

Verifying Stateful Programs with Substructural State and Hoare Types

Johannes Borgström*

Uppsala University
borgstrom@acm.org

Juan Chen

Microsoft Research, Redmond
juanchen@microsoft.com

Nikhil Swamy

Microsoft Research, Redmond
nswamy@microsoft.com

Abstract

A variety of techniques have been proposed to verify stateful functional programs by developing Hoare logics for the state monad. For better automation, we explore a different point in the design space: we propose using affine types to model state, while relying on refinement type checking to prove assertion safety.

Our technique is based on verification by translation, starting from FX, an imperative object-based surface language with specifications including object invariants and Hoare triple computation types, and translating into Fine, a functional language with dependent refinements and affine types. The core idea of the translation is the division of a stateful object into a pure value and an affine token whose type mentions the current state of the object. We prove our methodology sound via a simulation between imperative FX programs and their functional Fine translation.

Our approach enables modular verification of FX programs supported by an SMT solver. We demonstrate its versatility by several examples, including verifying clients of stateful APIs, even in the presence of aliasing, and tracking information flow through side-effecting computations.

1. Introduction

Several recent papers propose a verification methodology for stateful functional programs by developing Hoare logics for the state monad (Nanevski et al. 2006; Swierstra 2009; Borgström et al. 2010). Tools based on this approach are known to be powerful. For example, the Ynot tool has been used to carry out interactive proofs of correctness for programs that manipulate B+ trees (Malecha et al. 2010), a pointer structure with tricky sharing properties.

While such successes make a good case for interactive machine-assisted proofs in a Hoare logic for the state monad, this paper explores a different point in the design space of a program logic for stateful functional programs. We are motivated primarily by a desire to automate the modular verification of functional programs that make use of local state. We aim to develop a verification methodology based on classical first-order logic, for which powerful off-the-shelf solvers already exist. (In contrast, tools like Ynot work with separation logic to recover modularity, but automation for separation logic remains a significant challenge, although the situation is improving (Chlipala et al. 2009).)

Our work considers a well-known alternative to monadic state as the basis of a verification methodology: substructural state, by which we mean state modeled using linear (use-once) or affine (use-at-most-once) types (Wadler 1990). Our insight is that the problem of verifying stateful functional programs can be factored into two pieces. First, we can use an affine type system to model

the stateful behavior of the program in a purely functional style. Then, to prove the safety of assertions in the program, we can rely on automated refinement type checking for a first-order classical logic. Affine types partition the state of the program into a number of disjoint pieces, thus yielding a modular verification procedure without necessitating the use of the separation logic connectives.

In addition to improved automation, modeling state using affine types is attractive since it enables combinations of Hoare logic with other program verification disciplines. We work out one such example in detail in §5.1, where we combine our basic approach with a discipline of fractional capabilities to control aliasing.

Concretely, our work is grounded in the context of Fine (Swamy et al. 2010), a purely functional language with a type system based on a combination of dependent refinements and affine types. Our contributions include the following:

- We present FX (§3), an extension of the term syntax of Fine with stateful commands including object allocation, deletion and mutation. We also extend the type language with Hoare types (Nanevski et al. 2006) to give specifications (in the style of Hoare triples) to stateful code in FX.
- We show how to translate FX programs to Fine (§4), using Fine’s affine types to model state, and dependent refinements for the translation of Hoare types. The core idea is to divide a stateful object into a pure value and an affine token whose type mentions the current state of the object. We prove that our translation is a simulation, which implies that refinement type checking in Fine guarantees assertion safety for FX.
- We demonstrate the effectiveness of our approach on a variety of examples (§5), which, although simpler than the B+ trees verified interactively in Ynot, are still known to be challenging. Our examples include a conference management program previously verified for authorization properties by Swamy et al. (2010); a client of a stateful API of collections and iterators studied by Bierhoff and Aldrich (2007); and an encoding of information flow tracking suitable for use with programs that may leak information via side effects (an enhancement of a technique studied by Swamy (2008)).

Although the results of this paper are limited (primarily) to first-order programs, we discuss ongoing work that naturally extends our approach to the higher-order case in §5.3. The remainder of this section presents a short overview of our work, highlighting some of the technical challenges we solve.

1.1 Marrying substructural and dependent types

Linear types have been known as an effective tool for modeling state for the last 20 years (Wadler 1990), roughly as long as the monadic approach. Yet, with a few notable exceptions, verification tools and programming languages have only seldom adopted linear

*The first author was employed by Microsoft Research, Cambridge, for the duration of his work on this paper.

types as model of state. One reason for this, particularly in the type-based approach to verification, is the subtle interaction between linear and dependent types. The metatheory of an arbitrary mixture of these types is usually considered intractable. For example, Linear LF (Cervesato and Pfenning 2002), perhaps the canonical presentation of linearity and dependency, restricts linear values from appearing as indices in types. Fine has a similar restriction on affine values, placing seemingly severe limitations on expressiveness.

For a sense of the difficulty, consider implementing a function `incr` that increments an affine integer `x:alnt`. One might write it as $\lambda x:alnt. x+1$, consuming the original integer `x` and threading back the incremented version to the caller. We would like to give this function a more precise type, such as $x:alnt \rightarrow \{y:alnt \mid y > x\}$. This is the type of a dependent function from affine integers `x` to affine integers `y`, where the refinement type to the right of the arrow states that the returned value `y` is greater than `x`. However, this type is illegal in Linear LF and Fine, because the refinement formula $y > x$ is interpreted as a type, and the affine values `x` and `y` cannot be used at the type level.

A key technical contribution of our work is to surmount this difficulty by using a combination of value-indexed types and affine capabilities (Walker et al. 2000) to permit unrestricted refinements of stateful values. In our solution, the type of an increment function, say `incr2`, could be $x:int \rightarrow \text{Token } x \rightarrow (y:\{y:int \mid y > x\} * \text{Token } y)$, where `Token :: int \Rightarrow A` is a dependent type constructor that constructs an affine type from an `int`-typed value. (The kind of affine types is `A`.) The type of `incr2` indicates that to increment `x`, client code must present a token for `x`. The `incr2` function destructs the token for `x`, produces `y`, and a new token for `y`. By making `Token x` affine, we ensure that the client can no longer use `x` in a context that expects a mutable integer. By indexing the type of the token with the value of `x`, we ensure that a token of one value cannot be confused with a token for another. Finally, by separating `x` from its token, we can use `x` in refinement formulas, e.g., in the formula $y > x$ in the return type of `incr2`.

1.2 Shielding programmers from affinity

Clearly, requiring programmer to directly manipulate tokens is far from optimal—mutable values and their tokens have to be threaded through the program in a functional style. Additionally, exposing affinity to the programmer is also questionable, since it limits the use of library code—data structures that contain affine values become affine themselves, making it hard to use, say, a standard non-affine lists to hold affine values.

To counter these difficulties, we present FX, a surface syntax with imperative commands that hides the low-level plumbing of tokens and affinity from the programmer. To illustrate FX, consider a snippet of an FX implementation of ConfWeb, a verified conference management tool we had previously implemented directly in Fine.

```

1 type review = {txt:string; u:user}
2 val upd_review: r:review  $\rightarrow$  x:string  $\rightarrow$ 
3   {s1 T} ..unit {(s2) s2(r).txt=x && s2(r).u=s1(r).u}
4 let upd_review r x = r.txt := x

```

Here, we define a mutable object type `review` of reviews, with fields for review text and the reviewer `u`. Line 4 shows the implementation of a function `upd_review` that updates a review `r` by assigning to its `txt` field. FX programmers can use imperative commands like assignment rather than threading mutable objects through the program. Mutable objects can also be placed in standard data structures, without having to re-implement libraries to deal with affinity.

Specifications in FX are written using Hoare types, a notation resembling Hoare triples. Lines 2-3 show the specification of `upd_review`, with the form $x:t1 \rightarrow \{(s1) \Phi\} r:t2 \{(s2) \Psi\}$, where the formula Φ is a pre-condition on the function argument `x:t1` and pre-state `s1`; `r:t2` is the name and type of the return value; and Ψ is a

post-condition on the argument `x`, the pre-state `s1`, the return value `r`, and the post-state `s2`.

The main technical development of this paper is a translation from FX to Fine that introduces affine tokens and threads stateful values and their tokens through the program. The main results are Theorem 1, which establishes that the translation is a simulation, and Corollary 2, which establishes progress for FX programs that translate to well-typed Fine programs.

2. Overview

We begin with a brief review of the Fine programming language, then discuss ConfWeb, a conference management application we previously implemented in Fine. This application motivates, in part, the need for better support for imperative features in Fine. §2.2 presents the design of FX using examples from ConfWeb. §2.3 presents informally our methodology of verifying FX programs via translation to Fine.

2.1 A brief review of Fine

Fine is an experimental, purely functional programming language on the .NET platform. The principal novelty of Fine is in its type system, which is designed to support static verification of safety properties via a mixture of refinement and substructural types. Our work on Fine has, thus far, focused primarily on verifying security properties, including properties related to authorization policies and information flow controls. However, by design, the type system of Fine has little in it that is security-specific. This paper is a step towards putting Fine to use for the general-purpose verification of programs that mix both functional and imperative idioms. We discuss several elements of Fine’s design next—for details, consult our prior papers (Swamy et al. 2010; Chen et al. 2010).

Value-indexed types. Types in Fine can be indexed both by types (e.g., `list int`) as well as by values. For example, `array int 17` represent the type of an array of 17 integers, where the index `17:nat` is a natural number value. Fine prohibits non-value expressions from appearing as type indices, in effect forbidding type-level computations. This restriction limits expressiveness, but considerably simplifies Fine’s metatheory and implementation.

Dependent function types. Functions in Fine are, in general, given dependent function types, i.e., their range type depends on their argument. Dependent function types are written $x:t \rightarrow t'$, where the formal name `x` of the parameter of type `t` is in scope in `t'`. For example, the type of a function that allocates an array of `n` integers can be given the type $n:nat \rightarrow \text{array int } n$. When a function is non-dependent, we simply drop the formal name.

Refinement types. A refinement type in Fine is written $\{x:t \mid \phi\}$, where ϕ is a formula in which `x` may appear free. Formulas are drawn from the same syntactic category as types, although, for readability, we typeset formulas differently and use distinct meta-variables (ϕ versus `t`). In practice, types are refined by formulas from a first-order logic with equality, extended with user-defined predicates. For example, we may give the following type to a list permutation: $\forall \alpha. l:\text{list } \alpha \rightarrow \{m:\text{list } \alpha \mid \forall x:\alpha. \text{In } x \, l \Leftrightarrow \text{In } x \, m\}$. Universal quantifiers in formulas are represented using dependent arrows and existential quantifiers using dependent pairs; `In` is a user-defined predicate for list membership; and connectives like \Leftrightarrow are represented using indexed types (we elaborate below).

Refinement type checking. A refinement type $\{x:t \mid \phi\}$ is inhabited by values `v:t`, for which $\phi[v/x]$ is derivable. Derivability is defined with respect to assumptions induced by the program context (e.g., equalities due to pattern matching) as well as a set of assumptions that axiomatize user-provided predicates. We formalize this using an LCF-style (Milner 1979) kernel for Fine, which contains the inference rules for a classical first-order logic with equality, ex-

tended with constructors corresponding to user-provided axioms. As such, user-provided axioms must be used with care, since faulty axioms compromise soundness. Derivability is decided by relying on Z3 (de Moura and Bjorner 2008), an SMT solver. Formally, refinement types are viewed as Σ -types. However, we include a program transformation (somewhat similar to the coercions used by Sozeau (2006)) that systematically inserts pack/unpack operations so as to equip refinement types with a subtyping relation that allows programmers to view $\{x:t \mid \phi\}$ as a subtype of t .

Affine types. Dependent refinements in Fine bear close resemblance to constructs found in related languages like F7 (Bengtson et al. 2008) and Sage (Flanagan 2006), despite several technical differences. A substantial difference in Fine however, is the addition of affine types. The interaction between dependent and affine types in Fine is strictly regulated by the *no-affine-indices* restriction—this prohibits the use of values with affine types as type indices. As such, Fine is closely related to Linear LF (Cervesato and Pfenning 2002), which integrates linear and dependent types, with a similar restriction that prevents linear resources being mentioned in types. This paper shows that even with this restriction, the combination of dependent and affine types available in Fine provides a powerful set of primitives for verifying programs that use mutable local state.

Kind language. To enforce the no-affine-indices restriction (as well as to keep account of type constructors) Fine employs a system of kinds. The (slightly simplified) syntax of kinds is shown below:

$$\text{kinds } k ::= \star \mid A \mid \alpha :: k \Rightarrow k \mid x:t \Rightarrow k$$

Types in Fine are divided into two basic kinds: \star , the kind of normal non-affine types, and A , the kind of affine types. A \star -kinded type can be coerced to the A universe using the modality i , e.g., $\text{int}::\star$, while $j\text{int}::A$. Type constructors are given arrow kinds, which come in two flavors. The first, $\alpha::k \Rightarrow k'$ is the kind of type functions that construct a k' -kinded type from a k -kinded type α . Just as at the term level, type-level arrows are dependent—the type variable α can appear free in k' . Type functions that construct value-indexed types are given a kind $x:t \Rightarrow k$, where x names the formal of type t and x can appear free in k . In both cases, when the kind is non-dependent, we simply drop the formal name. The no-affine-indices manifests itself as a restriction on kinds $x:t \Rightarrow k$, where the domain type t must have kind \star . For example, the kind of `list` is $\star \Rightarrow \star$; the kind of the value-indexed `array` constructor is $\star \Rightarrow \text{nat} \Rightarrow \star$; the kind of the propositional connective `And` is $\star \Rightarrow \star \Rightarrow \star$; the kind of the user-defined predicate `In` is $\alpha::\star \Rightarrow \alpha \Rightarrow \text{list } \alpha \Rightarrow \star$.

Module system. Aside from the core type system, Fine has a simple module system allowing us to define certain types private to a module, which in effect forces clients of the module to view values of these private types abstractly. Fine uses the module system to prove secrecy and authenticity properties. In this paper, the module system comes in handy for the definition of unforgeable capabilities and the corresponding authenticity property plays a crucial role in proving the correctness of our translation from FX to Fine.

Proof-carrying compilation. Finally, it is worth mentioning that Fine is compiled in a proof-carrying style to DCIL, .NET bytecode enhanced with dependent and affine types in a backwards-compatible way. In addition to certification, compilation to DCIL allows Fine to interoperate easily with the other .NET languages. By virtue of the translation of the forthcoming sections, FX too enjoys a proof-carrying translation to DCIL.

