealed record for (**WRITE-FIELD**). Recall that (**WRITE-FIELD**) requires an *open*, *unsealed* record as input:

$$\frac{l : v \in \sigma}{\langle \{ \dots x : l \dots \}.x = v', \sigma \rangle \longrightarrow \langle v, \sigma[l \mapsto v'] \rangle}$$

$s, t ::=$	Terms	l	Location
$\lambda x.t$	term abstraction	$s, t ::=$	Runtime Terms
x	term variable	$\{x_1 : l_1, x_2 : l_2, \dots\}$	runtime record
$s(t)$	application	$v ::=$	Runtime Values
$\{f_1 : s_1, f_2 : s_2, \dots\}$	records	$\lambda x.t$	
$s.f$	field read	$\{f_1 : l_1, f_2 : l_2, \dots\}$	
$s.f = t$	field write	$\text{seal } \{f_1, l_1, f_2 : l_2, \dots\}$	
$\text{seal } s$	sealing		

$\langle (\lambda x.t)(v), \sigma \rangle \longrightarrow \langle t[x \mapsto v], \sigma \rangle$ (BETA-V)	$\frac{l : v \in \sigma}{\langle (\text{seal } \{\dots x : l \dots\}) .x, \sigma \rangle \longrightarrow \langle \text{seal } v, \sigma \rangle}$ (SEALED-FIELD)
$\frac{l_i \notin \sigma}{\langle \{x_i : v_i\}, \sigma \rangle \longrightarrow \langle \{x_i : l_i\}, (\sigma, l_1 : v_1, l_2 : v_2, \dots) \rangle}$ (RECORD-STORE)	$\langle \text{seal } (\lambda x.t), \sigma \rangle \longrightarrow \langle \lambda x.t, \sigma \rangle$ (SEAL-ELIM-ABS)
$\frac{l : v \in \sigma}{\langle \{\dots x : l \dots\} .x, \sigma \rangle \longrightarrow \langle v, \sigma \rangle}$ (FIELD)	$\langle \text{seal seal } v, \sigma \rangle \longrightarrow \langle \text{seal } v, \sigma \rangle$ (SEAL-ELIM-MULTIPLE)
$\frac{l : v \in \sigma}{\langle \{\dots x : l \dots\} .x = v', \sigma \rangle \longrightarrow \langle v, \sigma[l \mapsto v'] \rangle}$ (WRITE-FIELD)	$\frac{\langle s, \sigma \rangle \longrightarrow \langle t, \sigma' \rangle}{\langle E[s], \sigma \rangle \longrightarrow \langle E[t], \sigma' \rangle}$ (CONTEXT)

$E ::=$	Evaluation Context
$[]$	
$E(t) \mid v(E)$	
$\{x_0 : v_0, \dots, x_i : E, x_{i+1} : t_{i+1}, \dots\}$	
$E.x$	
$E.x = t \mid v.x = E$	
$\text{seal } E$	

Fig. 1. The syntax and semantics of λ_m .

The calculus does not contain any rule like the following, which would reduce writes on a sealed record:

$$\frac{l : v \in \sigma}{\langle (\text{seal } \{\dots x : l \dots\}) .x = v', \sigma \rangle \longrightarrow \langle v, \sigma[l \mapsto v'] \rangle}$$

So a term like:

$$\begin{aligned} \langle (\text{seal } \{x : 10\}) .x = 5, [] \rangle &\longrightarrow \langle \text{seal } (\{x : 0x0001\}) .x = 5, [0x0001 : 10] \rangle \\ &\longrightarrow \text{gets stuck.} \end{aligned}$$

Dynamic viewpoint adaptation: After reading a field from a sealed record, the semantics *seals* that value, ensuring transitive safety – see (SEALED-FIELD).

$$\frac{l : v \in \sigma}{\langle (\text{seal } \{\dots x : l \dots\}) .x, \sigma \rangle \longrightarrow \langle \text{seal } v, \sigma \rangle}$$

For example:

$$\begin{aligned} \langle (\text{seal } \{y : \{x : 10\}\}).y, [] \rangle &\longrightarrow \langle \text{seal } (\{y : \{x : 0x001\}\}).y, [0x001 : 10] \rangle \\ &\longrightarrow \langle \text{seal } (\{y : 0x002\}).y, [0x001 : 10, 0x002 : \{x : 0x001\}] \rangle \\ &\longrightarrow \langle \text{seal } (\{x : 0x001\}), [0x001 : 10, 0x002 : \{x : 0x001\}] \rangle \end{aligned}$$

Sealed references and *dynamic viewpoint adaptation* allow for a succinct guarantee of *dynamic transitive immutability safety* – that no value is ever mutated through a read-only reference or any other references transitively derived from it.

Aside from preventing writes through sealed references, we should show that sealing does not otherwise affect reduction. For this we need a definition that relates pairs of terms that are essentially equivalent except that one has more seals than the other.

Definition 3.1. Let s and t be two terms. We say $s \leq t$ if t can be obtained from s by repeatedly replacing sub-terms s' of s with sealed subterms $\text{seal } s'$.

This implies a similar definition for stores:

Definition 3.2. Let σ and σ' be two stores. We say $\sigma \leq \sigma'$ if and only if they have the same locations and for every location $l \in \sigma$, we have $\sigma(l) \leq \sigma'(l)$.

The following three lemmas formalize how reduction behaves for terms that are equivalent modulo seals. The first one is for a term t that is equivalent to a value – it states that if t reduces, the resulting term is still equivalent to the same value. It also shows that the resulting term has fewer seals than t , which we'll need later for an inductive argument.

Definition 3.3. Let s be a term. Then $|s|$ is the number of seals in s .

LEMMA 3.4. Let v be a value, σ_v be a store, t be a term such that $v \leq t$, and σ_t be a store such that $\sigma_v \leq \sigma_t$.

If $\langle t, \sigma_t \rangle \longrightarrow \langle t', \sigma'_t \rangle$ then $v \leq t'$, $\sigma_v \leq \sigma'_t$, and $|t'| < |t|$.

The next lemma is an analogue of Lemma 3.4 for terms. Given two equivalent terms s and t , if s steps to s' and t steps to t' , then either s and t' are equivalent or s' and t' are equivalent. Moreover, again, to show that reduction in t is equivalent to reduction in s , we have that $|t'| < |t|$ if $s \leq t'$.

LEMMA 3.5. Let s, t be terms such that $s \leq t$ and let σ_s, σ_t be stores such that $\sigma_s \leq \sigma_t$. If $\langle s, \sigma_s \rangle \longrightarrow \langle s', \sigma'_s \rangle$ and $\langle t, \sigma_t \rangle \longrightarrow \langle t', \sigma'_t \rangle$ then:

- (1) Either $s \leq t'$, $\sigma_s \leq \sigma'_t$, and $|t'| < |t|$, or
- (2) $s' \leq t'$ and $\sigma'_s \leq \sigma'_t$.

Together, Lemmas 3.4 and 3.5 relate how terms s and t reduce when they are equivalent modulo seals. Assuming that both s and t reduce, every step of s corresponds to finitely many steps of t , and they reduce to equivalent results as well. This shows that sealing is transparent when added onto references that are never written to, allowing for a succinct guarantee of immutability safety.

Finally, the last lemma states that erasing seals will never cause a term to get stuck. Seals can be safely erased without affecting reduction.

LEMMA 3.6. Let s, t be terms such that $s \leq t$ and let σ_s, σ_t be stores such that $\sigma_s \leq \sigma_t$. If $\langle t, \sigma_t \rangle \longrightarrow \langle t', \sigma'_t \rangle$ then:

- (1) Either $s \leq t'$, $\sigma_s \leq \sigma'_t$, and $|t'| < |t|$, or
- (2) There exists s' and σ'_s such that $\langle s, \sigma_s \rangle \longrightarrow \langle s', \sigma'_s \rangle$, $s' \leq t'$ and $\sigma'_s \leq \sigma'_t$.

From this we can derive the following multi-step analogue, after observing the following lemma:

LEMMA 3.7. *If s is a term and v is a value such that $s \leq v$, then s is also a value.*

Hence:

LEMMA 3.8. *Suppose s and t are terms such that $s \leq t$. If $\langle t, \sigma_t \rangle \longrightarrow^* \langle v_t, \sigma'_t \rangle$ for some value v_t , then for any $\sigma_s \leq \sigma_t$ we have $\langle s, \sigma_s \rangle \longrightarrow^* \langle v_s, \sigma'_s \rangle$ such that $v'_s \leq v'_t$ and $\sigma'_s \leq \sigma'_t$.*

Finally, it can be shown that the seals are to blame when two equivalent terms s and t reduce differently – in particular, when one reduces but the other gets stuck.

LEMMA 3.9. *Let s, t be terms such that $s \leq t$, and let σ_s, σ_t be stores such that $\sigma_s \leq \sigma_t$. If $\langle s, \sigma_s \rangle \longrightarrow \langle s', \sigma'_s \rangle$ and t gets stuck, then the reduction performed on s was a write to a record using rule (**WRITE-FIELD**).*

PROOF. (Sketch) As s cannot further reduce, the evaluation context of s and t must match; there are no extraneous seals that need to be discharged. As such, from inspection of the reduction rules, we see that in all cases except for (**WRITE-FIELD**), for every possible reduction that s could have taken, there is a possible reduction that t could have taken as well, as desired. \square

4 TYPING AND STATIC SAFETY

System λ_M provides a *dynamic guarantee* that a given program will never modify its sealed references, but it does not provide any static guarantees about the dynamic behavior of a given program. To do that, we need a type system for System λ_M that will reject programs like `access(seal Pair(3,5)).first = 10`, which we know will crash.

To ensure that well-typed programs do not get stuck, a type system for System λ_M needs a static analogue of sealing – a way to turn an existing type into a *read-only type*. Read-only types denote references that are *immutable* and that (transitively) *adapt* any other references read through them to be *immutable* as well.

Issues arise, however, when we introduce polymorphism.

4.1 Polymorphism

Recall our earlier example – a polymorphic `Pair` object.

```
case class Pair[X](var first: X, var second: X)
```

In a functional language, it is only natural to write higher-order functions that are polymorphic over the elements stored in the pair. Consider an in-place map function over pairs, which applies a function to each element in the pair, storing the result in the original pair. This naturally requires mutable access to a pair.

```
def inplace_map[X](pair: Pair[X], f: X => X): Unit = {
  pair.first = f(pair.first);
  pair.second = f(pair.second);
}
```

This is all well and good, but we may wish to restrict the behaviour of `f` over the elements of the pair. It may be safer to restrict the behaviour of `f` so that it could not mutate the elements passed to it. Note that we cannot restrict access to the pair, however, as we still need to mutate it.

```
// Is this well founded?
def inplace_map[X](pair: Pair[X], f: @readonly X => X): Unit = {
  pair.first = f(pair.first);
  pair.second = f(pair.second);
}
```


Now, such a definition requires the ability to further modify type variables with immutability qualifiers. This raises important questions – for example, is this operation even well founded? This depends on what X ranges over.

X ranges over an unqualified type: If type variables range over types which have not been qualified by `@readonly`, then this operation is clearly well founded – it is simply qualifying the unqualified type that X will eventually be substituted by with the `@readonly` qualifier. This approach has been used by ReIm for Java and for an immutability system for C# – [Gordon et al. 2012; Huang et al. 2012].

However, this raises the problem of polymorphism over immutability qualifiers as well – for example, a `Pair` should be able to store both immutable and mutable object references. The only natural solution is to then introduce a *mutability* qualifier binder to allow for polymorphism over immutability qualifiers, as thus:

```
case class Pair[M, X](var first: M X, var second: M X)
def inplace_map[M, X](pair: Pair[M, X], f: @readonly X => M X): Unit = {
  pair.first = f(pair.first);
  pair.second = f(pair.second);
}
```

Mutability qualifier binders have been used previously, most notably by [Gordon et al. 2012]. For one, updating the binding structure of a language is not an easy task – ReIm notably *omits* this sort of parametric mutability polymorphism [Huang et al. 2012]. However, this sort of solution has its downsides; in particular, existing higher-order functions need to be updated with immutability annotations or variables, as type variables no longer stand for a full type. For example, an existing definition of `List map` which appears as thus originally:

```
def map[X](l: List[X], f: X => X): List[X]
```

needs to be updated to read as the following instead:

```
def map[M, X](l: List[M X], f: M X => M X): List[M X]
```

Instead, we would like to have X range over fully qualified types as well, but as we will see that poses some issues as well.

X ranges over fully-qualified types: If type variables can range over types which have been already qualified by `@readonly`, then we can avoid introducing mutability binders in the definitions for `Pair`, `inplace_map`, and `map` above. A `Pair` can be polymorphic over its contents X without caring about the underlying mutability of X . However, this raises the question – how do we interpret repeated applications of the `@readonly` qualifier? For example, what if we applied `inplace_map` on a `Pair[@readonly Pair[Int]]`? Then `inplace_map` would expect a function `f` with type `@readonly (@readonly Pair[Int]) => @readonly Pair[Int]`. While our intuition would tell us that `@readonly (@readonly Pair[Int])` is really just a `@readonly Pair[Int]`, discharging this equivalence in a proof is not so easy.

One response is to explicitly prevent type variables from being further qualified. Calculi which take this approach include [Tschantz and Ernst 2005; Zibin et al. 2007]. However, this restriction prevents this version of `inplace_map` from being expressed. How can we address this?

Our approach, which we explain below, is to treat `@readonly` as a type operator that works over all types. Following the intuition that sealing removes setters from references, `@readonly` should be a type operator which removes setters from types. While this does cause complications, we show below how types like `@readonly @readonly Pair[Int]` can be dealt with, using *subtyping* and *type normalization*.

$s, t ::=$ $\lambda x. t$ $\Lambda(X <: S). t$ x $s(t)$ $s[T]$ $\{f_1 : s_1, f_2 : s_2, \dots\}$ $s.x$ $s.x = t$ seal s	Terms term abstraction type abstraction term variable application type application records field read field write sealing	$S, T ::=$ X $S \rightarrow T$ $\forall(X <: S). T$ $S \wedge T$ $\{f : T\}$ readonly T $\Gamma ::=$ \cdot $\Gamma, x : T$ $\Gamma, X <: T$	Types type variable function type for-all type intersection type record type readonly type Environment empty term binding type binding
l $s, t ::=$ $\{x_1 : l_1, x_2 : l_2, \dots\}$ $v ::=$ $\lambda x. t$ $\Lambda(X <: S). t$ $\{f_1 : l_1, f_2 : l_2, \dots\}$ seal $\{f_1, l_1, f_2 : l_2, \dots\}$	Location Runtime Terms runtime record Runtime Values	$\sigma ::=$ \cdot $\sigma, l : v$ $\Sigma ::=$ \cdot $\sigma, l : T$	Store empty cell l with value v Store Environment empty cell binding

Fig. 2. The syntax of System $F_{<,M}$.

4.2 System $F_{<,M}$

To address these issues, we introduce System $F_{<,M}$, which adds a type system in the style of System $F_{<}$ to System λ_M . The syntax of System $F_{<,M}$ is given in Figure 2; changes from System $F_{<}$ are noted in grey.

System $F_{<,M}$ is a straightforward extension of System $F_{<}$ with collections of mutable references – namely, records – and with two new extensions: *read-only* types and *sealed* references. To be close to existing functional languages with subtyping and records, records in System $F_{<,M}$ are modelled as intersections of single-element record types, to support record subsumption, as in [Amin et al. 2016] and [Reynolds 1997]. See Figures 4 and 5 for full subtyping and typing rules respectively.

Read-only types: The readonly type operator transforms an existing type to a read-only version of itself. Unlike the read-only mutability qualifier in Javari and ReIm, which is paired with a type to form a pair of a qualifier and a type, a read-only type in System $F_{<,M}$ is itself a type. The readonly operator can be seen as the static counterpart of sealing or of deleting setter methods from an object-oriented class type.

Any type T is naturally a subtype of its readonly counterpart **readonly** T , which motivates the choice of System $F_{<}$ as a base calculus. This subtyping relationship is reflected in the subtyping rule (**MUTABLE**). The (**SEAL**) typing rule gives a read-only type to sealed references.

Normal Forms

S, T	::=	Types in normal form	
		$\bigwedge_i (R_i)$	Intersection of components
R	::=	Normal form type components	
		\top	Top type
		$S \rightarrow T$	Normal function type
		$\forall (X <: S). T$	Normal for-all type
		$\{f : S\}$	Normal record type
		X	Type variable
		$\text{readonly } \{f : S\}$	Read-only normal record type
		$\text{readonly } X$	Read-only type variable

Fig. 3. Normal forms for System $F_{<}$.