2.2 ConfWeb: A first taste of FX

Our prior work on Fine included the development of an application ConfWeb, a conference management tool based on the Continue server (Krishnamurthi 2003). Continue’s behavior is governed by an authorization policy, modeled with liberal use of object-

```

1 type user=int
2 type NoConflict :: user ⇒ user ⇒ *
3 assume nc_sym: ∀u1, u2. NoConflict u1 u2 ⇒ NoConflict u2 u1
4 val checkNC: u1:user → u2:user → {b:bool | b=true ⇔ NoConflict u1 u2}
5 type review = {[txt:string; u:user]}
6 type paper = {[authors:list user; revs:list review | ∀u:user, r:review.
7             (In u self.authors && In r self.revs) ⇒ NoConflict u r.u]}
8 (* Allocation function with a Hoare type *)
9 val mk_review: txt:string → u:user → {(s1 ⊤) r:review
10                                     {(s2) s2(r).txt=txt && s2(r).u=u}}
11 let mk_review txt u = new review {[txt=txt; u=u]}
12 (* Pure library function can operate on a list of mutable objects *)
13 val for_all: ∀α::*, P:α ⇒ * .
14             f:(x:α → {b:bool | b=true ⇔ P x}) → !:list α →
15             {b:bool | b=true ⇔ ∀x:α. In x l ⇒ P x}
16 (* Impure function mutates a paper; also given a Hoare type *)
17 val add_review: p:paper → r:review →
18             {(s) not(In s(r).revs)} o:option review
19             {(t) (In s(r).revs || t(o)=Some s(r))&& ...}
20 let add_review p r =
21   if for_all <user, NoConflict r.u> (checkNC r.u) p.authors
22   then let tl = (p.revs := Nil) in p.revs := Cons r tl; None
23   else Some r

```

Figure 1. Invariants on mutable objects in ConfWeb

orientation and mutable state in Alloy (Jackson 2002). A basic data structure in the model is a *paper*, an object with mutable fields holding a paper submission and list of its associated reviews. An invariant on *paper* ensures that the reviews of an paper are from reviewers not conflicted with the paper’s authors. Our prior implementation of ConfWeb modeled many of these stateful features, but some features (including invariants on mutable objects like *paper*) proved to be too cumbersome to implement directly in Fine.

In this section, we introduce FX, a surface syntax for Fine with direct support for mutable objects. FX considerably simplifies programming with mutable objects and makes an implementation of ConfWeb more faithful to its original specification. Subsequent sections shows how mutable state can be systematically translated away via translation to purely functional Fine programs.

Figure 1 shows a fragment of ConfWeb written in FX. The key data structure in ConfWeb is `paper` (line 6), a mutable object with fields containing the authors of a paper and its reviews (in addition to other fields, not shown). We verify that this program preserves the no-conflict invariant on `paper`. The full program translates to a 109-line Fine program that is verified automatically in 8 seconds.

Mutable objects and invariants. FX extends Fine with a notation for records with mutable fields—we call these types *objects*, although our core calculus for FX does not include other OO features like inheritance. Object types are of the form $\{[f1:t1; \dots; fn:tn \mid \phi]\}$, where each f_i is a mutable field of type t_i , and the formula ϕ is an object invariant in which the name *self* refers to the object itself. Note that the values of non-object types like `string` and `int` are immutable.

The paper object and its invariant. Line 6 defines `paper`, an object with a field `authors` (a list of immutable integers, standing for user ids); and a field `revs`, a list of objects representing the reviews of a paper. The invariant on `paper` is a formula which uses two user-defined predicates. The first predicate, `In`, stands for list membership. The predicate `NoConflict` is a binary predicate to reflect the absence of conflict-of-interest between a pair of users. Line 2 shows the kind of `NoConflict` and line 3 shows a user-provided axiom that asserts that `NoConflict` is a symmetric relation. Rather than provide further axioms to define `NoConflict`, at line 4 we show the signature of a trusted external function `checkNC`, where we use dependent refinement types in Fine to assert that `checkNC` decides the `NoConflict` relation. The implementation of this function consults

a database that records conflict-of-interest declarations by authors and PC members of the conference—interoperability between Fine and the rest of .NET makes it easy to implement this function in F# and to call it from FX.

Impure functions and Hoare types. At line 11, we show the implementation of a function `mk_review`, which allocates a new `review` object. Its specification at lines 9-10 shows the type of an impure function, which, in general, has the form $x:t \rightarrow \{(s1) \Phi\} r:t' \{(s2) \Psi\}$, as explained in §1.2. For `mk_review`, the pre-condition Φ is trivial— \top , which always holds. The post-condition Ψ contains terms of the form $s2(r)$ —this stands for the value of the mutable object r in the state $s2$. In this case, $s2(r)$ is a record value of the object type `review`. The post-condition of `mk_review` states that the fields of the returned object r contain the arguments passed to the function.

Interacting with libraries. Next, at lines 13-15, we show the type of `for_all`, a library function on lists. This is a higher-order function whose argument f is a (pure) boolean-valued function over the α -typed elements of its second argument l , where f decides some property P of the list elements. The type we give to `for_all` is polymorphic in both α and P . The return type indicates that `for_all` returns `true` if, and only if, f returns `true` on each element of l . The type also indicates that both f and `for_all` are pure, since neither of them has a Hoare type. This pure function can be programmed in Fine and verified against its specification. The implementation is entirely standard—we omit it here for brevity.

Mutation and controlled aliasing. Finally, at lines 17-23, we show the type and implementation of a function `add_review` that tries to add a review to a paper. Informally, the implementation checks that the review’s author is not conflicted with any of the paper’s authors and, if the check succeeds, adds the review to the paper and returns `None`. Otherwise, `add_review` leaves the paper unchanged and returns the review to the caller. The type shows an impure function, where the precondition requires the review r to not be present in the paper p ’s `revs` list; the returned value is an `option review`, where the post-condition shows that the review r has been stored in the paper p , or that r is contained in returned option value.

The implementation of `add_review` reveals several subtleties of FX. First, to properly model stateful invariants on objects, FX forbids constructing aliases to mutable objects. (§5.1 shows how to relax this restriction.) If a mutable object o is stored in another object c ’s f field, then a field projection $c.f$ constructs an alias to the object o . As a result, FX does not permit unrestricted field projection. Instead, we provide other operations to emulate field projections (which get compiled to a small set of function calls in Fine).

At line 22, we call `for_all`, providing explicit instantiations for its type and predicate parameters, α and P respectively. We apply `checkNC` to $r.u$; project $p.authors$; and pass both to `for_all`. The projections here are safe since the type of `for_all` is pure, and hence does not create any unsafe aliases to p . Additionally, since the `user` type is immutable, it is always safe to project $r.u$. If the no-conflict check fails, we return r to the caller. Otherwise, we update the paper, adding the new review r to $p.revs$ and returning `None`. The update uses FX’s swap command, $c.f := v$, which replaces the f field of c with the value v and returns the original contents of $c.f$. We first swap out the original list of reviews from p ; add our new review to it; and then swap the extended list back in. Note that it would also be safe to project $p.revs$ and use it in a normal assignment $p.revs := Cons r p.revs$, since this creates no aliases. Projections of fields that contain primitive or immutable types are also safe.

In both branches of `add_review`, the mutable object r is stored within another object—in the `then`-branch, within p ; in the `else`-branch, in the returned option object. Since creating unrestricted aliases to mutable objects is forbidden, and the caller retains a reference to r , we need to indicate to the caller that its reference

r is invalidated. We use the absence of a sub-term $r(r)$ in the post-condition to indicate that r is invalidated in the post-state, i.e., that it may no longer be used by the caller. By contrast, the type of `upd_review` shows that the review r is still accessible to the caller by mentioning $s2(r)$ in its post-condition. Currently, FX cannot express that a reference has been conditionally consumed, e.g., returning a boolean success code b from `add_review` to indicate in the post-condition that r is consumed only if $b=true$.

2.3 Verifying FX programs via translation to Fine

This section sketches the translation of FX to Fine. The main idea is to thread record values corresponding to mutable objects through the program in a purely functional style. For this functional translation to be sound, we also thread affine capabilities for each object to render stale record values unusable. As discussed in §1, by making the capabilities affine, while giving record values \star -kinded types, we circumvent the no-affine-indices restriction and give precise refinement types (corresponding to object invariants, pre- and post-conditions) to the functional translation of imperative FX code. As a result, we can rely on Fine’s automated refinement type checking system to verify FX programs.

A module for FX primitives. The translation is organized around a Fine module, `Prims`, which collects definitions of FX primitives. Certain types will be private to `Prims`, and the metatheory of Fine (specifically, the value abstraction theorem provided by its module system) guarantees that clients of `Prims` treat these types abstractly.

Objects and constructors. An object with mutable fields in FX, say $r:review$, where `review` = $\{\{txt:string; u:user\}\}$ from Figure 1, is represented (roughly) in Fine as a pair of 1) a record value $r:review$ where `review` is a record type in Fine, $\{\{txt:string; u:user\}\}$; and 2) a token value $v:Token\ r$, used as a capability for r , where $Token::\alpha::\star \Rightarrow \alpha \Rightarrow A$, is a constructor of an affine, value-indexed type. Given an object type declaration in FX, we generate several public functions in `Prims`, corresponding to constructors, field updates, etc., as shown below. Imperative commands in FX get translated into calls to these functions.

```

1 private type Token::  $\alpha::\star \Rightarrow \alpha \Rightarrow A = | MkToken: x:\alpha \rightarrow Token\ x$ 
2 private type review =  $\{\{txt:string; u:user\}\}$ 
3 val mkReview: txt:string  $\rightarrow u:user \rightarrow unit \rightarrow$ 
4    $(r:rev * Token\ r * \{\_:unit | r.txt=txt \&\& r.u=u\})$ 
5 val reviewUpdTxt: r:review  $\rightarrow txt:string \rightarrow (Token\ r * unit) \rightarrow$ 
6    $(r':review * Token\ r' * \{\_:unit | r'.txt=txt \&\& r'.u=r.u\})$ 
7 val reviewReadTxt: r:review  $\rightarrow (Token\ r * unit) \rightarrow$ 
8    $(txt:string * Token\ r * \{\_:unit | r.txt=txt\})$ 

```

The function `mkReview` is the constructor function of `review`. In addition to the constructed record value r and its token, it returns a unit value refined with a formula recording information about the contents of r . The `reviewUpdTxt` function updates the `txt` field. Its arguments include r the object to be updated and its token; the return type shows r' and its token, where the refined unit shows that r' differs from r in only its `txt` field. We also show `reviewReadTxt`, a pure function that serves as a projection function for the `txt` field. Note that all three functions take seemingly redundant `unit`-typed arguments and return refined `units`—we include these because, as we shall see, they help keep our translation uniform.

Finally, the `Token` type, which has a single constructor `MkToken`, is declared private to `Prims` to ensure that its values cannot be manufactured directly by clients. Likewise, the `review` record type is also private. Stateful operations on $r:review$ (e.g., updating one of its fields) require a client to present a $v:Token\ r$ attesting to the validity of r ; the operation consumes the token, rendering r invalid for further use in stateful operations. Pure operations like `reviewReadTxt` thread the token back to the caller for further use.

Hoare types. The Fine type below is the translation of the FX type $r:review \rightarrow \{(s1) s1(r).u=1\} x:int \{(s2) s2(r).u=s1(r).u+x\}$.


```

1 r:review → (Token r * {_:unit | r.u=I}) →
2   (x:int * r':review * Token r' * {_:unit | r'.u=r.u+x})

```

This type shows a function that takes as arguments tokens for all the variables x where $s1(r)$ is a subterm in the pre- or post-condition; in this case a token for r , since only $s1(r)$ is mentioned. In addition to the token, we have a unit argument refined with the translation of the pre-condition. The return type of the function shows x :int the value returned by the impure FX computation. We also include in the return type all variables y where $s2(y)$ is a subterm of in the post-condition (i.e., r); tokens for these values; and, the translation of the post-condition formula expressed as a refinement on a unit value. The scoping rules of dependent functions, pairs (aka sums), and refinement types naturally allow the post-condition to describe properties of the initial value of r in the pre-state, the returned value x , and the value of r in the post-state, (i.e., r').

Adapting libraries to work with mutable objects. A significant advantage of our token-based translation is that it lends itself naturally to the re-use of types and libraries that were originally constructed for purely functional programs. For example, we show below how the standard list library can be adapted to work with lists of mutable objects. The `list` type shown below is the standard type for a list of non-affine values. FX programs are free to use the constructors of `list` directly to store immutable values within them, and to pattern match on lists to extract immutable values as well. But, we require a bit more care to store mutable objects within lists, since the constructed list would capture a reference to the object. For this, we provide functions in `Prims` that construct and destruct lists while managing tokens appropriately.

```

1 type list :: * ⇒ * = Nil : list α
2   | Cons : α → list α → list α
3 val mkCons: hd:α → tl:list α → (Token hd * Token tl * unit) →
4   (l:list α * Token l * {_:unit | l=Cons hd tl})
5 val unCons: l:list α → (Token l * {_:unit | ∃hd,tl. l=Cons hd tl}) →
6   (hd:α * tl:list α * Token hd * Token tl * {_:unit | l=Cons hd tl})
7 val mkNil: unit → (l:list α * Token l * {_:unit | l=Nil})

```

The function `mkCons` takes an object value $hd:\alpha$ and a list of object values $tl:list \alpha$; tokens for each of them; consumes the tokens for hd and tl and returns a new list l , a token for l , and a refined unit capturing the structure of l . While it is possible for a client program to pattern match against l to extract its components, to mutate any object in l , a client must destruct l by presenting l and its token to `unCons`; proving a pre-condition that l is indeed a `Cons`-cell; and receiving the components of l , tokens for each of them, and a post-condition relating the components to the original l .

By adopting a token-based approach, not only have we been able to circumvent the no-affine-indices restriction, but we have been able to store mutable objects in standard lists, operate over these lists using the standard implementations of `map`, `fold`, `for_all`, etc., since these functions are pure. Impure functions that need to explicitly manipulate the objects in a list, need to be written and verified in FX. Even if the no-affine-indices restriction were to somehow be lifted, representing mutable objects using affine types directly would lead to significant difficulties. For example, if our translation were to translate the mutable object `rev` in FX to an affine record `arev::A`, then we would need an alternative type of list, say `alist::A → A`, in which to store `arev::A` values and a duplicate standard library to operate over `alist \alpha` values (which would be much more cumbersome to write, since we would have to pay attention to the affinity of both the list and of the α -typed values it contains). Even in a system like F° (Mazurak et al. 2010), which supports sub-kinding between normal and linear kinds, we would have to construct separate libraries to work with linear and non-linear lists. Tokens make life much easier!

Nested objects and object invariants. Translating objects with nested objects does not add any further complexity. Object invariants too are easily translated using refinement types. The translation of the FX type `paper` (from Figure 1) is shown below. The constructor and destructor are similar to those for `list` and thus omitted.

```

1 private type paper = {self:{authors:list user; revs: list rev} |
2   ∀u:user, r:rev. In u self.authors && In r self.revs ⇒ NoConflict u r.u }
3 val swapRevs: p:paper → newrl:list rev → (Token p * Token newrl *
4   {_:unit | ∀u:user, r:rev. In u p.authors && In r newrl ⇒ NoConflict u r.u}) →
5   (oldrl:list rev * p':paper * Token oldrl * Token p' *
6   {_:unit | p'.authors=p.authors && p'.revs=newrl && p.revs=oldrl})

```

The refinement on the `paper` record shows the invariant from FX. Swapping a new value `newrl` into the `revs` field of a paper p calls `swapRevs` and has to prove that `newrl` satisfies the object invariant. The tokens of the old record p and the new field value `newrl` are consumed. The token for the old field value `oldrl` is available again.