Subtyping

 $\Gamma \vdash S <: T$

$\Gamma \vdash T <: T$	(REFL)	
$\frac{\Gamma \vdash R <: S \quad \Gamma \vdash S <: T}{\Gamma \vdash R <: T}$	(TRANS)	$\frac{\Gamma \vdash S <: T \quad \Gamma \vdash T <: S}{\Gamma \vdash \{x : S\} <: \{x : T\}}$ (RECORD)
$\frac{X <: T \in \Gamma}{\Gamma \vdash X <: T}$	(TVAR)	$\frac{\Gamma \vdash S <: T}{\Gamma \vdash \text{readonly } \{x : S\} <: \text{readonly } \{x : T\}}$ (READONLY-RECORD)
$\Gamma \vdash U <: \top$	(TOP)	$\Gamma \vdash S \wedge T <: S$ (INTER-LEFT)
$\frac{\Gamma \vdash T_1 <: S_1 \quad \Gamma \vdash S_2 <: T_2}{\Gamma \vdash S_1 \rightarrow S_2 <: T_1 \rightarrow T_2}$	(ARROW)	$\Gamma \vdash S \wedge T <: T$ (INTER-RIGHT)
$\frac{\Gamma \vdash T_1 <: S_1 \quad \Gamma, X <: T_1 \vdash S_2 <: T_2}{\Gamma \vdash \forall (X <: S_1). S_2 <: \forall (X <: T_1). T_2}$	(ALL)	$\frac{\Gamma \vdash S <: T_1 \quad \Gamma \vdash S <: T_2}{\Gamma \vdash S <: T_1 \wedge T_2}$ (INTER)
$\frac{\Gamma \vdash S <: T}{\Gamma \vdash \text{readonly } S <: \text{readonly } T}$	(READONLY)	
$\frac{\Gamma \vdash S <: T}{\Gamma \vdash S <: \text{readonly } T}$	(MUTABLE)	
$\frac{\Gamma \vdash nf(S) <: nf(T)}{\Gamma \vdash S <: T}$	(DENORMALIZE)	

Fig. 4. Subtyping rules of System $F_{<}$.

Typing and Runtime Typing

$$\Gamma \mid \Sigma \vdash t : T \text{ and } \Gamma \mid \Sigma \vdash \sigma$$

$$\begin{array}{c}
\frac{x : T \in \Gamma}{\Gamma \mid \Sigma \vdash x : T} \text{ (VAR)} \\
\frac{\Gamma, x : S \vdash t : T}{\Gamma \mid \Sigma \vdash \lambda x. t : S \rightarrow T} \text{ (ABS)} \\
\frac{\Gamma, X <: S \vdash t : T}{\Gamma \mid \Sigma \vdash \Lambda(X <: S). t : \forall(X <: S). T} \text{ (T-ABS)} \\
\frac{\Gamma \mid \Sigma \vdash t : S \rightarrow T \quad \Gamma \mid \Sigma \vdash s : S}{\Gamma \mid \Sigma \vdash t(s) : T} \text{ (APP)} \\
\frac{\Gamma \mid \Sigma \vdash t : \forall(X <: S). T \quad \Gamma \mid \Sigma \vdash S' <: S}{\Gamma \mid \Sigma \vdash t[S'] : T[X \mapsto S']} \text{ (T-APP)} \\
\frac{\Gamma \mid \Sigma \vdash s : S}{\Gamma \mid \Sigma \vdash \text{seal } s : \text{readonly } S} \text{ (SEAL)} \\
\frac{\Gamma \mid \Sigma \vdash s : \text{readonly } \{x : S\}}{\Gamma \mid \Sigma \vdash s.x : \text{readonly } S} \text{ (READONLY-RECORD-ELIM)} \\
\frac{\text{dom}(\sigma) = \text{dom}(\Sigma) \quad \forall l \in \text{dom}(\Sigma), \Gamma \mid \Sigma \vdash \sigma(l) : \Sigma(l)}{\Gamma \mid \Sigma \vdash \sigma} \text{ (STORE)} \\
\frac{\Gamma \mid \Sigma \vdash t_i : T_i}{\Gamma \mid \Sigma \vdash \{x_i : t_i \dots\} : \bigwedge_i \{x_i : T_i\}} \text{ (RECORD-INTRO)} \\
\frac{\Gamma \mid \Sigma \vdash t : \{x : T\}}{\Gamma \mid \Sigma \vdash t.x : T} \text{ (RECORD-ELIM)} \\
\frac{\Gamma \mid \Sigma \vdash s : \{x : T\} \quad \Gamma \vdash t : T}{\Gamma \mid \Sigma \vdash s.x = t : T} \text{ (RECORD-UPDATE)} \\
\frac{\Gamma \mid \Sigma \vdash s : S \quad \Gamma \vdash S <: T}{\Gamma \mid \Sigma \vdash s : T} \text{ (SUB)} \\
\frac{l_i : T_i \in \Sigma}{\Gamma \mid \Sigma \vdash \{x_i : l_i\} : \bigwedge_i \{x_i : T_i\}} \text{ (RUNTIME-RECORD)}
\end{array}$$

Fig. 5. Typing rules for System $F_{<:,M}$

Static viewpoint adaptation: The (READONLY-RECORD-ELIM) rule is a static counterpart of the (SEALED-FIELD) reduction rule. Given a reference s to a record with read-only type, it gives a read-only type to the result of a read $s.x$ of a field x from that reference. If S is the type of field x in the record type given to s , the rule viewpoint-adapts the type, giving $s.x$ the type $\text{readonly } S$.

4.2.1 Normal Forms for Types. In System $F_{<:,M}$, readonly is a type operator that can be applied to any type, which enables us to express types such as $\text{readonly } X$, where X is some type variable of unknown mutability. However, if X is itself instantiated with some readonly type $\text{readonly } T$, the type $\text{readonly } X$ becomes $\text{readonly } \text{readonly } T$, with two occurrences of the type operator. Intuitively, such a type should have the same meaning as $\text{readonly } T$.

Additionally, certain types should be equivalent under subtyping. For example, for both backwards compatibility and simplicity, arrow $S \rightarrow T$ and for-all types $\forall(X <: S). T$ should be equivalent under subtyping to their read-only forms $\text{readonly } (S \rightarrow T)$ and $\text{readonly } (\forall(X <: S). T)$, respectively, as well.

Normalization $nf(T)$ and $merge(T)$

$nf(T)$::=	Normalization
\top	=>	\top
X	=>	X
$S \rightarrow T$	=>	$nf(S) \rightarrow nf(T)$
$\forall(X <: S).T$	=>	$\forall(X <: nf(S)).nf(T)$
$S \wedge T$	=>	$nf(S) \wedge nf(T)$
$\{f : T\}$	=>	$\{f : nf(T)\}$
$readonly\ T$	=>	$merge(T)$

$merge(T)$::=	Merging
X	=>	$readonly\ X$
$\{f : T\}$	=>	$readonly\ \{f : T\}$
$S \wedge T$	=>	$merge(S) \wedge merge(T)$
$_$	=>	T

Fig. 6. Normalizing Types for System $F_{<M}$.

Having multiple representations for the same type, even infinitely many, complicates reasoning about the meanings of types and proofs of soundness. Therefore, we define a canonical representation for types as follows:

Definition 4.1. A type T is in normal form if:

- (1) T is the top type \top .
- (2) T is a function type $S_1 \rightarrow S_2$, where S_1 and S_2 are in normal form.
- (3) T is an abstraction type $\forall(X <: S_1).S_2$, where S_1 and S_2 are in normal form.
- (4) T is an intersection type $S_1 \wedge S_2$, where S_1 and S_2 are in normal form.
- (5) T is a record type $\{x : S\}$, where S is in normal form.
- (6) T is a read-only record type $readonly\ \{x : S\}$, where S is in normal form.
- (7) Type variables X and read-only type variables $readonly\ X$ are in normal form.

A type in normal form is simple – it is an intersection of function, abstraction, and record types, each possibly modified by a single `readonly` operator. For example, $\{x : X\} \wedge readonly\ \{y : Y\}$ is in normal form. The type $readonly\ (\{x : X\} \wedge \{y : Y\})$ is not. A grammar for types in normal form can be found in Figure 3.

This allows us to reason about both the shape of the underlying value being typed, and whether or not it has been modified by a `readonly` operator. Naturally we need a theorem which states that every type has a normal form and a function nf to compute that normal form. Such a function nf is shown in Figure 6. Normalization both computes a normal form and is idempotent – a type in normal form normalizes to itself.

LEMMA 4.2. *For any type T , $nf(T)$ is in normal form. Moreover, if T is in normal form, $nf(T) = T$.*

Moreover, types are equivalent to their normalized forms under the subtyping relationship.

LEMMA 4.3. $\Gamma \mid \Sigma \vdash nf(T) <: T$ and $\Gamma \mid \Sigma \vdash T <: nf(T)$.

PROOF. For one direction, note that $nf(nf(T)) = nf(T)$, and hence $nf(nf(T)) <: nf(T)$. Applying (DENORMALIZE) allows us to show that $nf(T) <: T$, as desired. The other case follows by a symmetric argument. \square

Not only does this allow us to simplify types to a normal form, this also allows us to state and prove canonical form lemmas and inversion lemmas, necessary for preservation and progress: Theorems 4.9 and 4.11. Below we give examples for record types. Similar lemmas exist and are mechanized for function types and type-abstraction types as well.