As an example, we show how to translate an FX expression `p.rev := Cons r tl; None` using `swapRevs`:

```

1 let newrl, tok_newrl = mkCons r tl (tok_r, tok_tl, ()) in
2 let oldrl, p', tok_oldrl, tok_p', _ = swapRevs p newrl (tok_p, tok_newrl, ()) in
3 let none, tok_none, () = mkNone () in
4 none, p', tok_none, tok_p', ()

```

At line 1, we construct the new review list `newrl` by calling `mkCons`. We then call `swapRevs` to swap `newrl` in and the old field value `oldrl` out, creating a new record p' . We create a record `none` corresponding to `None` at line 3. At line 4, in addition to returning `none`, we thread out p' and its token, since it represents an updated object that may be accessed later.

2.4 Understanding FX error messages

Errors that may be reported when verifying FX programs include type mismatch, e.g., passing an integer instead of a string to `up_review`, and assertion failures, i.e., assertions in FX programs are translated to unsatisfiable formulas in Fine. Another kind of errors relates to aliasing and is more subtle. Aliasing restrictions in FX are violated if a program tries to use an object that is unavailable, because the object has been deleted, or nested inside another object, or assigned to a local variable. More experience with FX programming and implementation in the future would give us more insights into user-friendly error reporting.

2.5 DCIL and the efficient execution of FX programs

The purely functional model of FX in terms of Fine, while convenient for verification and equational reasoning, leaves much to be desired in terms of producing an efficient execution strategy for FX programs. While not the main focus of this paper, in this section, we briefly consider how FX programs may be executed efficiently by translating them first to Fine and thence to DCIL—a dependently typed bytecode language to which we compile Fine in a type-preserving style. We begin with a few brief remarks to introduce DCIL.

DCIL is a typing discipline for a purely functional fragment of the .NET byte code language (CIL) (ECMA 2006). It extends the type system of CIL with value-indexed types and affine types, but not refinement types. DCIL is backwards compatible with CIL, and DCIL programs can be executed on standard .NET virtual machines. Conceptually, our type-preserving compilation strategy encodes refinement types in Fine using dependent pairs in DCIL. A value v of a refined type $\{x:t \mid \phi\}$ in Fine is translated to a t -typed value v and a proof of the refinement formula $\phi[v/x]$. A consequence of our type-preservation results for Fine and DCIL is that FX programs also enjoy compilation in a proof-carrying style to .NET byte code. However, after verification of DCIL byte code, we may choose to optimize the program by, say, erasing computationally irrelevant

fragments of the program such as proof terms. In what follows, we discuss a few further optimizations that aim, roughly, to undo the functional encoding of state wrought by the translation from FX to Fine—of course, we have yet to implement any of these optimizations.

Erasure of tokens and units. The first inefficiency in our translation of FX to Fine is the introduction of affine tokens and refined unit values. Tokens are threaded through computation to witness the accuracy of the functional model of state, and the unit values appear as placeholders which witness the validity of pre- and post-conditions—both of these uses are computationally irrelevant. Since both tokens and units are identifiable by their types, we expect a type-based erasure of these fragments of the program to be a simple first optimization.

In-place updates. A convenient aspect of our translation is that the parts of the program that perform field updates are demarcated as separate functions within the `Prims` module, e.g., field updates and swaps are modeled by functions like `reviewUpdTxt` and `swapRevs`. These functions are expensive since they copy their arguments rather than performing in-place updates. While Fine itself has no facilities to model in-place updates, DCIL programs, because they are also CIL programs do not have this restriction. A straightforward strategy, then, is for functions like `reviewUpdTxt` to be implemented in CIL using in-place updates. We can then rely on the .JIT compiler to inline the function call.

Elimination of threading. A final inefficiency of our translation is that mutable objects are always threaded through computations; although, clearly, once mutation is implemented using in-place updates, threading objects is redundant. A pass to eliminate threading would be a nice optimization. However, in contrast to the two previous optimizations, elimination of threading requires a somewhat more sophisticated whole-program transformation.

3. Formalization of core FX

This section formalizes the syntax and the dynamic semantics of a core subset of FX. As mentioned previously, FX is considered a surface syntax for Fine. FX has no static semantics of its own. Instead, we introduce the notion of *type shapes* for FX and formalize a simple syntactic analysis of FX to compute these shapes. Type shapes serve only to facilitate the translation of FX to Fine (§4), where the types of an FX program can be interpreted using Fine’s type system. We show that this interpretation is sound by proving a simulation between FX programs and their Fine counterparts.

Core FX is a lambda calculus augmented with imperative commands for object allocation, deletion, and mutation. It uses a monomorphic typing discipline based around dependent functions, Hoare types and objects with mutable fields. Its syntax is below.

Core syntax of FX

| | |
|--|-------------|
| $\tau ::= TC \mid x:\tau \rightarrow C \mid \{\overline{f:\tau} \mid \phi\}$ | types |
| $C ::= \{(s)\Phi\} r:\tau \{(s')\Psi\}$ | Hoare types |
| $\phi, \Phi, \Psi ::= P \overline{p} \mid \forall x:\tau. \Phi \mid \exists x:\tau. \Phi$ $\mid \Phi \wedge \Phi' \mid \Phi \vee \Phi' \mid \neg \Phi \mid \top \mid \perp$ | formulas |
| $S ::= \text{assume } \phi \mid P :: k \mid S; S'$ | signature |
| $p ::= v \mid s(\eta) \mid p.f$ | atoms |
| $v ::= \eta \mid c \mid \text{error} \mid \lambda x:\tau. e$ | values |
| $\eta ::= x \mid \ell$ | names |
| $e ::= v \mid v v' \mid \text{let } x = e \text{ in } e' \mid e:C \mid \text{assert } (s)\Phi$ $\mid \text{new}_\tau \{\overline{f = v}\} \mid \text{destruct } v \text{ as } \{\overline{f = x}\} \text{ in } e \mid v.f ::= v'$ | terms |

FX kinds are the same as those in Fine (see §2.1) and omitted here. FX types τ include nullary type constants TC , like `unit`. Abstractions are given impure, dependent function types, $x:\tau \rightarrow C$, where x names the formal and is bound in the computation type C , whose syntax was described in §1.2. (Pure function types $x:\tau \rightarrow \tau'$

are sugar for $x:\tau \rightarrow \{(s)\top\} \tau' \{(t)\top\}$, since the pre- and post-conditions \top do not mention any objects.) We also have object types $\{\overline{f:\tau} \mid \phi\}$, which can be thought of as records with mutable fields $f_1:\tau_1, \dots, f_n:\tau_n$. The formula ϕ specifies an object invariant, referring to each of the fields using `self.f1, ..., self.fn`.

Formulas are ranged over by ϕ, Φ , and Ψ , and are first order formulas defined over a set of n -ary base predicates P over atoms p . The atoms include FX values v ; field projections (which are not themselves values in FX); and terms of the form $s(\eta)$, which stands for the value of an object named η in the state s . We use ϕ for object invariants, which cannot refer to state variables, while Φ and Ψ may contain state variables.

A signature S collects top-level user-specified assumptions `assume` ϕ and declarations of predicates $P :: k$ of an FX program. The assumptions are global and do not refer to state variables. The predicate declarations specify kinds (k) of predicates (P).

Values in FX are names (variables x or memory locations ℓ); constants c , where the type of each constant c is denoted TC_c ; `error`, used only during execution to indicate failures; or lambda abstractions. Expressions e are required to be in A-normal form (Flanagan et al. 1993). Dependent typing in FX and in Fine only permits values to index types—A-normal form helps ensure that expressions do not escape into the type level. The expression forms include the standard function application $v v'$ and let-binding `let` $x = e \text{ in } e'$; the form $e:C$ ascribes a computation type C to e ; and `assert` $(s)\Phi$ asserts formula Φ of the current state s .

The imperative fragment of FX includes two core constructs for object allocation and deletion. Allocation is `new` _{τ} $\{\overline{f = v}\}$, which creates a new object of type τ initializing its fields f_1, \dots, f_n to v_1, \dots, v_n . Expression `destruct` $v \text{ as } \{\overline{f = x}\} \text{ in } e$ destructs an object v , and binds the contents of its fields f_1, \dots, f_n to fresh variables x_1, \dots, x_n in e . The fresh variables x_i are in scope within e , while v is no longer accessible in e or later in the program. Using these two constructs alone, we can encode field updates, field swaps, etc. However, we include field swaps, $v.f ::= v'$, as a primitive as well, since we use this in our examples. This instruction assigns the new value v' to the field f of v and returns the old value of the field.

3.1 Dynamic semantics

Runtime configurations in FX are expressions e paired with a store σ , a finite map from memory locations ℓ to object values $\{\overline{f = v}\}$. We show the dynamic semantics of FX below. Congruence rules are elided.

The semantics has the form $(e, \sigma) \rightarrow (e', \sigma')$, and is mostly as one would expect, except for two subtleties. First, to specify progress for FX, it is convenient to ensure that FX programs never get stuck. Instead, we mark certain evaluation steps as erroneous and define safe FX programs as those that are guaranteed to reduce without taking erroneous steps. Specifically, when evaluating `assert` $(s)\Phi$ we check whether the current values in the store satisfy the formula Φ , under the assumptions in the signature S . Subterms $s(\ell)$ in Φ are replaced with \hat{v}_ℓ , which is a location-free object value corresponding to $\sigma(\ell)$ (where $\sigma \vdash \ell \downarrow \hat{v}_\ell$ is in §4). Assertions evaluate to the error value if the check fails (the last rule). Although some care has to be taken to translate the two-state predicate Ψ properly, the translation of ascriptions ($e:C$) to assertions is straightforward.

The other subtlety is that bindings in the store σ are tagged with a superscript t , indicating whether a location represents allocated (\oplus) or freed (\ominus) memory. These tags have no impact on the reduction of FX programs, but the translation of FX to Fine uses these tags to prove that if an FX program is translated to a well-typed Fine program, it never references freed memory. The reduction rule for `new` creates a new memory location ℓ and tags it with \oplus . The rule for `destruct` substitutes the bound variables with the object contents

in e and tags the destructed location with \ominus . Finally, the rule for swap updates the object and returns the old value, as expected.

Expressions with failures such as assertion failures or missing fields evaluate to the error value.

Selected dynamic semantics of FX: $(e, \sigma) \rightarrow (e', \sigma')$

| | |
|--|--|
| $\sigma ::= \sigma + \{\ell \mapsto \{\{f = \bar{v}\}\}^t \mid \bullet \text{ where } t ::= \oplus \mid \ominus$ | |
| $(\mathbf{let } x = v \mathbf{ in } e, \sigma) \rightarrow (e[v/x], \sigma)$ | LET |
| $(\mathbf{let } x = v : C \mathbf{ in } e, \sigma) \rightarrow (e[v/x], \sigma)$ | LETC |
| $((\lambda x : \tau. e) v_2, \sigma) \rightarrow (e[v_2/x], \sigma)$ | FUN |
| $(\mathbf{assert } (s)\Phi, \sigma) \rightarrow ((\cdot), \sigma)$ | if $S \models \Phi[\hat{v}_\ell/s(\ell), \forall \ell \in FV(\Phi)]$ where $\sigma \vdash \ell \downarrow \hat{v}_\ell$ |
| $((e : \{(s)\Phi\} r : \tau \{(t)\Psi\}), \sigma) \rightarrow$ $(\mathbf{let } _ = \mathbf{assert } (s)\Phi \mathbf{ in let } x = e \mathbf{ in}$ $\mathbf{let } _ = \mathbf{assert } (t)\Psi[x/r, \sigma(\ell)/s(\ell)] \mathbf{ in } x, \sigma)$ | |
| $(\mathbf{new } \tau \{\{f = \bar{v}\}\}, \sigma) \rightarrow (\ell, \sigma + \{\ell \mapsto \{\{f = \bar{v}\}\}^\oplus\}) \quad \ell \notin \text{dom}(\sigma)$ | |
| $(\mathbf{destruct } \ell \mathbf{ as } \{\{f = \bar{x}\}\} \mathbf{ in } e, \sigma + \{\ell \mapsto \{\{f = \bar{v}\}\}^t\}) \rightarrow$ $(e[v_i/x_i]_{i=1}^n, \sigma + \{\ell \mapsto \{\{f = \bar{v}\}\}^\ominus\})$ | |
| $(\ell. f_i := v'_i, \sigma + \{\ell \mapsto \{\{f = \bar{v}\}\}^t\}) \rightarrow$ $(v_i, \sigma + \{\ell \mapsto \{\{\dots; f_i = v'_i; \dots\}\}^t\})$ | |
| $(e, \sigma) \rightarrow (\mathbf{error}, \sigma) \quad e \neq v \text{ otherwise}$ | |

3.2 Computing type shapes

In preparation for the translation of FX to Fine, we define a simple syntactic analysis of FX programs to compute *type shapes*. Type shapes describe an equivalence relation on computation types, where types that describe expressions with the same footprint (i.e., the set of locations read or written) are placed in the same equivalence class. The footprint of an expression (and hence the shape of its type) guides the translation of §4 since it determines the values that we must thread into and out of a translated FX expression.

Formally, we define two functions, iv (input variables) and ov (output variables), from computations types C to sets of names $X, Y \subseteq 2^\eta$. Given a computation type $C = \{(s)\Phi\} r : \tau \{(t)\Psi\}$, $\text{iv}(C) = \{\eta \mid s(\eta) \in \text{Subterms}(\Phi, \Psi)\}$ and $\text{ov}(C) = \{\eta \mid t(\eta) \in \text{Subterms}(\Psi)\} \setminus \{r\}$. We write $C \sim \{X\} r : \tau \{Y\}$, when $X = \text{iv}(C)$, $Y = \text{ov}(C)$, and $\tau \sim \tau'$. The last relation, $\tau \sim \tau'$ is a structural congruence on type shapes lifted into function types, i.e., $\tau \sim \tau'$ and $x : \tau_1 \rightarrow C_1 \sim x : \tau_2 \rightarrow C_2$ when $\tau_1 \sim \tau_2$ and $C_1 \sim C_2$, where we write $C_1 \sim C_2$ when $C_1 \sim \{X\} r : \tau \{Y\}$ and $C_2 \sim \{X\} r : \tau \{Y\}$.

Intuitively, the input variables of an expression e (strictly, e 's type C) represent the objects that may be read, modified, deleted, or stored inside other objects by e ; the output variables represent the subset of the input variables (excluding the returned value r) that are still accessible after e has been evaluated.

The first judgment $\Gamma \vdash v \sim \tau$ says that the value v (if it type checks in Fine) has a type that is in the same shape equivalence class as τ , where Γ maps names to their types. Likewise, the judgment $\Gamma \vdash e \sim C$ says that an expression (if it type checks in Fine) has a type that is in the same equivalence class as C . In some cases, we overload notation and write $\Gamma \vdash e \sim \{X\} r : \tau \{Y\}$, with the obvious meaning, i.e., the type of e is in the equivalence class described by $\{X\} r : \tau \{Y\}$. We also write $\eta : \tau, X$ to conditionally add a name to an input or output variable set— $\eta : \tau, X$ means η, X when τ is an object type, and X otherwise.

Computing type shapes: $\Gamma \vdash v \sim \tau$ and $\Gamma \vdash e \sim C$

| | |
|---|---|
| $\Gamma \vdash \eta \sim \Gamma(\eta)$ | $\Gamma \vdash c \sim TC_c$ |
| $\Gamma, x : \tau \vdash e \sim C$ | $\Gamma \vdash v_1 \sim (x : \tau' \rightarrow C) \quad \Gamma \vdash v_2 \sim \tau'$ |
| $\Gamma \vdash \lambda x : \tau. e \sim x : \tau \rightarrow C$ | $\Gamma \vdash v_1 v_2 \sim C[v_2/x]$ |

| | |
|---|--|
| $\frac{\Gamma \vdash v \sim \tau}{\Gamma \vdash v \sim \{v : \tau\} r : \tau \{\cdot\}}$ | $\frac{X = \{\eta \mid s(\eta) \in \text{Subterms}(\Phi)\}}{\Gamma \vdash \mathbf{assert } (s)\Phi \sim \{X\} r : \tau \{\cdot\}}$ |
| $\frac{\Gamma \vdash e \sim C' \quad C' \sim C}{\Gamma \vdash (e : C) \sim C}$ | $\frac{\tau' = \{\{f : \bar{\tau} \mid \phi\}\} \quad \Gamma \vdash \bar{v} \sim \bar{\tau}}{\Gamma \vdash (\mathbf{new } \tau' \{\{f = \bar{v}\}\}) \sim \{\bar{v} : \bar{\tau}\} r : \tau' \{\cdot\}}$ |
| $\frac{\Gamma \vdash v_1 \sim \{\{f : \bar{\tau} \mid \phi\}\} \quad \Gamma \vdash v_2 \sim \tau_k}{\Gamma \vdash (v_1. f_k := v_2) \sim \{v_1, v_2 : \tau_k\} r : \tau_k \{v_1\}}$ | |
| $\frac{\Gamma \vdash v \sim \{\{f : \bar{\tau} \mid \phi\}\} \quad \Gamma, \bar{z} : \bar{\tau} \vdash e \sim \{X\} r : \tau \{Y\}}{\Gamma \vdash \mathbf{destruct } v \mathbf{ as } \{\{f = \bar{z}\}\} \mathbf{ in } e \sim \{X \setminus \{\bar{z}\} \cup v\} r : \tau \{Y \setminus \bar{z}\}}$ | |
| $\frac{\Gamma \vdash e_1 \sim \{X_1\} x : \tau_1 \{Y_1\} \quad \Gamma, x : \tau_1 \vdash e_2 \sim \{X_2\} r : \tau_2 \{Y_2\}}{\Gamma \vdash \mathbf{let } x = e_1 \mathbf{ in } e_2 \sim \{X_1 \cup (X_2 \setminus \{x\})\} r : \tau_2 \{Y_2 \cup (Y_1 \setminus X_2) \setminus \{x\}\}}$ | |

The first four rules compute shapes for values and function application. An object value v , when used in an expression context, can be given a computation type that includes itself as an input variable, and no output variables. We give assert expressions $\mathbf{assert } (s)\Phi$ a shape where the input and output variables are the same, namely those referenced by Φ , i.e., all the objects mentioned in an assertion are still accessible after the assertion has been evaluated. An expression e can be ascribed the type C , (using $e : C$) so long as the ascription does not change the shape of the type.

When computing type shapes for the imperative operations, we pay attention to the aliasing relationships induced by each operation. However, we do not attempt to check here that FX programs respect the aliasing constraints on the object graph necessary for assertion safety—all checking is left to Fine. With this principle in mind, consider the rule for **new**. It states that the input variables are the stateful field values \bar{v} , while the output variable set is empty. Since references to each of the \bar{v} are captured by the newly allocated object, the empty output variable set indicates that the \bar{v} should no longer be accessed directly by the program. The swap statement is similar. Its input variables include the object being updated v_1 and the new contents of the field v_2 (if its type τ_k is an object type); the output variables only include v_1 since the reference to v_2 has been captured by v_1 . For a **destruct** expression, the input variables include v , the object being destructed, and e 's input variables but not the \bar{z} since these are locally bound; for the same reason, \bar{z} are excluded from the output variables. Note that e (or some enclosing expression) may attempt to use v after it has been destructed—this situation will be trapped by the Fine type checker. Finally, we have **lets**. Here, the input variables are the union of the input variables of e_1 and e_2 , excluding the let-bound variable x . The output variables are the outputs of e_2 (excluding x) and those outputs of e_1 that are not in the footprint of e_2 .

4. Translating FX to Fine

As sketched in §2.3, FX programs are translated to functional Fine programs. We formalize this translation here, and prove that the translation of FX programs preserve their semantics. Since Fine programs can be verified using refinement type checking, we also obtain the result that the safety of FX programs can be established by the Fine type checker. Our description of the translation starts by describing the role of tokens and their types in our translation. We then describe the translation of types and finally the translation of terms. The section concludes with discussion of our main results.

4.1 FX objects, tokens, and ownership

For safety in the presence of mutation and aliasing, we require FX programs to respect an ownership discipline on the object store. Affine tokens play a central role in our translation, and serve primarily to enforce this ownership discipline. To illustrate this point, we first discuss how to represent FX object values in Fine. As

we have seen, FX object values are locations ℓ , where a store σ maintains a mapping from locations to records holding the object's fields. In Fine, we represent object values as immutable records. The translation of objects is described by the judgment below.

$$\frac{\sigma \vdash \ell \downarrow \{f = v\}^\ell \quad \sigma \vdash v_i \downarrow \hat{v}_i \quad \forall i}{\sigma \vdash \ell \downarrow \{f_1 : \hat{v}_1, \dots, f_n : \hat{v}_n\}^\ell}$$

This rule looks up the binding for a location ℓ in σ , and produces a tuple of Fine values, where each \hat{v}_i , recursively, is the translation of the components of the object stored at ℓ .¹ Fine values that correspond to FX objects are tagged with the location ℓ from which the value was derived—we explain its significance shortly.

Now, given a store $\sigma = \ell_1 \mapsto \{f = 0\}^\ell$, consider the simple FX program and its translation to Fine shown below (eliding irrelevant unit-values from the Fine versions for clarity).

$$\begin{array}{ll} \text{FX (1)} & \text{let } x = \text{new}_\tau \{g = \ell_1\} \text{ in let } y = \ell_1.f :=: 1 \text{ in } x \\ \text{Fine (1)} & \text{let } x, \text{tok}_x = \text{mk}_\tau \{f = 0\}^{\ell_1} \text{ tok}_{\ell_1} \text{ in} \\ & \text{let } y, \text{tok}_y = \text{swap}_\tau \{f = 0\}^{\ell_1} 1 \text{ tok}_{\ell_1} \text{ in } x, \text{tok}_x \end{array}$$

The FX program allocates a new object x storing the object ℓ_1 inside it, then updates ℓ_1 and returns x . The program above violates ownership, since after a single step of execution, the store contains a new location $\ell_2 \mapsto \{g = \ell_1\}^\ell$, meaning that a reference to ℓ_1 has been captured by ℓ_2 , and the very next step mutates ℓ_1 directly.

The image of the FX program under the translation is the Fine program shown directly beneath it. The Fine program contains calls to the functions mk_τ and swap_τ (discussed in §2.3), passing in Fine values corresponding to the object ℓ_1 (i.e., $\{f = 0\}$), and, importantly, the variable tok_{ℓ_1} representing a token or capability to use the immutable record value in a stateful way. Tokens are, of course, affine and since tok_{ℓ_1} is used more than once, the Fine program fails to type check, and the FX program is dismissed as potentially unsafe.

Not all violations of ownership are as easily detected as in this first example. One problematic case is when a reference to an object is captured by a closure, and somewhere else in the program the captured object is updated via an alias, as in the example shown below (with the same initial store σ).

$$\begin{array}{ll} \text{FX (2)} & \text{let } g = \lambda _.\text{unit}.\ell_1 \text{ in let } _ = \ell_1.f :=: 1 \text{ in } g \\ \text{Fine (2)} & \text{let } g = \lambda _.\text{unit}.\lambda x:\text{Token} \{f = 0\}^{\ell_1} .(\{f = 0\}^{\ell_1}, x) \text{ in} \\ & \text{let } _ = \text{swap}_\tau \{f = 0\}^{\ell_1} 1 \text{ tok}_{\ell_1} \text{ in } g \end{array}$$

In this case, we translate the thunk g by adding an additional argument for the token of the captured object value $\{f = 0\}^{\ell_1}$. Notice that the type of the token variable ($\text{Token} \{f = 0\}^{\ell_1}$) is indexed by the value for which it is a capability. Any operation that mutates ℓ_1 consumes its token, producing a new token with a type indexed by the new value. Thus, after an update the thunk g becomes unusable since a token with the appropriate type cannot be passed in. (Unless the update does not change the value of the object, in which case it is still perfectly safe to use g). Since this program never attempts to use g , it is safe in FX and is also type correct in Fine. In other words, in addition to tokens being affine, giving them value-indexed types is also crucial for safety.

Token threading, for all its benefits (recall also the discussion of §2.3), also makes it harder to prove that our translation from FX to Fine is a simulation. Informally, we wish to show that when an FX program e_1 is translated to a well-typed Fine program \hat{e}_1 , and if e_1 steps to e_2 , then \hat{e}_1 steps (using one or more steps) to some \hat{e}_2 , where \hat{e}_2 is semantically equivalent to the image of e_2 under the translation. If we inspect the example program FX (2) above, we see that after several steps it reduces

¹Throughout this section, syntactic elements from Fine are indicated using the hatted version of the corresponding elements in FX.

to $\lambda _.\text{unit}.\ell_1$, in a store where $\ell_1 \mapsto \{f = 1\}^\ell$, which translates to $\lambda _.\text{unit}.\lambda x:\text{Token} \{f = 1\}^{\ell_1} .(\{f = 1\}^{\ell_1}, x)$. However, this Fine lambda term is syntactically distinct from the value of g in the program Fine (2) above— g contains a stale value for ℓ_1 . In proving our simulation result we show that, despite the syntactic differences, these two values are in fact semantically equivalent. We use Fine's module system (formalized and proved sound previously using the colored brackets of Grossman et al. (2000)) to show that stale values that result from the reduction of Fine programs are semantically irrelevant since these values must always be treated abstractly. The ℓ -superscripts on values facilitate this proof.

4.2 Type translation

The translation of FX types τ and computation types C to Fine types $\hat{\tau}$ (\hat{C} is a synonym for $\hat{\tau}$) is shown below. Throughout the translation, we use variables named tok_η for η 's token. We also use $\text{toks}(X)$ to mean the set of tokens for variables in X , i.e., $\text{toks}(\eta.X) = \text{tok}_\eta, \text{toks}(X)$ with $\text{toks}(\cdot) = \cdot$.

Translation of types: $\Gamma \vdash \tau \rightsquigarrow \hat{\tau}$ and $\Gamma \vdash C \rightsquigarrow \hat{C}$

$$\frac{}{\Gamma \vdash TC \rightsquigarrow TC} \quad \frac{\forall i. \Gamma \vdash \tau_i \rightsquigarrow \hat{\tau}_i \quad \Gamma \vdash \phi \rightsquigarrow \hat{\phi}}{\Gamma \vdash \{f:\tau \mid \phi\} \rightsquigarrow \{\text{self}:\{f:\hat{\tau}} \mid \hat{\phi}\}}$$

$$\frac{\Gamma \vdash \tau \rightsquigarrow \hat{\tau} \quad \Gamma; x:\tau \vdash C \rightsquigarrow \hat{C} \quad C \rightsquigarrow \{X\}r:\tau\{Y\}}{\Gamma \vdash x:\tau \rightarrow C \rightsquigarrow x:\hat{\tau} \rightarrow (\{\text{Token } x_j\}_{x_j \in X} * \{ _:\text{unit} \mid \text{Pre}(C)_\Gamma \}) \rightarrow \hat{C}}$$

$$\frac{C \rightsquigarrow \{X\}r:\tau\{Y\} \quad \Gamma \vdash \tau \rightsquigarrow \hat{\tau} \quad \forall y_j \in Y, \Gamma \vdash y_j:\tau_j \quad \Gamma \vdash \tau_j \rightsquigarrow \hat{\tau}_j}{\Gamma \vdash C \rightsquigarrow (r:\hat{\tau} * \{y'_j:\hat{\tau}_j\}_{y_j \in Y} * \{\text{Token } y'_j\}_{y_j \in Y} * \{ _:\text{unit} \mid \text{Post}(C)_\Gamma \})}$$

$$\text{where, for } C = \{(s)\Phi\}r:\tau\{(t)\Psi\}$$

$$\text{Pre}(C)_\Gamma = \hat{\Phi} \quad \text{if } \Gamma \vdash \Phi[\eta/s(\eta)] \rightsquigarrow \hat{\Phi}$$

$$\text{Post}(C)_\Gamma = \hat{\Psi} \quad \text{if } \Gamma \vdash \Psi[\eta/s(\eta)][\eta'/t(\eta)] \rightsquigarrow \hat{\Psi}$$

Type constructors TC remain the same. Object types in FX are translated to refined record types in Fine, with each field type translated to $\hat{\tau}_j$, and the invariant translated to $\hat{\phi}$ as described shortly. Function types in FX $x:\tau \rightarrow C$ are translated to Fine function types that bind the same parameter x (with translated type $\hat{\tau}$), tokens for each input variable of C , and a unit value refined with the translated precondition $\text{Pre}(C)_\Gamma$. The function returns a value with the translated computation type \hat{C} . A computation type C is translated to a dependent tuple, consisting of the return value r (with the translated return type $\hat{\tau}$), the output variables y'_j and their tokens, and a unit value whose type is refined with the translated post-condition $\text{Post}(C)_\Gamma$. The output variables y_j are rebound to y'_j in the tuple type, to stand for the values of y_j at the function exit point.

Translation of FX formulas is straightforward and is defined by the judgments $\Gamma \vdash \Phi \rightsquigarrow \hat{\Phi}$ and $\Gamma \vdash \phi \rightsquigarrow \hat{\phi}$. The rules in these judgments are congruences over the structure of formulas, where, in the base case, value-indexed predicates $P\bar{v}$ are translated to $P\hat{\bar{v}}$, for $\Gamma \vdash v_i \rightsquigarrow \hat{v}_i$, the value translation judgment of the next section. We use two macros $\text{Pre}(C)_\Gamma$ and $\text{Post}(C)_\Gamma$ for translating pre- and post-conditions of computations to refinement formulas in Fine. The refinement formula replaces the stateful terms in a formula, $s(\eta)$ and $t(\eta)$, with the corresponding name bindings for the immutable values in Fine that correspond to η .