LEMMA 4.4 (INVERSION OF RECORD SUBTYPING). *If S is a subtype of $\{f : T'\}$, and S is in normal form, then at least one of its components is a type variable X or a record type $\{f : S'\}$, where $\Gamma \vdash T' <: S' <: T'$.*

LEMMA 4.5 (CANONICAL FORMS FOR RECORDS). *If v is a value and $\emptyset \mid \Sigma \vdash v : \{f : T\}$, then v is a record and f is a field of v that maps to some location l .*

LEMMA 4.6 (INVERSION OF READ-ONLY RECORD SUBTYPING). *If S is a subtype of $\text{readonly } \{f : T'\}$, and S is in normal form, then at least one of its components is a type variable X , read-only type variable $\text{readonly } X$, a record type $\{f : S'\}$ where $\Gamma \vdash T' <: S' <: T'$, or a read-only record type $\text{readonly } \{f : S'\}$ where $\Gamma \vdash T' <: S' <: T'$.*

LEMMA 4.7 (CANONICAL FORMS FOR READ-ONLY RECORDS). *If v is a value and $\emptyset \mid \Sigma \vdash v : \text{readonly } \{f : T\}$, then v is a record or a sealed record and f is a field of v that maps to some location l .*

Note that normalization is necessary to state the inversion lemmas for read-only records, as $\text{readonly } \{f : T'\}$, $\text{readonly } \text{readonly } \{f : T'\}$, etc, give an infinite series of syntactically inequivalent but semantically equivalent types describing the same object – a read-only record where field f has type T' .

4.2.2 *Operational Safety.* Operationally, we give small-step reduction semantics coupled with a store to System $F_{<:\mathbb{M}}$ in Figure 7.

Again, these rules are a straightforward extension of System $F_{<:}$ with mutable boxes and records, with additional rules for reducing sealed records. To prove progress and preservation theorems, we additionally need to ensure that the store σ itself is well typed in the context of some store typing environment Σ – see rule (STORE).

The crux of preservation for System $F_{<:\mathbb{M}}$ is to show that sealed records are never given a non-read-only type, so that the typing rule for reading from a mutable record – (RECORD-ELIM) – cannot be applied to sealed record values.

LEMMA 4.8. *Suppose $\Gamma \mid \Sigma \vdash \text{seal } r : T$ for some record r . If T is in normal form, then the components of T are:*

- The top type \top , or
- a read-only record type $\text{readonly } \{f : T'\}$.

From this key result we can show that preservation holds for System $F_{<:\mathbb{M}}$.

THEOREM 4.9 (PRESERVATION OF SYSTEM $F_{<:\mathbb{M}}$). *Suppose $\langle s, \sigma \rangle \longrightarrow \langle t, \sigma' \rangle$. If $\Gamma \mid \Sigma \vdash \sigma$ and $\Gamma \mid \Sigma \vdash s : T$ for some type T , then there is some environment extension Σ' of Σ such that $\Gamma \mid \Sigma' \vdash \sigma'$ and $\Gamma \mid \Sigma' \vdash t : T$.*

Conversely, values given a non-read-only record type must be an unsealed collection of references.

Evaluation

$$\boxed{\langle s, \sigma \rangle \longrightarrow \langle t, \sigma' \rangle}$$

$$\begin{array}{c}
\langle (\lambda x.t)(v), \sigma \rangle \longrightarrow \langle t[x \mapsto v], \sigma \rangle \text{ (BETA-V)} \\
\hline
\langle \{x_i : v_i\}, \sigma \rangle \longrightarrow \langle \{x_i : l_i\}, (\sigma, l_1 : v_1, l_2 : v_2, \dots) \rangle \\
\text{(RECORD-STORE)} \\
\hline
\langle \{\dots x : l \dots\}.x, \sigma \rangle \longrightarrow \langle v, \sigma \rangle \text{ (FIELD)} \\
\hline
\langle \{\dots x : l \dots\}.x = v', \sigma \rangle \longrightarrow \langle v, \sigma[l \mapsto v'] \rangle \\
\text{(WRITE-FIELD)} \\
\langle (\Lambda(X <: S).t)[T], \sigma \rangle \longrightarrow \langle t[X \mapsto T], \sigma \rangle \\
\text{(BETA-T)} \\
\hline
E ::= [] \mid E(t) \mid v(E) \mid E[T] \\
\mid \{x_0 : v_0, \dots, x_i : E, x_{i+1} : t_{i+1}, \dots\} \\
\mid E.x \\
\mid E.x = t \mid v.x = E \\
\mid \text{seal } E
\end{array}$$

$$\begin{array}{c}
\langle s, \sigma \rangle \longrightarrow \langle t, \sigma' \rangle \\
\hline
\langle (\text{seal } \{\dots x : l \dots\}).x, \sigma \rangle \longrightarrow \langle \text{seal } v, \sigma \rangle \\
\text{(SEALED-FIELD)} \\
\langle \text{seal } (\lambda x.t), \sigma \rangle \longrightarrow \langle \lambda x.t, \sigma \rangle \\
\text{(SEAL-ELIM-ABS)} \\
\langle \text{seal } (\Lambda(X <: S).t), \sigma \rangle \longrightarrow \langle \Lambda(X <: S).t, \sigma \rangle \\
\text{(SEAL-ELIM-TABS)} \\
\langle \text{seal seal } v, \sigma \rangle \longrightarrow \langle \text{seal } v, \sigma \rangle \\
\text{(SEAL-ELIM-MULTIPLE)} \\
\langle s, \sigma \rangle \longrightarrow \langle t, \sigma' \rangle \\
\hline
\langle E[s], \sigma \rangle \longrightarrow \langle E[t], \sigma' \rangle \text{ (CONTEXT)}
\end{array}$$

Evaluation ContextFig. 7. Reduction rules for System $F_{<:M}$

LEMMA 4.10. *Suppose $\emptyset \mid \Sigma \vdash v : \{f : T\}$ for runtime value v . Then v is an unsealed runtime record where field f maps to some location l .*

This lemma is needed to prove progress.

THEOREM 4.11 (PROGRESS FOR SYSTEM $F_{<:M}$). *Suppose $\emptyset \mid \Sigma \vdash \sigma$ and $\emptyset, \Sigma \vdash s : T$. Then either s is a value or there is some t and σ' such that $\langle s, \sigma \rangle \longrightarrow \langle t, \sigma' \rangle$.*

5 STATIC IMMUTABILITY SAFETY

Armed with Progress and Preservation, we can state immutability safety for full System $F_{<:M}$. System λ_M allows us to show that sealed records are never used to mutate their underlying referenced values. System $F_{<:M}$ shows that well-typed programs using seals do not get stuck. To prove immutability safety for System $F_{<:M}$, one problem still remains – System $F_{<:M}$ allows records that are not sealed to be given a read-only type. We still need to show that records with such a type are not used to mutate their values. In other words, we need to show that records with a read-only type *could be* sealed, and that the resulting program would execute in the same way.

We will do this by showing that, given an original, well-typed System $F_{<:M}$ program s , we can add seals to its read-only subterms to obtain a new, well-typed System $F_{<:M}$ program t , and furthermore that t behaves the same way as s , up to having additional seals in the resulting state.

The first step is to show that sealing does not disturb the typing judgment for terms.

LEMMA 5.1. *Suppose $\Gamma \mid \Sigma \vdash t : \text{readonly } T$. Then $\Gamma \mid \Sigma \vdash \text{seal } t : \text{readonly } T$.*

PROOF. By (SEAL), $\Gamma \mid \Sigma \vdash \text{seal } t : \text{readonly readonly } T$. Then since $\text{readonly readonly } T < \text{readonly } T$, by (SUB), $\Gamma \mid \Sigma \vdash \text{seal } t : \text{readonly } T$, as desired. \square

From this, given a term s and a typing derivation for s , $D = \Gamma \mid \Sigma \vdash s : T$, we can seal those subterms of s that are given a read-only type in D .

LEMMA 5.2. *Let C be a term context with n holes, and let $s = C[s_1, s_2, s_3, \dots, s_n]$ be a term. Suppose D is a typing derivation showing that $\Gamma \mid \Sigma \vdash s : T$. Suppose also that D gives each subterm s_i of s a type $\text{readonly } T_i$. Then $s' = s[\text{seal } s_1, \text{seal } s_2, \dots, \text{seal } s_n]$ has the following properties:*

- (1) $s \leq s'$, and
- (2) *There exists a typing derivation D' showing that $\Gamma \mid \Sigma \vdash s' : T$ as well.*

PROOF. (1) is by definition. As for (2), to construct D' , walk through the typing derivation D showing that $\Gamma \mid \Sigma \vdash s : T$. When we reach the point in the typing derivation that shows that s_i is given the type $\text{readonly } T_i$, note that $\text{seal } s_i$ can also be given the type $\text{readonly } T_i$ by the derivation given by Lemma 5.1. Replace the sub-derivation in D with the derivation given by Lemma 5.1 to give a derivation in D' for $\text{seal } s_i$, as desired. \square

This motivates the following definition.