Translation of signatures. Predicate declarations $P :: k$ are translated to type constructors in Fine with the same name P and the same kind k . Assumptions $\text{assume } \phi$ are translated to data constructors in Fine representing proofs of $\hat{\phi}$ if $\Gamma \vdash \phi \rightsquigarrow \hat{\phi}$. These data constructors are included in Fine's LCF-style proof kernel, along with the standard first-order logic inference rules (see §2.1).

4.3 Value and expression translation

This section describes the translation of FX expressions. An FX expression e is translated to a Fine terms \hat{e} , where the free variables of \hat{e} correspond to the object locations ℓ in e . For translation of source code the translated term \hat{e} is closed, since store locations do not occur in the source code of a program.

The value translation judgment $\Gamma \vdash v \rightsquigarrow \hat{v}$ states that an FX value v is translated to a Fine value \hat{v} . The translation of FX names η (variables or locations) is trivial—these are translated to Fine variables with the same names, which may be substituted away when the Fine expression is closed. The rule for translation of functions is shown below.

$$\frac{\Gamma, x:\tau_1; \bullet \vdash (e:C) \rightsquigarrow \hat{e} \quad \Gamma \vdash \tau_1 \rightsquigarrow \hat{\tau}_1 \quad \hat{\tau}_2 = (\{\text{Token } \eta\}_{\eta \in \text{iv}(C)} * \{_:\text{unit} \mid \text{Pre}(C)\}_{\Gamma})}{\Gamma \vdash \lambda x:\tau_1.(e:C) \rightsquigarrow \lambda x:\hat{\tau}_1.\lambda y:\hat{\tau}_2.\text{let } \overline{\text{tok}}_{\eta, -} = y \text{ in } \hat{e}}$$

This rule follows the translation of function types. We add a parameter $y:\hat{\tau}_2$ to receive the tokens for the function’s input variables. The function body $e:C$ is translated to \hat{e} , preceded by a **let** that unpacks the components of the tuple y , giving token variables their required names and types, $\text{tok}_{\eta}:\text{Token } \eta$. Our formalization does not attempt to infer types, or pre- and post-conditions for functions. As such, the rule above expects both the formal parameter x to be decorated with a type and for the function body to be ascribed a computation type. In practice, we expect weakest precondition inference to alleviate some of this difficulty, although, of course, annotations will be expected at least on recursive functions. Additionally, while our simulation result holds for translations of arbitrary FX programs (that are well-typed in Fine), in practice we expect FX programs to be closure-converted prior to translation.

The judgment $\Gamma; \dot{Y} \vdash e \rightsquigarrow \hat{e}$ (below) translates an FX expression e to a Fine expression \hat{e} . The environment includes \dot{Y} , either empty (\bullet) or a set of names Y representing the output variables of a computation. We define an operation \odot on output sets: $\dot{Y} \odot Y'$ returns Y' if $\dot{Y} = \bullet$, and \dot{Y} otherwise. Intuitively, \dot{Y} corresponds to locations that may have been modified in expressions preceding e in a block of **lets**. Fine values (and tokens) corresponding to locations in this set need to be “threaded out” in the translation.

$$\begin{array}{c} \text{Translation of expressions, } \Gamma; \dot{Y} \vdash e \rightsquigarrow \hat{e} \\ \hline \text{T-Return} \frac{\Gamma \vdash v \rightsquigarrow \tau \quad \Gamma \vdash v \rightsquigarrow \hat{v}}{\Gamma; \dot{Y} \vdash v \rightsquigarrow (\hat{v}, Y, \text{toks}(v:\tau, Y), ())} \\ \text{T-Let} \frac{\Gamma \vdash e_1 \rightsquigarrow \{X_1\}x:\tau_1 \{Y_1\} \quad \Gamma; \bullet \vdash e_1 \rightsquigarrow \hat{e}_1 \quad \Gamma \vdash \text{let } x = e_1 \text{ in } e_2 \rightsquigarrow C \quad \Gamma, x:\tau_1; \dot{Y} \odot \text{ov}(C) \vdash e_2 \rightsquigarrow \hat{e}_2}{\Gamma; \dot{Y} \vdash \text{let } x = e_1 \text{ in } e_2 \rightsquigarrow \text{let } x, Y_1, \text{tok}_x, \text{tok}_{Y_1, -} = \hat{e}_1 \text{ in } \hat{e}_2} \\ \text{T-Assert} \frac{\Gamma \vdash \Phi[\eta/s(\eta)] \rightsquigarrow \hat{\Phi} \quad X = \text{iv}_s(\Phi)}{\Gamma; \bullet \vdash \text{assert } (s)\Phi \rightsquigarrow \text{ASRT}_X(\hat{\Phi})} \\ \text{T-Chk} \frac{\Gamma; \bullet \vdash e \rightsquigarrow \hat{e} \quad C \sim \{X\}r:\tau \{Y\} \quad C = \{(s)\Phi\}r:\tau \{(t)\Psi\} \quad X_1 = \text{iv}_s(\Phi) \quad X_2 = \text{iv}_t(\Psi)}{\Gamma; \bullet \vdash (e:C) \rightsquigarrow \text{let } _ , X_1, \text{tok}_{X_1, -} = \text{ASRT}_{X_1}(\text{Pre}(C)_{\Gamma}) \text{ in} \\ \text{let } r, Y, \text{toks}(r:\tau, Y), _ = \hat{e} \text{ in} \\ \text{let } _ , X_2, \text{tok}_{X_2, -} = \text{ASRT}_{X_2}(\text{Post}(C)_{\Gamma}) \text{ in} \\ (r, Y, \text{toks}(r:\tau, Y), ())} \\ \text{T-App} \frac{\Gamma \vdash v_1 v_2 \rightsquigarrow C \quad \Gamma \vdash v_1 \rightsquigarrow \hat{v}_1 \quad \Gamma \vdash v_2 \rightsquigarrow \hat{v}_2}{\Gamma; \bullet \vdash v_1 v_2 \rightsquigarrow \hat{v}_1 \hat{v}_2 (\text{toks}(\text{iv}(C)), ())} \\ \text{T-New} \frac{\Gamma \vdash (\text{new}_{\tau} \{f = v\}) \rightsquigarrow C \quad \forall i. \Gamma \vdash v_i \rightsquigarrow \hat{v}_i}{\Gamma; \bullet \vdash (\text{new}_{\tau} \{f = v\}) \rightsquigarrow \text{mk}_{\tau} \hat{v}_1 \cdots \hat{v}_n (\text{toks}(\text{iv}(C)), ())} \end{array}$$

$$\begin{array}{c} \text{T-Upd} \frac{\Gamma \vdash (v_1.f_k := v_2) \rightsquigarrow C \quad \Gamma \vdash v_1 \rightsquigarrow \hat{v}_1 \quad \Gamma \vdash v_2 \rightsquigarrow \hat{v}_2}{\Gamma; \bullet \vdash (v_1.f_k := v_2) \rightsquigarrow \text{upd}_{f_k}(\hat{v}_1, \hat{v}_2, \text{toks}(\text{iv}(C)), ())} \\ \text{T-Destr} \frac{\Gamma \vdash v \rightsquigarrow \{f:\tau\} \quad \Gamma \vdash v \rightsquigarrow \hat{v} \quad \Gamma, z_i:\tau_i; \dot{Y} \vdash e \rightsquigarrow \hat{e}}{\Gamma; \dot{Y} \vdash \text{destruct } v \text{ as } \{f = z\} \text{ in } e \rightsquigarrow \\ \text{let } \overline{z}_i, \overline{\text{tok}}_{z_i, -} = \text{destr}_{\tau} \hat{v} \text{ tok}_v \text{ in } \hat{e}} \end{array}$$

To illustrate the significance of \dot{Y} , consider the FX expression $\text{let } y = (\ell.f := 1) \text{ in } y$. The type shape computed for this expression is $\{\ell\}r:\text{int } \{\ell\}$, indicating that ℓ was modified in the expression and is still accessible after the expression has been evaluated. When translating the body of the **let** (i.e., y), we need to return y as well as the value and token corresponding to the updated value of ℓ . The rule (T-Return) does just this. It translates values v that appear in expression contexts to tuples containing the translated value \hat{v} , any additional values dictated by the output set Y (if any), tokens for the object values, and finally a unit for the post-condition.

The other aspect of the \dot{Y} environment is illustrated in the translation of let expressions $\text{let } x = e_1 \text{ in } e_2$. The first premise computes the shape of the let-bound expression e_1 for its set of output variables Y_1 , and binds names and tokens for these in the translation. To translate the body e_2 , we compute the shape C of the entire let expression for its set of outputs $Y = \text{ov}(C)$ —these are locations that may have been modified in either e_1 or e_2 and are still live, and so, must be threaded out when translating e_2 . We do this by translating e_2 in a context $\dot{Y} \odot Y$, which ensures that Y is threaded out of e_2 , unless \dot{Y} is non-empty, in which case the output variables of some enclosing let are threaded out of e_2 .

The remaining rules are mostly straightforward. To translate assertions, we define the following macros: $\text{iv}_s(\Phi) = \{\eta \mid s(\eta) \in \text{Subterms}(\Phi)\}$ and $\text{ASRT}_X(\hat{\Phi}) = ((, X, \text{toks}(X), ((: \{ _ : \text{unit} \mid \hat{\Phi} \})))$. Assertions reduce to unit in FX—the first unit value in the tuple corresponds to this. The last unit value with a refined type corresponds to a check of the assertion formula. Additionally, we include values and tokens for each object value $\ell \in X$. This ensures that assertions in Fine never refer to stale values, thereby keeping the behavior of Fine assertions in correspondence with FX. The ascription form $(e:C)$ is translated, as expected, using a pair of assertions. The tokens used in the assertion are rebound so that the capabilities they stand for are not consumed.

Applications $(v_1 v_2)$ are translated to pass in tokens (computed using the type shape judgment) for the input variables of v_1 , as well as a unit value for the precondition (T-App). The Fine type checker must prove that the unit value passed in can be given a refined type corresponding to the precondition of v_1 . The translation of **new** uses the constructor mk_{τ} of the record type τ (T-New). (The function mkRev in §2.3 is an instance of such a constructor.) The constructor takes the initial values $\hat{v}_1, \dots, \hat{v}_n$ for the fields, consumes the tokens for the input variables (again computed using type shapes), and returns the new record r and its token tok_r . The translation of swaps using (T-Upd) is similar. The translation of a destruct expression calls the destructor destr_{τ} of the record type τ (T-Destr). The destructor consumes the token tok_v for the record to destruct, and returns the field values and their tokens. These are let-bound and in scope for \hat{e} , the translated body.

The main result of this section is Theorem 1, which shows that the translation from FX to Fine preserves the semantics of the FX program, if the resulting Fine program is well-typed. More precisely, the translation is a weak simulation modulo strong bisimilarity. We do not give meaning to FX programs that translate to ill-typed Fine programs since ill-typed Fine programs are never executed.

From this result we derive Corollary 2, a progress result for FX that shows that FX programs that translate to well-typed Fine programs do not reach the **error** state. An additional result is that the computation types given to FX programs describe small footprints. §A gives formalizations of all lemmas and proofs.

Setting up for the simulation. Using the rules of §4.3, an FX expression e produces \hat{e} , a Fine program that contains free variables corresponding to locations ℓ and tokens for locations tok_ℓ . Given a store σ that maps locations ℓ to values v^ℓ , free location variables ℓ are eliminated via substitution with the corresponding translated Fine values \hat{v}^ℓ .

The semantics of Fine also makes it convenient to handle free token variables, since it permits affine values to be held in a memory, M , rather than inlined in the program. We translate the FX store σ to a Fine memory M , which maps free token variables tok_ℓ in \hat{e} to Fine values $\text{MkToken } \hat{v}^\ell$. The reduction of Fine programs has the form $(M, \hat{e}) \xrightarrow{\text{fine}} (M', \hat{e}')$. Reads from M are destructive, and Fine’s metatheory establishes that affine values are used at most once.

We present a simplified version of our simulation lemma below, and illustrate its structure graphically. To type check the Fine programs \hat{e}_1 and \hat{e}_2 in the statement below, we use an environment including a signature \hat{S}_{base} for FX primitives (corresponding to the **Prims** module of §2.3), and affine capabilities for every token in M .

| FX | | Fine |
|-------------------|--------------------|--|
| (e_1, σ_1) | \rightsquigarrow | (M_1, \hat{e}_1) |
| \downarrow | | \downarrow_+ |
| (e_2, σ_2) | \rightsquigarrow | $(M_2, \hat{e}_2) \cong (M_2, \hat{e}'_2)$ |

Theorem 1 (Fine programs simulate FX programs). *If an FX runtime configuration (e_1, σ_1) translates to a well-typed Fine configuration (M_1, \hat{e}_1) , and (e_1, σ_1) steps to (e_2, σ_2) , then (e_2, σ_2) translates to a well-typed Fine configuration (M_2, \hat{e}'_2) , and (M_1, \hat{e}_1) takes multiple steps to a configuration (M_2, \hat{e}_2) such that \hat{e}'_2 is strongly bisimilar to \hat{e}_2 .*

Theorem 1 establishes that our translation relation is a simulation up to strong bisimilarity, between \hat{e}_2 and \hat{e}'_2 above. The relation $M_2 \vdash \hat{e}_2 \cong \hat{e}'_2$ states that \hat{e}_2 and \hat{e}'_2 reduce in “lock step”, and is derived from a value abstraction theorem provided by Fine’s module system. In addition to this core simulation result, the full version of Theorem 1 establishes several other key properties. Among these are 1. that FX programs that translate to well-typed Fine programs respect aliasing and ownership constraints on the object graph; and 2. that the types given to effectful expressions only mention locations that they access or modify.

Theorem 1 also yields an important corollary. Since the **error** state cannot be translated to Fine, FX programs that translate to well-typed Fine programs do not reach the **error** state.

Corollary 2 (Progress for FX). *Given a closed FX program e . If $\cdot; \bullet \vdash e \rightsquigarrow \hat{e}$, where $\hat{S}_{base}; \cdot; \vdash \hat{e} : \hat{\tau}$, then for any reduction sequence $(e, \cdot) \xrightarrow{\text{fx}}^* (e', \sigma')$, e' is guaranteed to not be **error**.*

5. Examples

This section illustrates the use of FX using further examples. We have applied the translation of §4 (manually, at present) to extended versions of these examples (and the ConfWeb example of §2), and verified the resulting programs using the Fine type checker.

5.1 Verifying stateful APIs in the presence of aliasing

In this section, we show how to verify two properties of clients of a stateful API of collections and iterators, even in the presence of aliasing. First, we ensure that the collection underlying an iterator is never modified while an iteration over the collection is in progress. Second, we ensure that a client never attempts to extract

```

1 (* Fragment of a library implementing fractional permissions *)
2 private type trk  $\alpha = \{v:\alpha; p:\text{rat}\}$ 
3 type Aliases ::  $(\alpha :: * \Rightarrow \alpha \Rightarrow \alpha \Rightarrow *)$ 
4 assume A_refl:  $\forall x:\text{trk } \alpha, y:\text{trk } \alpha. x.v=y.v \Leftrightarrow \text{Aliases } x \ y$ 
5 val split:  $x:\text{trk } \alpha \rightarrow \{(s) \top\} y:\text{trk } \alpha$ 
6    $\{(t) \text{Aliases } t(x) \ s(x) \ \&\& \ \text{Aliases } t(x) \ t(y) \ \&\& \ t(y).p=(x).p+s(x).p/2 \}$ 
7 val join:  $x:\text{trk } \alpha \rightarrow y:\text{trk } \alpha \rightarrow \{(s) \text{Aliases } s(x) \ s(y)\} \text{unit}$ 
8    $\{(t) \text{Aliases } t(x) \ s(x) \ \&\& \ t(x).p=s(x).p+s(y).p \}$ 
9 (* Fragment of a safe wrapper for the Collections API *)
10 type collection  $\alpha$ 
11 type iterator  $\alpha$ 
12 type coll  $\alpha = \text{trk } (\text{collection } \alpha)$ 
13 type istate = HasMore | Unknown
14 type iter  $\alpha = \text{trk } \{i:\text{iterator } \alpha; c:\text{coll } \alpha; \text{st}:\text{istate}\}$ 
15 val newColl:  $\text{unit} \rightarrow \{(s) \top\} c:\text{coll } \alpha \{t(t).p=I \}$ 
16 val add:  $c:\text{coll } \alpha \rightarrow y:\text{trk } \alpha \rightarrow \{(s) \ s(c).p=I \ \&\& \ s(y).p > 0 \} \text{unit}$ 
17    $\{(t) \ t(c).p=s(c).p \ \&\& \ \text{Aliases } t(c) \ s(c) \ \&\& \ t(y).p=s(y).p/2 \}$ 
18 val iterator:  $c:\text{coll } \alpha \rightarrow \{(s) \ s(c).p > 0 \} i:\text{iter } \alpha$ 
19    $\{(t) \ t(c).p=s(c).p/2 \ \&\& \ \text{Aliases } t(c) \ s(c) \ \&\& \ t(i).p=I \ \&\& \ \text{Aliases } t(c) \ t(i).v.c \ \&\& \ t(i).v.c.p = t(c).p \}$ 
20
21 val finalize:  $c:\text{coll } \alpha \rightarrow i:\text{iter } \alpha \rightarrow \{(s) \ s(i).p=I \ \&\& \ \text{Aliases } c \ s(i).v.c \} \text{unit}$ 
22    $\{(t) \ t(i).p=0 \ \&\& \ \text{Aliases } t(c) \ s(c) \ \&\& \ t(c).p = s(c).p + s(i).v.c.p \}$ 
23 val next:  $i:\text{iter } \alpha \rightarrow \{(s) \ s(i).p=I \ \&\& \ s(i).v.\text{st}=\text{HasMore}\} y:\text{trk } \alpha$ 
24    $\{(t) \ t(i).p=s(i).p \ \&\& \ t(y).p > 0 \ \&\& \ \text{Aliases } s(i) \ t(i) \ \&\& \ t(i).v.i=s(i).v.i \}$ 
25 val hasNext:  $i:\text{iter } \alpha \rightarrow \{(s) \ s(i).p > 0 \} b:\text{bool}$ 
26    $\{(t) \ t(i).p=s(i).p \ \&\& \ \text{Aliases } t(i) \ s(i) \ \&\& \ (b=\text{true} \Leftrightarrow t(i).v.\text{st}=\text{HasMore}) \}$ 
27 (* A client of the Collections API *)
28 val client:  $c:\text{coll } \alpha \rightarrow \{(s) \ s(c).p=I \} \text{unit} \{t(t).p = s(c).p \ \&\& \ t(c).v=s(c).v \}$ 
29 let client c =
30   let it1 = iterator c in
31   let rec loop1 (c:coll  $\alpha$ ) (it1:iter) :
32      $\{(s) \ \text{Aliases } s(it1).v.c \ s(c) \ \&\& \ s(it1).p=I \ \&\& \ s(c).p > 0 \} \text{unit}$ 
33      $\{(t) \ t(it1).p=s(it1).p \ \&\& \ t(it1).v.c=s(it1).v.c \ \&\& \ t(c).p=s(c).p \} =$ 
34     if hasNext it1 && ... then
35       let rec loop2 (c:coll  $\alpha$ ) (it2:iter) :
36          $\{(s) \ \text{Aliases } s(it2).v.c \ s(c) \ \&\& \ s(it2).p=I \} \text{unit}$ 
37          $\{(t) \ t(c).p=s(c).p+s(it2).v.c.p \} =$ 
38         if hasNext it2 then let a = next it2 in ... loop2 c it2
39       else finalize c it2 in
40       let it2 = iterator c in loop2 c it2;
41       let b = next it1 in ... loop1 c it1
42     else () in
43   loop1 c it1; finalize c it1

```

Figure 2. Controlled aliasing using fractional permissions

elements from an iterator that has been exhausted. Our example is adapted from the work of Bierhoff and Aldrich (2007), who develop a special-purpose type system that uses linear logic to check type-state properties in the presence of aliasing, and apply it to clients of the Java collections library. A variant of this example has also been studied by Krishnaswami et al. (2009), who verify a similar program using interactive proofs in a higher-order separation logic. In contrast, we enforce an aliasing discipline by developing a library of fractional permissions (Boyland 2003), and rely on an SMT solver for assertion checking.

We highlight three elements of our solution. First, the library-based approach to aliasing illustrates the flexibility of our approach—variations on our permission scheme could be implemented in a similar fashion. This library also illustrates the value of substructural state—affine tokens introduced by the translation ensure proper use of the library. Next, we leverage local state, a feature of FX. Permissions are mutable fields within structured, alias-controlled objects. A similar approach using monadic state would involve explicitly modeling a map from alias-controlled objects to their permissions. Finally, we show how the local state of permissions for collections and iterators can be easily combined, and for these permissions to be used with a state machine that

tracks whether or not an iterator has been exhausted. The modularity afforded by our approach makes these combinations natural.

The program in Figure 2 contains three parts. First, we have the interface of our permissions library. Next, we show the interface of a wrapper to an underlying implementation of the .NET collections library. Although not shown, the wrapper is programmed in FX and contains calls to the permissions library to manage aliases. Finally, we show a client program that uses the wrapped collections library. Although we have yet to formalize this, we conjecture that for clients that verify against the wrapper, calls to the wrapper can be replaced with direct calls to the underlying library.

Permissions library. Lines 1-9 of Figure 2 show the interface of our fractional permissions library. At a high level, our library defines a type `trk α` that associates a permission `p:rat` (a rational number) with an α -typed value. Using `split`, a client can construct an alias `y` to a tracked object `x`, but the permission previously associated with `x` is split between `x` and `y`. The `join` function allows a client to destroy `y`, an alias of `x`, coalescing the permission previously associated with `y` into `x`'s new permission. The library implementation (not shown) directly manipulates the private permission field, and uses `A.refl` to generate its post-condition `Aliases`. Two tracked values are `Aliases` if they refer to the same underlying object.

Wrapper of the collections library. Lines 10-27 show the interface of our wrapper of the .NET collection API. The abstract types `collection α` and `iterator α` correspond to the types implemented by the .NET library. We then define `coll α`, the tracked version of a `collection α`. The type `iter α` is the tracked version of an iterator and associates a permission with a record of three fields, and illustrates how local state in FX can be easily composed. The first field, `i`, is the iterator itself; the second field `c:coll α` holds a tracked reference to the collection from which the iterator was derived; the last field represents a state machine to track when an iterator is exhausted.

At line 16 we show the signature of `newColl`, a function that allocates a new collection. Its post-condition shows that the returned collection has a full permission. The `add` function (17-18) allows an object `y` to be added to the collection `c`. Since this mutates the underlying collection, the pre-condition of `add` requires `c` to hold a full permission. Since `c` captures a reference to `y`, the post-condition shows the permission of `y` halved.

Lines 19-20 shows show the signature of `iterator`, which produces an iterator `i` over a collection `c`. Calling `iterator` requires only a non-zero permission on `c`, but the post-condition shows that it consumes half of `c`'s permission, since the iterator captures a reference to the collection from which it was derived. By halving the permission of `c`, we ensure that `c` cannot be modified while the iterator is still live. The implementation of `iterator` (not shown) calls `split` to stash a reference to `c` in the returned object. The `finalize` operation (22-23) allows an iterator `i` to be destroyed, so long as there are no other aliases to `i`. Its post-condition restores the permission to the collection from which the iterator was derived—the implementation of `finalize` calls to `join`.

The function `next` (24-25) allows a client to extract a tracked object from an iterator. The pre-condition of `next` requires the iterator to be in a state where it has more elements. In the post-condition, the state of the iterator is updated to indicate that it may or may not have more elements. In order to establish the pre-condition of `next`, clients can call and test the result of `hasNext`.

A client program. Lines 29-44 show a client function that takes a collection `c` with a full permission, extracts an iterator `it1` (line 31), and iterates over `it1` using `loop1` and finalizes the iterator afterward (line 44), leaving `c` with a full permission. At each iteration of `loop1`, we extract another iterator `it2` from `c` (line 41) and iterate over it in `loop2` and finalize `it2` before exiting `loop2` (line 40). Concurrent modifications to the collection `c` (say, by calling `add`)

```

1 module InfoFlow
2 type level = High:level | Low:level | J:level → level → level
3 type LEQ :: level ⇒ level ⇒ *
4 assume Lattice_assumptions: LEQ Low High, ...
5 (* The program counter sensitivity level, initially Low *)
6 private type pc = {v:level}
7 let pc=new pc{ v=Low}
8 (* An abstract monad for leveled data *)
9 private type L :: * ⇒ level ⇒ * = MkL : α → l:level → L α l
10 val return: l:level → α → L α l
11 val bind: l:level → m: {m:level | LEQ l m} → pc1:level → x:L α l →
12     f:(α → {(s2) s2(pc).v=J l pc1 } L β m {(s2') s2'(pc)=s2(pc)}) →
13     {(s1) LEQ s1(pc).v pc1 } L β m {(s1') s1'(pc)=s1(pc)}
14 let bind l m pc1 (MkL x _) f = let tmp = pc.v in pc.v := J l pc1;
15     let r = f x in pc.v := tmp; r
16 (* Channels on which to send data (side effects) *)
17 type Ch :: * ⇒ level ⇒ *
18 val write: α → l:level → Ch α l → {(s1) LEQ s1(pc).v l } unit
19     {(s2) s2(pc)=s1(pc)}
20 end
21 val client: L str Low → L str High → Ch str High →
22     {(s1) s1(pc).v=Low } L str High {(s2) s2(pc)=s1(pc)}
23 let client lx1 lx2 chan = bind Low High Low lx1 (fun x1 →
24     bind High High High lx2 (fun x2 →
25     let x = strcat x1 x2 in (write x High chan); return High x))

```

Figure 3. Tracking information flow through side-effects

fail to type check, since while there are still iterators extant, `c` has less than a full permission.

Notice that aside from the annotation with loop invariants, the `client` function is fairly direct. Pleasingly, `client` contains no explicit calls to `split` or `join`. All these operations are factored into the implementation of the wrapped collection API, and the specifications of this API in effect manage the implicit aliasing behavior of the client program. Clients can also call `split` and `join` directly to explicitly manage aliases, should the need arise.

Translating to Fine. Translating Figure 2 to Fine is relatively straightforward. We first closure-convert the nested loops, hoisting them to the top-level. Thereafter, the translation follows the rules of §4 directly (which generalizes readily to handle type abstraction and application). Our implementation uses a more complex model for permissions that maintains separate fractions for read and write permissions; a larger API for collections, including functions to remove elements from collections and to query the size of a collection; and finally, a more complex client program. The resulting Fine program is 244 lines long, and its verification involves proving 29 goals, which we discharge in 15 seconds.

5.2 Tracking implicit information flows through impure code

Our second example develops an information-flow tracking library suitable for programs with side-effects. The basic idea is to combine a monadic treatment of dependency (as in DCC (Abadi et al. 1999)) with the program-counter technique used (originally) by Fenton (1973) in his Data Mark Machine. In short, we represent α -typed data protected at security level `l` as values of the type `L α l`, where the type `L` represents a family of monads arranged in a lattice according to the ordering on security levels. Additionally, we maintain a global stateful value `pc` that accounts for the influence of protected data on reaching the current program point. Effectful operations (such as writing a message to a channel) have preconditions (expressed as constraints on `pc`) to ensure that they are not control dependent on secret data.

We prove two properties of client programs. First, that `Low` channels never carry data that is marked `High`-security; and, secondly, that `Low` channels are never used at `High` security-level program points. Using these two properties, a syntactic proof of a

noninterference property is possible. Although we have not proved noninterference for the program of Figure 3 per se, a similar encoding was proved noninterferent by Swamy (2008).

Note that this program does not make use of local state—there is a single global piece of state called `pc`. As such, this program could easily be modeled and verified using a monadic treatment of state. However, this example highlights two features of FX. First, it illustrates the ability of FX to handle programs that manipulate global state. Second, it serves as initial evidence that our approach generalizes to higher order programs. Our information flow library example makes use of a second-order function `bind`, where the type of its function argument is parametrized by the preceding arguments to the `bind` function. We discuss further generalizations to higher order programs in §5.3.

The API. Figure 3 shows a module `InfoFlow` (lines 1-20) which defines the monad `L α l` and exports it abstractly to the client (21-25). `InfoFlow` begins with a definition of the security levels using the type `level`, which stands for a two-point lattice, ordered by the relation `LEQ`. The ordering is axiomatized by user-provided assumptions—we show one such assumption at line 4; the others are standard. At line 6, we define the type `pc`, a mutable object with a single field, `v`, which will hold the sensitivity of the program counter—we initialize `pc` to `Low` at line 7. We define the `L α l` type at line 9 and the type of function `return` at line 10, which allows any value to be injected into the monad using any level.

To appreciate the type of `bind` in Figure 3, it is instructive to consider a type of `bind` for flow controls in purely functional code: `l:level → m:{m:level | LEQ l m} → x:L α l → f:(α → L β m) → L β m`. This version of `bind` allows `f` to view the protected data `x:L α l` at the unprotected type `α`, but the constraint `LEQ l m` ensures that the result `f` computes from `x` is at least as secret as `x` itself. However, this model is only sound if `f` is a pure function—nothing prevents `f` from writing its unprotected `α`-typed argument to a public channel.

The purpose of the program-counter `pc` is to prevent such leaks, and it plays a role in two parts of `InfoFlow`. Consider, first, our model of channels at lines 16-19. Here, we define `Ch α l`, the type of a channel on which to send (or receive) `α`-typed data to parties privileged to view data at secrecy level `l`. To protect against leaks via implicit flows, the pre-condition of `write` states that the sensitivity of the program counter `pc` must not be greater than `l`. Next, consider the type of `bind` at line 11—we give the function `f` passed to `bind` the type `α → {(s2) s2(pc).v=J l pc l} L β m {(s2') s2'(pc)=s2(pc)}`. This type allows `f` to view the protected data `x:L α l` at its underlying type `α`. But, the pre-condition on `f`'s type says that the state of `pc` is elevated to be the join of `pc1` (the value of `pc` at the time `write` is called) and `l` (the level on the secret data `x`). This ensures that `f` cannot satisfy the pre-condition of `write` if `f` calls `write` using a channel `c:Ch α l2`, for some `l2` less secret than `l`.