Definition 5.3. Let s be a term and let $D = \Gamma \mid \Sigma \vdash s : T$ be a typing derivation for s . Define $\text{crest}(s, D)$ to be the term constructed from s by replacing all subterms s_i of s given a read-only type in D by $\text{seal } s_i$.

A crested term essentially seals any sub-term of the original term that is given a read-only type in a particular typing derivation. By definition, for any term s and typing derivation D for s , we have $s \leq \text{crest}(s, D)$. Moreover, a crested term can be given the same type as its original term as well.

LEMMA 5.4. *Let s be a term and let $D = \Gamma \mid \Sigma \vdash s : T$ be a typing derivation for s . Then $s \leq \text{crest}(s, D)$, and there exists a typing derivation showing that $\Gamma \mid \Sigma \vdash \text{crest}(s, D) : T$ as well.*

Now by progress – Theorem 4.11 – we have that for any well typed term s with typing derivation $D = \emptyset \mid \Sigma \vdash s : T$, its protected – crested – version $\text{crest}(s, D)$ will also step. By preservation – Theorem 4.9 – we have that $\text{crest}(s, D)$ either eventually steps to a value or runs forever, but never gets stuck. It remains to relate the reduction steps of $\text{crest}(s, D)$ to those of s , and specifically to show that if one reduces to some specific value and store, then the other also reduces to an equivalent pair of value and store.

We will do so by using the dynamic immutability safety properties proven in Section 3. System $F_{<:M}$ satisfies the same sealing-equivalence properties as System λ_M – seals do not affect reduction, except perhaps by introducing other seals. The following are analogues of Lemmas 3.4, 3.5, and 3.6 for System $F_{<:M}$.

LEMMA 5.5. *Let v be a value, σ_v be a store, t be a term such that $v \leq t$, and σ_t be a store such that $\sigma_v \leq \sigma_t$.*

If $\langle t, \sigma_t \rangle \longrightarrow \langle t', \sigma'_t \rangle$ then $v \leq t'$, $\sigma_v \leq \sigma'_t$, and $|t'| < |t|$.

LEMMA 5.6. *Let s, t be terms such that $s \leq t$ and let σ_s, σ_t be stores such that $\sigma_s \leq \sigma_t$. If $\langle s, \sigma_s \rangle \longrightarrow \langle s', \sigma'_s \rangle$ and $\langle t, \sigma_t \rangle \longrightarrow \langle t', \sigma'_t \rangle$ then:*

- (1) *Either $s \leq t'$, $\sigma_s \leq \sigma'_t$, and $|t'| < |t|$, or*
- (2) *$s' \leq t'$ and $\sigma'_s \leq \sigma'_t$.*

LEMMA 5.7. *Let s, t be terms such that $s \leq t$ and let σ_s, σ_t be stores such that $\sigma_s \leq \sigma_t$. If $\langle t, \sigma_t \rangle \longrightarrow \langle t', \sigma'_t \rangle$ then:*

- (1) Either $s \leq t'$, $\sigma_s \leq \sigma'_t$, and $|t'| < |t|$, or
- (2) There exists s' and σ'_s such that $\langle s, \sigma_s \rangle \longrightarrow \langle s', \sigma'_s \rangle$, $s' \leq t'$ and $\sigma'_s \leq \sigma'_t$.

Stepping back, we can see using Lemma 5.6 that one step of s to a term s' corresponds to *finitely many* steps of $\text{crest}(s, D)$; every step that $\text{crest}(s, D)$ takes either removes a seal or corresponds to a reduction step that s originally took. Hence $\text{crest}(s, D)$ eventually steps to a term t' such that $s' \leq t'$, preserving the desired equivalence of reduction between s and $\text{crest}(s, D)$. The following is a generalization of the previous statement to two arbitrarily chosen well-typed terms s and t satisfying $s \leq t$.

LEMMA 5.8. *Suppose $\emptyset, \Sigma \vdash \sigma_s$ and $\emptyset, \Sigma \vdash s : T$. Suppose $\langle s, \sigma_s \rangle \longrightarrow \langle s', \sigma'_s \rangle$. For $\sigma_s \leq \sigma_t$, and $s \leq t$, such that $\Gamma, \Sigma \vdash \sigma_s$ and $\Gamma, \Sigma \vdash t : T$, we have that $\langle t, \sigma_t \rangle \longrightarrow^* \langle t', \sigma'_t \rangle$ where $s' \leq t'$ and $\sigma'_s \leq \sigma'_t$.*

PROOF. From Theorem 4.11 we have that there exists a t' and σ'_t that $\langle t, \sigma_t \rangle \longrightarrow \langle t', \sigma'_t \rangle$. By Lemma 5.6 we have that either $s \leq t'$, $\sigma_s \leq \sigma'_t$, and $|t'| < |t|$, or that $s' \leq t'$ and $\sigma'_s \leq \sigma'_t$. If $s' \leq t'$ and $\sigma'_s \leq \sigma'_t$ we are done. Otherwise, observe that since $|t'| < |t|$, a seal was removed. This can only occur a finite number of times, as t and t' have at most a finite number of seals, so we can simply loop until we obtain a t' and σ'_t such that $s' \leq t'$ and $\sigma'_s \leq \sigma'_t$. Note that Preservation – Theorem 4.9 allows us to do so as each intermediate step t' can be given the same type $\Gamma \mid \Sigma \vdash t' : T$. \square

Finally, when s eventually reduces to a value v , we can use Lemma 5.5 to show that $\text{crest}(s, D)$ reduces to a similar value v' as well. Again, the following is a generalization of the previous statement to two arbitrarily chosen well-typed terms s and t satisfying $s \leq t$.

LEMMA 5.9. *Suppose $\emptyset, \Sigma \vdash \sigma_s$ and $\emptyset, \Sigma \vdash s : T$ such that s eventually reduces to a value v_s – namely, that $\langle s, \sigma_e \rangle \longrightarrow^* \langle v_s, \sigma'_s \rangle$ for some σ'_s .*

Then for any t such that $s \leq t$ and $\emptyset, \Sigma \vdash t : T$, we have that t eventually reduces to some value v_t , – namely $\langle t, \sigma_e \rangle \longrightarrow^ \langle v_t, \sigma'_t \rangle$, such that $v_s \leq v_t$ and $\sigma'_s \leq \sigma'_t$.*

PROOF. For each step in the multi-step reduction from $\langle s, \sigma_e \rangle \longrightarrow^* \langle v_s, \sigma'_s \rangle$ we can apply Lemma 5.8 to show that $\langle t, \sigma_t \rangle$ eventually reduces to $\langle t', \sigma'_t \rangle$ where $v_s \leq t'$ and $\sigma'_s \leq \sigma'_t$. Now by Theorem 4.11 and Lemma 5.5 we have that either t' is a value, in which case we are done, or that $\langle t', \sigma'_t \rangle$ steps to $\langle t'', \sigma'_s \rangle$ where $v_s \leq t''$. Again, we can only take a finite number of steps of this fashion as the rule which reduces $t' \longrightarrow t''$ can only be one that removed a seal, so eventually we obtain a value v_s such that $\langle t, \sigma_s \rangle \longrightarrow^* \langle v_t, \sigma'_t \rangle$ with $v_s \leq v_t$, and $\sigma'_s \leq \sigma'_t$, as desired. Again, note that Preservation – Theorem 4.9 allows us to do so as each intermediate step t' can be given the same type $\Gamma \mid \Sigma \vdash t' : T$. \square

Now from Lemma 5.9 we obtain our desired immutability safety results as a consequence – namely, given a well-typed term s that reduces to a value v_s , any references in s with a readonly type are never actually mutated, since they can be transparently sealed (which does not change the typing) to no ill effect. Formally, our main result is:

THEOREM 5.10. *Suppose s is a term, $D = \emptyset \mid \Sigma \vdash s : T$ is a typing derivation for s , and let σ_s be some initial store such that $\emptyset \mid \Sigma \vdash \sigma_s$. Then:*

- $\text{crest}(s, D)$ can be given the same type as s – $\emptyset \mid \Sigma \vdash \text{crest}(s, D) : T$.