A proof of noninterference for this scheme relies crucially on `f` treating `pc` values abstractly—we achieve this by marking the `pc` type `private`. This ensures, for example, that `f` cannot mutate the `pc` itself. The only mutation of `pc` occurs in the implementation of `bind`, which elevates the `pc` before calling `f`; then restores it before returning the value computed by `f`.

The `client` program concatenates `Low` and `High` strings, writes the result on a `High` channel, and returns a `High` string. The explicit `level` arguments to `bind` (e.g., the three occurrences of `High` on line 24) lead to some syntactic noise—this could easily be eliminated with some inference for implicit parameters.

Translating to Fine. As previously, translating Figure 3 to Fine begins with a simple closure-conversion step, turning the global variable `pc` into an argument of every stateful function. The resulting Fine program is 81 lines long, and produces 7 verification goals that are discharged automatically in 6 seconds.

```

1 type pt = {x:real; y:real}
2 val div_x_by_y: p:pt → {(s) s(p).y i=0} unit {(t) t(p).x = s(p).x/s(p).y}
3 let div_x_by_y p = p.x := p.x/p.y
4 val iter < α:* , PreF::State ⇒ α ⇒ *, PostF::State ⇒ State ⇒ α ⇒ *,
5     Inv::State ⇒ list α ⇒ *, Post::State ⇒ State ⇒ list α ⇒ * > :
6     f:(x:α → {(s) PreF s x} unit {(t) PostF s t x}) → l:list α →
7     {(s) Inv s l && ∀s1:State, s2:State, hd:α, tl:list α.
8         ((Inv s1 Nil ⇒ Post s1 s2 Nil) &&
9          ((Inv s1 (Cons hd tl) ⇒ PreF s1 s2 hd) &&
10         (Inv s1 (Cons hd tl) && PostF s1 s2 hd ⇒ Inv s2 tl) &&
11         (PostF s1 s2 hd && Post s1 s2 tl ⇒ Post s1 s2 (Cons hd tl)))) } unit
12     {(t) Post s t l}
13 let rec iter f l = match l with Nil → () | Cons hd tl → f hd; iter f tl
14 prop XDiv :: State ⇒ State ⇒ list pt ⇒ *
15 assume XD_nil: ∀s1:State, s2:State, XDiv s1 s2 Nil
16 assume XD_cons: ∀s1:State, s2:State, hd:pt, tl:list pt.
17     ((s2(hd).x = s1(hd).x/s1(hd).y) && (XDiv s1 s2 tl) ⇒ XDiv s1 s2 (Cons hd tl))
18 prop YNEQ0 :: State ⇒ list pt ⇒ *
19 assume YNEQ0_Nil: ∀ s:(YNEQ0 s Nil)
20 assume YNEQ0_Cons: ∀s:State, hd:pt, tl:list pt.
21     (s(hd).y i=0) && (YNEQ0 s tl) ⇒ YNEQ0 s (Cons hd tl)
22 val div_all: pl:list pt → {(s) YNEQ0 s pl} unit {(t) XDiv s t pl}
23 let div_all pl = iter < α=pt, PreF=λs.λp. s(p).y i=0,
24     PostF=λs1.λs2.λp. s2(p).x=s1(p).x/s1(p).y,
25     Inv=YNEQ0, Post=XDiv > div_x_by_y pl

```

Figure 4. Mutating a list of objects using a higher-order iterator

5.3 Extension to higher-order programs

Although our token-threading translation works equally well for both first- and higher-order programs, we have yet to formalize an extension of the assertion language of FX to work with higher-order stateful code. However, this extension is fairly natural, since the kind language of Fine makes it easy to write types for higher-order programs. A simple example of this was shown in §2.2, where the type of `for_all` abstracted over the predicate `P` decided by its argument `f`. However, in the case of `for_all`, the argument `f` was required to be a pure function. Below is a detailed example showing how extend this scheme to higher-order stateful code. We intend to elaborate on the application of FX to higher-order programs in future work.

Our final example illustrates how FX generalizes to the verification of higher-order stateful programs. We use extensively polymorphism which is provided by the implementation of Fine. This example uses a technique due to Régis-Gianas and Pottier (2008) for purely functional programs. A recent extension of this work to higher-order stateful programs has been developed in the context of the Who verifier (Kanig and Filiâtre 2009), requiring a custom effect analysis and assuming a strict no-aliasing discipline. Our example shows that Fine's general-purpose affine and refinement types are up to the task too, and, although we have yet to combine them, we conjecture that the fractional permissions library of §5.1 can be used in conjunction with higher-order stateful computations to relax Who's restrictions on aliasing.

At a high-level, the program in Figure 4 works with points, objects with two mutable fields, `x` and `y`, called `pt`. Given a list of points `pl`, the client program `div_all` calls a library function `iter` on `pl`, passing in a stateful function `div_x_by_y` which updates the `x` field of each point in the list by dividing it by the `y` field. Given that `pl` is a list in which every point has a non-zero `y` field, we would like to prove that after updating `pl`, the `x` field of each point in the updated list has been divided by the corresponding `y` field, and that no divide-by-zero errors occur.

The type and definition of iter. The main complexity in this example is in the type of `iter`, shown at lines 4-12. The traditional arguments of `iter` are a type argument `α:*` for the type of the list elements; a function argument `f` of type `α → unit`; and a list `l:list α`.

The implementation of `iter` (line 13) applies `f` to each element of `l` and returns `unit`. To verify `iter` independently of its uses, we give it a specification that generalizes over the specification of `f`—the additional type arguments to `iter` (`PreF` etc.) serve this purpose.

First, the predicates `PreF` and `PostF` correspond to the pre- and post-conditions of the function `f`. We use an FX-specific type `State` for the state variables bound in computation types. (The `State` type has no counterpart in Fine—they are erased in the translation.) `PreF` has kind $State \Rightarrow \alpha \Rightarrow *$, meaning a one-state predicate applied to the α -typed values, i.e., for a state variable `s` and $x:\alpha$, `PreF s x` is a valid type. `PostF` has kind $State \Rightarrow State \Rightarrow \alpha \Rightarrow *$, a two-state predicate on α -typed values. We use these predicates in `f`'s type: $(x:\alpha \rightarrow \{s\} \text{PreF } s \ x) \ \text{unit} \ \{(t) \text{PostF } s \ t \ x\}$. `PreF` allows us to speak about the value of `x` in the pre-state `s`; `PostF` relates the value of `x` in the pre-state `s` to its value in the post-state `t`. In addition to the pre- and post-conditions of `f`, `iter` is also parametrized by an inductive invariant `Inv` on the list `l`; as well as a post-condition `Post` for `iter` itself. `Inv`, like `PreF`, is a single-state predicate for the state of the list just prior to each (recursive) call to `iter`. The post-condition `Post` relates the state of the list before and after each call to `iter`.

Then, at lines 7-12, we see the computation type for the body of `iter`. The pre-condition expresses a number of constraints between the four quantified predicates. First, we require the inductive invariant to be true of the list argument `l` in the initial state (`Inv s l` on line 7); line 8 is the base case of the induction; line 9 establishes the pre-condition of `f` on each iteration; line 10 re-establishes the inductive invariant; line 11 is the induction step; the post-condition on line 12 is straightforward.

A client program. At lines 14-17, we define an inductive, two-state predicate `XDiv` on lists of points to say, roughly, that each point in the list is updated so that its `x` field is divided by its `y` field. At lines 18-21, we define another inductive (one-state) predicate, `YNEQ0`, to say that every point in a list has a non-zero `y` field. At line 22 we show the type of `div_all`. In the implementation of `div_all`, we call `iter` instantiating each of its type and predicate parameters. The instantiation of α to `pt` is easy. But, to instantiate the predicate parameters, we rely on predicate literals of the form $\backslash x.P$, where `x` is an abstracted value variable in the body of the proposition `P`. Using this notation, we instantiate `PreF` and `PostF` to predicates corresponding to the pre- and post-conditions of `div_x.by.y`; `Inv` and `Post` to the invariant and the expected post-condition of `div_all`. Finally, we pass in the function argument and the list itself.

Translating to Fine. Admittedly, our example program uses some rather powerful machinery to describe a type for `iter`—machinery that is considerably beyond the formalism of prior sections. However, an extension of our translation to account for these features is relatively straightforward. Indeed, a significant virtue of our approach of verifying stateful FX programs via translation to functional Fine programs is that much of the complexity associated with quantifying over state predicates is simply eliminated (although, we still rely on the ability in Fine to quantify over predicates on values). Our translation produces an 89 line program in Fine, with 13 proof obligations that we discharge in about 7 seconds.

As stated, our formalization of translation in Section 4 is sufficiently expressive to handle higher-order stateful functions abstractions—threading tokens through such computations “just works”. However, without the ability to quantify over predicates, higher-order stateful functions are not particularly useful. The main limitation of our formalism is that it lacks polymorphism and an associated language of kinds. Extending our formalism to handle polymorphism over $*$ -kinded types is trivial; quantification over non-stateful predicates, as in the type of `for_all` from §2 also poses no additional challenges. However, translating quantification over stateful predicates does take some care—we discuss this below.

Each n -ary one-state predicate of kind $State \Rightarrow t_1 \Rightarrow \dots \Rightarrow t_n \Rightarrow *$, is translated to a predicate in Fine by erasing the `State` argument, i.e., $t_1 \Rightarrow \dots \Rightarrow t_n \Rightarrow *$. No translation of the types $t_1 \dots t_n$ is needed either, since object types in FX are translated to records of the same name and fields in Fine. Applications of one-state predicates are translated trivially—we simply drop the state variable argument and the other arguments remain unchanged. For each n -ary two state predicate, $State \Rightarrow State \Rightarrow t_1 \Rightarrow \dots \Rightarrow t_n \Rightarrow *$ in FX, we generate a $2n$ -ary predicate in Fine of kind $t_1 \Rightarrow t_1 \Rightarrow \dots \Rightarrow t_n \Rightarrow t_n \Rightarrow *$. To translate applications of two-state predicates we follow a pattern similar to the translation of post-conditions in computation types, i.e., the $Post(\Psi, C)$ thing from Section 4. We used primed variables etc. For example, the instantiation of `PostF` at line 27 in Figure 4 becomes $\backslash p.\backslash p'.p'.x=p.x/p.y$; the return type of `iter` ($r:\text{unit} * l:\text{list } \alpha * \text{Valid } l' * \{ _:\text{unit} | \text{Post } l' \}$); etc. The simplicity of this translation is a direct consequence of our functional simulation—it allows us to simply erase state variable arguments and to relate the value of a variable in the pre-state to its value in the post-state using a pair of immutable values.

6. Related work and conclusions

Our work is perhaps most closely related to the work of Charguéraud and Pottier (2008), who show how to functionalize imperative programs using a translation based on linear capabilities. Like us, they prove their translation sound via a simulation argument. However, their work does not include a program logic, although motivated by the desire to verify programs; our work includes the use of Hoare types and their translation to refinement types. Additionally, in the absence of a logic, Charguéraud and Pottier embed a specific aliasing discipline into their calculus, (based on the adoption and focus constructs of Fähndrich and DeLine (2002)). In contrast, because of the expressiveness of refinement types, we show how aliasing controls can be encoded using a library of fractional permissions.

Our use of tokens in the translation of FX to Fine, is closely related to Walker et al.'s calculus of capabilities (Walker et al. 2000). Whereas Walker et al. prove a syntactic type soundness property for the capability calculus (which is sufficient for their domain of safe memory management), we prove the correctness of our token-based translation using a simulation argument. Instead of targeting just memory safety properties, the ATS language (Zhu and Xi 2005) combines stateful views with indexed types for full verification, where stateful views are described using linear logic. This gives ATS the ability to reason directly about pointer manipulations, but at a price, since linear logic is hard to automate. FX is also related to HOOP (Flanagan et al. 2006), a language that uses dependent types to express refinements on imperative objects. However, unlike FX, refinements in HOOP can only mention immutable data.

Finally, our work is closely related to the work of Borgström et al. (2010), who use Hoare types for a state monad to verify stateful computations. However, as argued throughout this paper, the state monad suffers from a lack of modularity—a monolithic state is threaded through the entire program, making it difficult (without resorting to separation logic, as in Ynot), to reason locally about parts of the state or to mix stateful idioms. In contrast, FX programs use local state in the form of mutable objects, and, as illustrated by the example of §5.1, permits multiple stateful idioms to be combined in a modular way.

In summary, we have presented FX, a functional language with support for mutable objects. We show how FX programs can be translated to functional Fine programs, using affine types to model state and refinement type checking for verification. Our approach extends the scope of automatic, modular verification for programs that use mutable local state and aliasing.

References

- M. Abadi, A. Banerjee, N. Heintze, and J. G. Riecke. A core calculus of dependency. In *POPL*, 1999.
- S. Arun-Kumar and M. Hennessy. An efficiency preorder for processes. *Acta Inf.*, 29(9):737–760, 1992. ISSN 0001-5903.
- J. Bengtson, K. Bhargavan, C. Fournet, A. D. Gordon, and S. Maffei. Refinement types for secure implementations. In *CSF*, 2008.
- K. Bierhoff and J. Aldrich. Modular typestate checking of aliased objects. *OOPSLA*, 2007.
- J. Borgström, A. Gordon, and R. Pucella. Roles, stacks, histories: A triple for Hoare. *J. Funct. Program.*, 2010. To appear.
- J. Boyland. Checking interference with fractional permissions. In *SAS*, pages 55–72. Springer, 2003.
- I. Cervesato and F. Pfenning. A linear logical framework. *Inf. Comput.*, 179(1), 2002.
- A. Charguéraud and F. Pottier. Functional translation of a calculus of capabilities. In *ICFP '08*, 2008.
- J. Chen, R. Chugh, and N. Swamy. Type-preserving compilation of end-to-end verification of security enforcement. In *PLDI '10*. ACM, 2010.
- A. Chlipala, G. Malecha, G. Morrisett, A. Shinnar, and R. Wisnesky. Effective interactive proofs for higher-order imperative programs. In *ICFP*, 2009.
- L. de Moura and N. Björner. Z3: An efficient SMT solver. In *TACAS*, 2008.
- ECMA. Standard ecma-335: Common language infrastructure (cli), 2006.
- M. Fähndrich and R. DeLine. Adoption and focus: practical linear types for imperative programming. In *PLDI*, 2002.
- J. Fenton. *Information Protection Systems*. PhD thesis, U. Cambridge, 1973.
- C. Flanagan. Hybrid type checking. In *POPL*, 2006.
- C. Flanagan, A. Sabry, B. F. Duba, and M. Felleisen. The essence of compiling with continuations. In *PLDI*, 1993.
- C. Flanagan, S. N. Freund, and A. Tomb. Hybrid types, invariants, and refinements for imperative objects. In *FOOLWOOD '06*, 2006.
- D. Grossman, G. Morrisett, and S. Zdancewic. Syntactic type abstraction. *ACM TOPLAS*, 22(6):1037–1080, 2000. ISSN 0164-0925.
- D. Jackson. Alloy: a lightweight object modelling notation. *TOSEM*, 11(2), 2002.
- J. Kanig and J.-C. Filliâtre. Who: a verifier for effectful higher-order programs. In *ML*. ACM, 2009.
- S. Krishnamurthi. The Continue server. In *PADL*, 2003.
- N. R. Krishnaswami, J. Aldrich, L. Birkedal, K. Svendsen, and A. Buisse. Design patterns in separation logic. In *TLDI*, 2009.
- G. Malecha, G. Morrisett, A. Shinnar, and R. Wisnesky. Toward a verified relational database management system. In *POPL*, 2010.
- K. Mazurak, J. Zhao, and S. Zdancewic. Lightweight linear types in System F°. In *TLDI*, 2010.
- R. Milner. LCF: A way of doing proofs with a machine. In *MFCS*, 1979.
- A. Nanevski, G. Morrisett, and L. Birkedal. Polymorphism and separation in Hoare type theory. In *ICFP*, 2006.
- Y. Régis-Gianas and F. Pottier. A Hoare logic for call-by-value functional programs. In *MPC*, 2008.
- M. Sozeau. Subset coercions in Coq. In *TYPES*. Springer-Verlag, 2006.
- N. Swamy. *Language-based Enforcement of User-defined Security Policies*. PhD thesis, University of Maryland, College Park, August 2008.
- N. Swamy, J. Chen, and R. Chugh. Enforcing stateful authorization and information flow policies in Fine. In *ESOP*, 2010.
- W. Swierstra. A Hoare logic for the state monad. In *TPHOLs*, 2009.
- P. Wadler. Linear types can change the world! In *PROGRAMMING CONCEPTS AND METHODS*. North, 1990.
- D. Walker, K. Crary, and G. Morrisett. Typed memory management via static capabilities. *ACM TOPLAS*, 22(4), 2000.
- D. Zhu and H. Xi. Safe programming with pointers through stateful views. In *PADL*, 2005.