Moreover, if $\langle s, \sigma_s \rangle \longrightarrow^* \langle v_s, \sigma'_s \rangle$, for some value v_s , then:

- $\text{crest}(s, D)$ will reduce to a value v_t – $\langle \text{crest}(s, D), \sigma_e \rangle \longrightarrow^* \langle v_t, \sigma'_t \rangle$, such that
- v_t and σ'_t are equivalent to v_s and σ'_s , modulo additional seals – namely, that $v_s \leq v_t$ and $\sigma'_s \leq \sigma'_t$.

Finally, it is useful to show that the converse result is also true; seals can be safely *removed* without affecting reduction. First note that seals themselves can be transparently removed without affecting the types assigned to the term.

LEMMA 5.11. *Suppose $\Gamma \mid \Sigma \vdash \text{seal } s : T$. Then $\Gamma \mid \Sigma \vdash s : T$.*

Moreover, the following analogue of Lemma 3.8 holds in System $F_{<:M}$.

LEMMA 5.12. *Suppose s and t are terms such that $s \leq t$. If $\langle t, \sigma_t \rangle \longrightarrow^* \langle v_t, \sigma'_t \rangle$ for some value v_t , then for any $\sigma_s \leq \sigma_t$ we have $\langle s, \sigma_s \rangle \longrightarrow^* \langle v_s, \sigma'_s \rangle$ such that $v_s \leq v_t$ and $\sigma'_s \leq \sigma'_t$.*

While Lemma 5.12 is enough to show when $s \leq t$, if t reduces to a value then so does s , we need Lemma 5.13 to reason about the types of s and v_s .

LEMMA 5.13. *Suppose s and t are terms such that $s \leq t$. If $\langle t, \sigma_t \rangle \longrightarrow^* \langle v_t, \sigma'_t \rangle$ for some value v_t , then for any $\sigma_s \leq \sigma_t$ we have $\langle s, \sigma_s \rangle \longrightarrow^* \langle v_s, \sigma'_s \rangle$ for some value v_s such that $v_s \leq v'_s$ and $\sigma'_s \leq \sigma'_t$. Moreover, $\Gamma \mid \Sigma \vdash s : T$ and $\Gamma \mid (\Sigma', \Sigma) \vdash v_s : T$ for some Σ' as well.*

PROOF. By Lemma 5.11 we can show that $\Gamma \mid \Sigma \vdash s : T$. By Lemma 5.12 we have that v reduces to some value v_s . By preservation – Theorem 4.9 we have that v_s has type T , as desired. \square

6 MECHANIZATION

Our mechanization of System $F_{<:M}$ is based on the mechanization of System $F_{<:}$ by Aydemir et al. [2008]. Our mechanization is a faithful model of System $F_{<:M}$ as described in this paper except for one case. To facilitate mechanization, reduction in our mechanization is done via explicit congruence rules in each reduction rule instead of an implicit rule for reducing inside an evaluation context, similar to how Aydemir et al. [2008] originally mechanize System $F_{<:}$ as well.

Proofs for all lemmas except for Theorem 5.10 and Lemmas 3.9, 5.2, and 5.4 have been mechanized using Coq 8.15 in the attached artifact. Theorem 5.10 and Lemmas 5.2, 5.4, and 5.13 have not been mechanized as they require computation on typing derivations which is hard to encode in Coq as computation on Prop cannot be reflected into Set. Lemma 3.9 has been omitted from our mechanization as it is hard to formally state, let alone prove, in a setting where reduction is done by congruence, though it almost follows intuitively from how the reduction rules are set up.

As the proofs of Lemmas 5.5, 5.6, 5.7, and 5.12 do not rely on any extra structure present in System $F_{<:M}$ over System λ_M , proofs for their System λ_M analogues Lemmas 3.4, 3.5, 3.6, and 3.8 have been omitted, as they can be recovered by erasing the appropriate cases from their System $F_{<:M}$ analogues.

7 RELATED AND FUTURE WORK

7.1 Limitations – Parametric Mutability Polymorphism

Unlike other systems, System $F_{<:M}$ does not support directly mutability polymorphism, neither through a restricted @polyread modifier as seen in Huang et al. [2012], nor through explicit mutability variables as seen in Gordon et al. [2012].

This is a true limitation of System $F_{<:M}$, however, we note that it is possible to desugar parametric mutability polymorphism from a surface language into a core calculus like System $F_{<:M}$. As Huang et al. [2012] point out in their work, parametric mutability polymorphism can be desugared via *overloading*, noting that overloading itself can be dealt with in a surface language before desugaring into a base calculus, as seen before with Featherweight Java [Igarashi et al. 2001].

For example, consider the following top-level parametric function, access, which is parametric on mutability variable M :

```
def access[M](z: [M] Pair[Pair[Int]]): M Pair[Int] = { z.first }
```

This function can be rewritten instead as two functions with the same name `access`, one taking in a regular, mutable pair, and one taking in a readonly pair:

```
def access(z: Pair[Pair[Int]]): Pair[Int] = { z.first }
def access(@readonly z: Pair[Pair[Int]]): @readonly Pair[Int] = { z.first }
```

Nested and first-class functions are a little trickier but one can view a polymorphic, first-class function value as a read-only record packaging up both overloads.

```
{
  access: (z: Pair[Pair[Int]]) => { z.first },
  access: (@readonly z: Pair[Pair[Int]]) => { z.first }
}
```

It would be interesting future work to see how one could add parametric mutability polymorphism to System $F_{<:M}$.

7.2 Future Work – Algorithmic Subtyping

The subtyping rules of System $F_{<:M}$ are fairly involved and it is difficult to see if an algorithmic subtyping system could be devised. We would conjecture that one could do so, using techniques from [Muehlboeck and Tate \[2018\]](#)'s integrated subtyping work, but nonetheless algorithmic subtyping for System $F_{<:M}$ remains an interesting and open problem.

7.3 Viewpoint Adaptation

Viewpoint adaptation has been used in reference immutability systems to denote the type-level adaptation which is enforced to guarantee transitive immutability safety. When a field $r.f$ is read from some record r , the mutability of the resulting reference needs to be adapted from *both* the mutability of r and from the type of f in the record itself. While this notion of adaptation was known as early as Javari [[Tschantz and Ernst 2005](#)], the term “viewpoint adaptation” was first coined by [Dietl et al. \[2007\]](#). They realized that this notion of adaptation could be generalized to arbitrary qualifiers – whether or not the type of a field read $r.f$ should be qualified by some qualifier $@q$ should depend on whether or not f 's type is qualified and whether or not r 's type is qualified as well – and used it to implement an *ownership* system for Java references in order to tame *aliasing* in Java programs.

7.4 Reference Immutability

Reference immutability has long been studied in the context of existing object-oriented languages such as Java and C#, and more recently has been studied in impure, functional languages like Scala.

roDOT [[Dort and Lhoták 2020](#)]: roDOT extends the calculus of Dependent Object Types [[Amin et al. 2016](#)] with support for reference immutability. In their system, immutability constraints are expressed through a type member field $x.M$ of each object, where x is mutable if and only if $M \leq \perp$, and x is read-only if and only if $M \geq \top$. Polymorphism in roDOT is out of all reference immutability systems closest to how polymorphism is done in System $F_{<:M}$. Type variables quantify over full types, and type variables can be further restricted to be read-only as in System $F_{<:M}$. Constructing a read-only version of a type, like how we use `readonly` in System $F_{<:M}$, is done in roDOT by taking an intersection with a bound on the type member M . For example, `inplace_map` from before could be expressed in roDOT using an intersection type to modify immutability on the type variable X :

```
def inplace_map[X](Pair[X]: pair, f: (X & {M :=> Any}) => X): Unit
```

Dort et al. also prove that roDOT respects immutability safety, but with different techniques than how we show immutability safety in System $F_{<:M}$. Instead of giving operational semantics with

special forms that guard references from being mutated, and relying on progress and preservation to imply static safety, they take a different approach and show instead that values on the heap that change during reduction must be reachable by some statically-typed mutable reference in the initial program. roDOT is a stronger system than System $F_{<,M}$, as in particular mutabilities can be combined. For example, one could write a generic getF function which reads a field f out of any record that has f as a field polymorphic over *both* the mutabilities of the record x and the field f :

```
def getF[T](x: {M: *, f : T}) : T & {M :> x.M} = x.f
```

Here, the return type of getF will give the proper, tightest, viewpoint-adapted type for reading $x.f$ depending on both the mutabilities of x and f . This is not directly expressible in System $F_{<,M}$ and can only be expressed using overloading:

```
def getF[T](x: @readonly {f : T}): @readonly T = x.f
def getF[T](x: {f : T}) : T = x.f
```

However, in contrast, roDOT is significantly more complicated than System $F_{<,M}$.

Immutability for C# [Gordon et al. 2012]: Of all the object calculi with reference immutability the calculus of Gordon et al. [2012] is closest to that of roDOT in terms of flexibility. Polymorphism is possible over *both* mutabilities and types in Gordon’s system, but must be done separately; type variables instead quantify over base types that have not been qualified with some immutability annotation, whether that be read-only or mutable. The inplace_map function that we discussed earlier would be expressed with both a base-type variable as well as a mutability variable:

```
def inplace_map[M, X](Pair[M X]: pair, f: @readonly X => M X): Unit
```

Like roDOT, Gordon’s system also allows for mutability annotations to be combined in types, in effect allowing viewpoint adaptation to be expressed at the type level using the mutability operator $\sim>$. For example, getF could be written as the following in Gordon’s system:

```
def getF[MS, MT, T, S <: {f : MT T}](x: MS S) : (MS  $\sim>$  MT) T = x.f
```

Unlike roDOT however, which allows for inferences to be drawn about the mutability of the type $(T \& \{M :> x.M\}).M$ depending on the bounds on T and S , the only allowable judgment we can draw about $MS \sim> MT$ is that it can be widened to @readonly. We cannot conclude, for example, that $MS \sim> MT <: M$ in the following, even though both $MS <: M$ and $MT <: M$:

```
def getF[M, MS <: M, MT <: M, T, S <: {f : MT T}](x: MS S) : (MS  $\sim>$  MT) T = x.f
```

Gordon et al. also demonstrate the soundness and immutability safety of their system but through an embedding into a program logic [Dinsdale-Young et al. 2013].

Javari [Tschantz and Ernst 2005]: Reference immutability was first modelled in the context of Java; Javari is the earliest such extension. In Javari’s formalization, Lightweight Javari, type variables X stand in for either other type variables, class types, and readonly-qualified class types. Unlike roDOT and System $F_{<,M}$, in Lightweight Javari, type variables **cannot** be further qualified by the readonly type qualifier. Lightweight Javari, however, does support parametric mutability polymorphism for class types, but does not support parametric mutability polymorphism directly on methods. Instead, limited parametric mutability method polymorphism in Javari, denoted with the keyword romaybe, is desugared using overloading into the two underlying methods handling the read-only case and the mutable case replacing romaybe in the source. Our earlier example, getF, can be written using romaybe as follows:

```

class HasF<T> {
    T f;
    maybe T getF() maybe { return f; }
}

```

However, this example is inexpressible in the core calculus Lightweight Javari, as `@readonly T` is ill-formed. As for safety, immutability safety is done in Lightweight Javari through a case analysis on how typed Lightweight Javari program terms can reduce. [Tschantz and Ernst 2005] claim that the soundness of Lightweight Javari reduces to showing the soundness of Lightweight Java, but no formal proof is given.

ReIm: [Huang et al. 2012]: ReIm simplifies Javari to enable fast, scalable mutability inference and analysis. Like Javari, ReIm supports two type qualifiers – `readonly` and `polyread`, where `readonly` marks a read-only type and `polyread` is an analogue of `maybe` from Javari. Like Lightweight Javari, and unlike `roDOT` and System $F_{<,M}$, ReIm restricts how qualifiers interact with generics. ReIm’s polymorphism model is similar to that of Gordon et al. [2012] – type variables range over unqualified types. However, ReIm has no mechanism for mutability polymorphism, and therefore `getF` cannot be written in ReIm at all. Unlike other related work, neither soundness nor immutability safety is proven to hold for ReIm.

Immutability Generic Java: [Zibin et al. 2007]: Immutability Generic Java is a scheme for expressing immutability using Java’s existing generics system. The type `List<Mutable>` denotes a *mutable* reference to a `List`, whereas the type `List<ReadOnly>` denotes a *read-only* reference to a `List`. Viewpoint adaptation is not supported, and transitive immutability must be explicitly opted into. For example, in the following snippet, the field `value` of `C` is *always* mutable. Transitive immutability must be explicitly opted into by instantiating `List` with the immutability parameter `ImmutOfC`.

```

class C<ImmutOfC> {
    List<Mutable /* ImmutOfC for transitivity */, Int> value;
}

```

Moreover, transitive immutability cannot be expressed at all over fields given a generic type. Type variables by the nature of how immutability is expressed in IGJ range over fully qualified types, and there is no mechanism for re-qualifying a type variable with a new immutability qualifier. For example, the mutability of `value` in any `Box` below depends *solely* on whether or not `T` is mutable. Hence the `value` field of a `Box` is mutable even if it was read through a read-only `Box` reference – that is, a reference of type `Box<ReadOnly>`.

```

class Box<ImmutOfBox, T> {
    T value;
}

Box<ReadOnly, List<Mutable,Int>> b = new Box(...);
b.value.add(10); // OK -- even though it mutates the underlying List.

```

7.5 Languages with Immutability Systems

Finally, some languages have been explicitly designed with immutability in mind.

C++: `const`-qualified methods and values provide limited viewpoint adaptation. Reading a field from a `const`-qualified object returns a `const`-qualified field, and C++ supports function and method dispatching based on the constness of its arguments [Stroustrup 2007]. Mutability

polymorphism is not explicitly supported but can be done with a combination of templates and overloading.

```

struct BoxedInt {
    int v{0};
};
template<typename T> struct HasF<T> {
    T f;
    T& getF() { return f; }
    const T& getF() const { return f; }
}

const HasF<BoxedInt> x;
x.getF() // Calls const qualified getF()
const BoxedInt& OK = x.f; // OK, as x.f is of type const BoxedInt.
BoxedInt& Bad = x.f; // Bad, discards const-qualifier.

```

In this example a C++ compiler would disallow `Bad` because the type of `x.f` has been adapted to a l-value of `const BoxedInt`. However, viewpoint adaptation does not lift to reference or pointer types in C++. For example, if instead we had a pointer-to-T in `HasF`:

```

template<typename T> struct HasF<T> {
    T* f;
}

BoxedInt b{5};
const HasF<BoxedInt> x{&b};
BoxedInt* NotGreat = x.f; // OK, as x stores a constant pointer to a mutable
    BoxedInt
NotGreat->v = 10; // Modifies b!

```

C++'s limited viewpoint adaptation gives `x.f` the type `BoxedInt * const`, which is a constant pointer to a mutable `BoxedInt`, not the type `BoxedInt const * const`, which would be a constant pointer to a *constant* `BoxedInt`. This allows the underlying field to be mutated.

D: In contrast to C++, where `const` becomes useless for pointer and reference fields, D supports full reference immutability and viewpoint adaptation with a *transitive* `const` extended to work for pointer and reference types [Bright et al. 2020]. Again, mutability polymorphism is not directly supported but can be encoded with D's compile-time meta-programming system.

Rust: In Rust, references are either *mutable* or *read-only*, and only one *mutable* reference can exist for any given value. Read-only references are transitive, like they are in System $F_{<,M}$, roDOT, and other reference immutability systems, and unlike C++. Here, in this example, we cannot write to `s3.f` as it `s3` is an read-only reference to `s2`, even though `s2.f` has type `&mut String`.

```

struct HasF<T> {
    f: T
}

fn main() {
    let mut s1 = String::from("hello");
    let s2 = HasF { f: &mut s1 };
    s2.f.push_str("OK");
    let s3 = &s2;
    s3.f.push_str("BAD");
}

```

```
}

```

Unlike other languages, though, the mutability of a reference is an intrinsic property of the reference type itself. Instead of having a type operator `readonly` that, given a reference type `T`, creates a read-only version of that reference type, Rust instead defines `&` and `&mut`, type operators that, given a type `T`, produce the type of a read-only reference to a `T` and the type of a mutable reference to a `T`, respectively. Here, in the following example, `s1` is a `String`, `s2` is a mutable reference to a `s1` – `&mut String`, and `s3` is a read-only reference to `s2` – `& (&mut String)`, where all three of `s1`, `s2`, and `s3` are stored at distinct locations in memory.

```
let s1 = String::from("hello");
let mut s2 = &s1;
let s3 = &s2;
```

As such, in Rust, one cannot create a read-only version of an existing reference type. This makes higher-order functions over references that are polymorphic over mutability, like `inplace_map` from above, inexpressible in Rust. However, if we instead had a `Pair` that owned its elements, we could write the following version of `inplace_map`:

```
struct Pair<T> {
    fst: T,
    snd: T
}

fn inplace_map<T>(p: &mut Pair<T>, f: fn (&T) -> T) {
    p.fst = f(&p.fst);
    p.snd = f(&p.snd);
}
```

Note, though, that in this setting, the elements `p.fst` and `p.snd` are embedded in the pair `p` and owned by it.

7.6 Type Qualifiers and Polymorphism

Foster et al. [1999] formalize a system for enriching types with qualifiers with support for polymorphism over both ground, unqualified types and qualifiers themselves. In this setting, `readonly` can be viewed as a type qualifier, similar to how C++’s `const` can be viewed as a qualifier in [Foster et al. 1999]. The resulting calculus which arises is similar to the calculus of [Gordon et al. 2012] restricted only to reference immutability qualifiers.

7.7 Contracts

Our approach to sealing references is similar to and was inspired by practical programming experience with Racket contracts – [Strickland et al. 2012]. Sealing, in particular, can be viewed as attaching a *chaperone* contract which raises an exception whenever the underlying chaperoned value is written to, and attaches a similar chaperone to every value read out of the value. For example, a dynamic reference immutability scheme for Racket vectors could be implemented with the following chaperone contract:

```
(define (chaperone-read vec idx v)
  (seal v))
(define (chaperone-write vec idx v)
  (error 'seal "Tried to write through an immutable reference."))

(define (seal v)
```

```
(cond
  [(vector? v) (chaperone-vector vec chaperone-read chaperone-write)]
  [else v]))
```

Strickland et. al. prove that chaperones can be safely erased without changing the behaviour of the underlying program when it reduces to a value. Our results on dynamic safety, Lemmas 3.4, 3.5, and 3.6 can be viewed as an analogue of [Strickland et al. 2012, Theorem 1] specialized to reference immutability. In this setting, our static immutability safety results show that a well-typed program will never raise an error by writing to a chaperoned vector.

8 CONCLUSION

We contributed a *simple* and *sound* treatment of reference immutability in System $F_{<}$. We show how a simple idea, *sealing* references, can provide dynamic immutability safety guarantees in an untyped context – System λ_M – and how soundness and System $F_{<}$ -style polymorphism can be recovered in a typed calculus System $F_{<,M}$ which builds on both System λ_M and System $F_{<}$. Our hope is to enable reference immutability systems in functional languages via this work, by giving simple soundness foundations in a calculus (System $F_{<}$) which underpins many impure functional languages today.

ACKNOWLEDGMENTS

We thank Yaoyu Zhao for his interesting discussions on reference immutability. We thank Alexis Hunt and Hermann (Jianlin) Li for their useful feedback on early drafts of this work. This work was partially supported by the Natural Sciences and Engineering Research Council of Canada and by an Ontario Graduate Scholarship. No seals were clubbed in the creation of this paper.

REFERENCES

- Nada Amin, Samuel Grütter, Martin Odersky, Tiark Ropf, and Sandro Stucki. 2016. The essence of dependent object types. *A List of Successes That Can Change the World: Essays Dedicated to Philip Wadler on the Occasion of His 60th Birthday* (2016), 249–272.
- Brian Aydemir, Arthur Charguéraud, Benjamin C. Pierce, Randy Pollack, and Stephanie Weirich. 2008. Engineering Formal Metatheory. In *Proceedings of the 35th Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages* (San Francisco, California, USA) (POPL '08). Association for Computing Machinery, New York, NY, USA, 3–15. <https://doi.org/10.1145/1328438.1328443>
- Walter Bright, Andrei Alexandrescu, and Michael Parker. 2020. Origins of the D Programming Language. *Proc. ACM Program. Lang.* 4, HOPL, Article 73 (jun 2020), 38 pages. <https://doi.org/10.1145/3386323>
- Luca Cardelli, Simone Martini, John C. Mitchell, and Andre Scedrov. 1991. An Extension of System F with Subtyping. In *Theoretical Aspects of Computer Software, International Conference TACS '91, Sendai, Japan, September 24-27, 1991, Proceedings (Lecture Notes in Computer Science, Vol. 526)*, Takayasu Ito and Albert R. Meyer (Eds.). Springer, 750–770. https://doi.org/10.1007/3-540-54415-1_73
- Werner Dietl, Sophia Drossopoulou, and Peter Müller. 2007. Generic Universe Types. In *ECOOP 2007 – Object-Oriented Programming*, Erik Ernst (Ed.). Springer Berlin Heidelberg, Berlin, Heidelberg, 28–53.
- Thomas Dinsdale-Young, Lars Birkedal, Philippa Gardner, Matthew Parkinson, and Hongsok Yang. 2013. Views: Compositional Reasoning for Concurrent Programs. *SIGPLAN Not.* 48, 1 (jan 2013), 287–300. <https://doi.org/10.1145/2480359.2429104>
- Vlastimil Dort and Ondřej Lhoták. 2020. Reference Mutability for DOT. In *34th European Conference on Object-Oriented Programming (ECOOP 2020) (Leibniz International Proceedings in Informatics (LIPIcs), Vol. 166)*, Robert Hirschfeld and Tobias Pape (Eds.). Schloss Dagstuhl–Leibniz-Zentrum für Informatik, Dagstuhl, Germany, 18:1–18:28. <https://doi.org/10.4230/LIPIcs.ECOOP.2020.18>
- Mattias Felleisen and D. P. Friedman. 1987. A Calculus for Assignments in Higher-Order Languages. In *Proceedings of the 14th ACM SIGACT-SIGPLAN Symposium on Principles of Programming Languages* (Munich, West Germany) (POPL '87). Association for Computing Machinery, New York, NY, USA, 314. <https://doi.org/10.1145/41625.41654>

- Jeffrey S. Foster, Manuel Fähndrich, and Alexander Aiken. 1999. A Theory of Type Qualifiers. In *Proceedings of the ACM SIGPLAN 1999 Conference on Programming Language Design and Implementation* (Atlanta, Georgia, USA) (PLDI '99). Association for Computing Machinery, New York, NY, USA, 192–203. <https://doi.org/10.1145/301618.301665>
- Colin S. Gordon, Matthew J. Parkinson, Jared Parsons, Aleks Bromfield, and Joe Duffy. 2012. Uniqueness and Reference Immutability for Safe Parallelism. In *Proceedings of the ACM International Conference on Object Oriented Programming Systems Languages and Applications* (Tucson, Arizona, USA) (OOPSLA '12). Association for Computing Machinery, New York, NY, USA, 21–40. <https://doi.org/10.1145/2384616.2384619>
- Philipp Haller and Ludvig Axelsson. 2017. Quantifying and Explaining Immutability in Scala. *Electronic Proceedings in Theoretical Computer Science* 246 (apr 2017), 21–27. <https://doi.org/10.4204/eptcs.246.5>
- Wei Huang, Ana Milanova, Werner Dietl, and Michael D. Ernst. 2012. ReIm and ReImInfer: Checking and Inference of Reference Immutability and Method Purity. In *Proceedings of the ACM International Conference on Object Oriented Programming Systems Languages and Applications* (Tucson, Arizona, USA) (OOPSLA '12). Association for Computing Machinery, New York, NY, USA, 879–896. <https://doi.org/10.1145/2384616.2384680>
- Atsushi Igarashi, Benjamin C. Pierce, and Philip Wadler. 2001. Featherweight Java: A Minimal Core Calculus for Java and GJ. *ACM Trans. Program. Lang. Syst.* 23, 3 (may 2001), 396–450. <https://doi.org/10.1145/503502.503505>
- Fabian Muehlboeck and Ross Tate. 2018. Empowering Union and Intersection Types with Integrated Subtyping. *Proc. ACM Program. Lang.* 2, OOPSLA, Article 112 (oct 2018), 29 pages. <https://doi.org/10.1145/3276482>
- John C. Reynolds. 1997. *Design of the Programming Language Forsythe*. Birkhäuser Boston, Boston, MA, 173–233. https://doi.org/10.1007/978-1-4612-4118-8_9
- T. Stephen Strickland, Sam Tobin-Hochstadt, Robert Bruce Findler, and Matthew Flatt. 2012. Chaperones and Impersonators: Run-Time Support for Reasonable Interposition. In *Proceedings of the ACM International Conference on Object Oriented Programming Systems Languages and Applications* (Tucson, Arizona, USA) (OOPSLA '12). Association for Computing Machinery, New York, NY, USA, 943–962. <https://doi.org/10.1145/2384616.2384685>
- Bjarne Stroustrup. 2007. *The C++ programming language - special edition* (3. ed.). Addison-Wesley.
- Matthew S. Tschantz and Michael D. Ernst. 2005. Javari: Adding Reference Immutability to Java. *SIGPLAN Not.* 40, 10 (oct 2005), 211–230. <https://doi.org/10.1145/1103845.1094828>
- Yoav Zibin, Alex Potanin, Mahmood Ali, Shay Artzi, Adam Kiezun, and Michael D. Ernst. 2007. Object and Reference Immutability Using Java Generics. In *Proceedings of the the 6th Joint Meeting of the European Software Engineering Conference and the ACM SIGSOFT Symposium on The Foundations of Software Engineering* (Dubrovnik, Croatia) (ESEC-FSE '07). Association for Computing Machinery, New York, NY, USA, 75–84. <https://doi.org/10.1145/1287624.1287637>