Syntax of instrumented Fine expressions

| | | | |
|-----------|-------|---|-------------|
| \hat{e} | $::=$ | $\hat{v} \mid \hat{v}_1 \hat{v}_2 \mid \mathbf{let} \bar{x} = \hat{e}_1 \mathbf{in} \hat{e}_2$ | Expressions |
| \hat{u} | $::=$ | $x \mid \ell \mid tok_\ell \mid D\hat{v}_1, \dots, \hat{v}_n \mid \lambda x: \hat{\tau}. \hat{e} \mid (\hat{v}_1, \hat{v}_2)$ | pre-values |
| \hat{v} | $::=$ | $\hat{u} \mid \hat{u}^{\ell: \tau}$ | values |

Extending typing relation in Fine for l -annotated values

$$\frac{\hat{S}; \hat{\Gamma}; \cdot \vdash \hat{u} : \tau \quad \hat{S} \ni (\text{prims}, \tau)}{\hat{S}; \hat{\Gamma}; X \vdash \hat{u}^{\ell: \tau} : \tau} \text{ T-Loc} \quad \text{Note the empty } X\text{-set in the premise of this rule. Values corresponding to objects are token-free.}$$

Reduction for Fine expressions

| | | | |
|---|--------|--|---------|
| $(M; \lambda x. \hat{e} \hat{v}) \xrightarrow[\text{pub}]{\text{fine}} (M; \hat{e}[\hat{v}/x])$ | R-Beta | $(M; \mathbf{let} x = \hat{v} \mathbf{in} \hat{e}) \xrightarrow[\text{pub}]{\text{fine}} (M; \hat{e}[\hat{v}/x])$ | R-Let |
| $(M; \hat{e}_1) \xrightarrow[\text{p}]{\text{fine}} (M'; \hat{e}'_1)$ | R-Ctx | $(M; \mathbf{let} x = \hat{e}_1 \mathbf{in} \hat{e}_2) \xrightarrow[\text{p}]{\text{fine}} (M'; \mathbf{let} x = \hat{e}'_1 \mathbf{in} \hat{e}_2)$ | |
| $(M; \mathbf{let} x, y = (\hat{v}_1, \hat{v}_2) \mathbf{in} \hat{e}) \xrightarrow[\text{pub}]{\text{fine}} (M; \hat{e}[\hat{v}_1/x][\hat{v}_2/y])$ | R-Pair | | |
| $M = M_0, \text{TokenVal } (\hat{v}_1), \dots, \text{TokenVal } (\hat{v}_n) \quad \hat{v}_{\text{base}} = (\hat{v}_1, \dots, \hat{v}_n)^{\ell: \tau}$ | | $M' = M_0, (tok_\ell, \text{MkToken } \hat{v}_{\text{base}})$ | R-New |
| $(M; \text{mk}_\tau \hat{v}_1 \dots \hat{v}_n (\text{tokens}(\hat{v}_1, \dots, \hat{v}_n), ())) \xrightarrow[\text{prims}]{\text{fine}} (M'; \hat{v}_{\text{base}}, tok_\ell, ())$ | | | |
| $\hat{v}_{\text{base}} = (\hat{v}_1, \dots, \hat{v}_k, \dots, \hat{v}_n)^{\ell: \tau} \quad M = M_0, \text{TokenVal } \hat{v}_{\text{base}}, \text{TokenVal } \hat{v}'_k$ $\hat{v}'_{\text{base}} = (\hat{v}_1, \dots, \hat{v}'_k, \dots, \hat{v}_n)^{\ell: \tau} \quad M' = M_0, \text{TokenVal } \hat{v}'_{\text{base}}, \text{TokenVal } \hat{v}_k$ | | $(M; \text{upd}_{f_k}(\hat{v}_{\text{base}}, \hat{v}'_k, \text{tokens}(\hat{v}_{\text{base}}, \hat{v}'_k), ())) \xrightarrow[\text{prims}]{\text{fine}} (M'; \hat{v}_k, \hat{v}'_{\text{base}}, \text{tokens}(\hat{v}_k, \hat{v}'_{\text{base}}), ())$ | R-Upd |
| $\hat{v}_{\text{base}} = (\hat{v}_1, \dots, \hat{v}_n)^{\ell: \tau} \quad M = M_0, \text{TokenVal } \hat{v}_{\text{base}}$ $M' = M_0, \text{TokenVal } \hat{v}_1, \dots, \text{TokenVal } \hat{v}_n$ | | $(M; \text{destr}_\tau \hat{v}_{\text{base}} tok_{\hat{v}_{\text{base}}}) \xrightarrow[\text{prims}]{\text{fine}} (M'; \hat{v}_1, \dots, \hat{v}_n, \text{tokens}(\hat{v}_1, \dots, \hat{v}_n), ())$ | R-Destr |

where

1. $\text{TokenVal } \hat{v} = (tok_l, \text{MkToken } \hat{v})$ when $\hat{v} = (\hat{v}_1)^{\ell: \tau}$
otherwise, $\text{TokenVal } \hat{v} = \cdot$
2. $\text{token}(\hat{v}) = tok_l$ when $\hat{v} = (\hat{v}_1)^{\ell: \tau}$
otherwise $\text{token}(\hat{v}) = \cdot$
3. $\text{tokens}(\hat{v}_1, \hat{v}) = \text{token}(\hat{v}_1), \text{tokens}(\hat{v})$
with $\text{tokens}() = \cdot$

Figure 5. Syntax and semantics of Fine (Reproduced with minor revisions from ESOP '10)

A. Soundness of the translation

The translation of runtime configurations of FX programs entails one additional technicality that amount to further book-keeping for tokens. The subtlety that needs to be handled is that when a program $e = \mathbf{let} x = x.f := 1; e_1 \mathbf{in} e_2$ reduces to $e' = \mathbf{let} x = e_1 \mathbf{in} e'_2$, the type shape computed for $(x.f := 1; e_1)$ is, in general, different from the shape computed for e_1 alone; e.g., the shape for e_1 alone would omit x as an output variable. Since the shape of the types govern the structure of the token threading, we need to account for the tokens threaded when translating e when translating the reduction of e , i.e., e' .

We accomplish this in the translation judgment (Figure 6) $\Gamma; \mathbf{Y} \vdash e \Longrightarrow C \rightsquigarrow \hat{e}$. This judgment is a reproduction of the rules of §4, with two changes. First, the judgment is indexed by a structure \mathbf{Y} that specifies the set of output variables to be threaded out of each sub-expression in the program. Second, we find it convenient in our proofs to inline the definition of the type-shape judgment into the definition of translation judgment—thus, in addition to translating e to \hat{e} , we compute a type shape C for e while accounting explicitly for the tokens that, according to \mathbf{Y} , must be threaded through.

This additional judgment is simply a technical device. We prove (Lemma 13) that whenever $\Gamma; \dot{Y} \vdash e \rightsquigarrow \hat{e}$ then $\exists \mathbf{Y}, C. \Gamma; \mathbf{Y} \vdash e \Longrightarrow C \rightsquigarrow \hat{e}$.

Definition 3 (Base signature). *A type declaration in FX is of the form $td = (\tau = \{f_1 : \tau_1; \dots; f_n : \tau_n\})$ and an FX signature is a sequence of type declarations, $S ::= td, S | \cdot$.*

$\boxed{\Gamma; \mathbf{Y} \vdash e \Longrightarrow C \rightsquigarrow \hat{e}}$ Translation of FX runtime configurations, where $\mathbf{Y} ::= Y \mid (1; \mathbf{Y}, 2; \mathbf{Y})$

$$\begin{array}{c}
\text{T-Return} \frac{\Gamma \vdash v \sim \tau \quad \Gamma \vdash v \rightsquigarrow \hat{v} \quad C \sim \{v: \tau\} r: \tau \{Y\}}{\Gamma; \mathbf{Y} \vdash v \Longrightarrow C \rightsquigarrow (\hat{v}, Y, \text{tok}_s(v: \tau, Y), ())} \quad \text{T-Let} \frac{\Gamma; \mathbf{Y}.1 \vdash e_1 \Longrightarrow C_1 \rightsquigarrow \hat{e}_1 \quad C_1 \cong \{X_1\} x: \tau_1 \{Y_1\} \quad \Gamma, x: \tau_1; \mathbf{Y}.2 \vdash e_2 \Longrightarrow C_2 \rightsquigarrow \hat{e}_2 \quad C_2 \cong \{X_2\} s: \tau_2 \{Y_2\} \quad C \cong \{X_1 \cup X_2 \setminus \{x\}\} s: \tau_2 \{Y_2\}}{\Gamma; \mathbf{Y} \vdash \text{let } x = e_1 \text{ in } e_2 \Longrightarrow C \rightsquigarrow \text{let } x, Y_1, \text{tok}_x, \text{tok}_{Y_1}, - = \hat{e}_1 \text{ in } \hat{e}_2} \\
\text{T-Assert} \frac{\Gamma \vdash \Phi[\eta/s(\eta)] \rightsquigarrow \hat{\Phi} \quad X = \text{iv}_s(\Phi) \quad C \sim \{X\} \text{.unit} \{X\}}{\Gamma; \bullet \vdash \text{assert } (s)\hat{\Phi} \Longrightarrow C \rightsquigarrow \text{ASRT}_X(\hat{\Phi})} \quad \text{T-App} \frac{\Gamma \vdash v_1 v_2 \sim C \quad \Gamma \vdash v_1 \rightsquigarrow \hat{v}_1 \quad \Gamma \vdash v_2 \rightsquigarrow \hat{v}_2}{\Gamma; \bullet \vdash v_1 v_2 \Longrightarrow C \rightsquigarrow \hat{v}_1 \hat{v}_2 (\text{tok}_s(\text{ov}(C)), ())} \\
\text{T-Chk} \frac{\Gamma; \mathbf{Y} \vdash e \Longrightarrow C' \rightsquigarrow \hat{e} \quad C \sim \{X\} r: \tau \{Y\} \sim C' \quad C = \{(s)\Phi\} r: \tau \{(t)\Psi\} \quad X_1 = \text{iv}_s(\Phi) \quad X_2 = \text{iv}_t(\Psi)}{\Gamma; \mathbf{Y} \vdash (e; C) \Longrightarrow C \rightsquigarrow \text{let } _ , X_1, \text{tok}_{X_1}, - = \text{ASRT}_{X_1}(\text{Pre}(C)_\Gamma) \text{ in let } r, Y, \text{tok}_s(r: \tau, Y), - = \hat{e} \text{ in let } _ , X_2, \text{tok}_{X_2}, - = \text{ASRT}_{X_2}(\text{Post}(C)_\Gamma) \text{ in } (r, Y, \text{tok}_s(r: \tau, Y), ())} \\
\text{T-New} \frac{\Gamma \vdash (\text{new}_\tau \{\overline{f=v}\}) \sim C \quad \forall i. \Gamma \vdash v_i \rightsquigarrow \hat{v}_i}{\Gamma; \bullet \vdash (\text{new}_\tau \{\overline{f=v}\}) \Longrightarrow C \rightsquigarrow \text{mk}_\tau \hat{v}_1 \cdots \hat{v}_n (\text{tok}_s(\text{iv}(C)), ())} \quad \text{T-Upd} \frac{\Gamma \vdash (v_1.f_k := v_2) \sim C \quad \Gamma \vdash v_1 \rightsquigarrow \hat{v}_1 \quad \Gamma \vdash v_2 \rightsquigarrow \hat{v}_2}{\Gamma; \bullet \vdash (v_1.f_k := v_2) \Longrightarrow C \rightsquigarrow \text{upd}_{f_k}(\hat{v}_1, \hat{v}_2, \text{tok}_s(\text{iv}(C)), ())} \\
\text{T-Destr} \frac{\Gamma \vdash v \sim \{\overline{f: \tau}\} \quad \Gamma \vdash v \rightsquigarrow \hat{v} \quad \Gamma, z_i: \tau_i; \mathbf{Y} \vdash e \Longrightarrow C' \rightsquigarrow \hat{e} \quad C' \sim \{X\} r: \tau \{Y\} \quad C \sim \{v, X\} r: \tau \{Y\}}{\Gamma; \mathbf{Y} \vdash \text{destruct } v \text{ as } \{\overline{f=z}\} \text{ in } e \Longrightarrow C \rightsquigarrow \text{let } \overline{z_i}, \overline{\text{tok}_{z_i}}, - = \text{destr}_\tau \hat{v} \text{ tok}_v \text{ in } \hat{e}}
\end{array}$$

$\boxed{\sigma \vdash v \downarrow \hat{v}; L}$ Translating FX objects to Fine values

$$\frac{\sigma(\ell) = \{\overline{f=v}\}^\dagger \quad \sigma \vdash v_i \downarrow \hat{v}_i; L \quad \forall i}{\sigma \vdash \ell \downarrow (\hat{v}_1, \dots, \hat{v}_n)^{\ell; \tau; L, \ell}} \quad \frac{\cdot \vdash \lambda x: \tau. e \Rightarrow \tau' \rightsquigarrow \hat{v}}{\sigma \vdash \lambda x: \tau. e \downarrow \hat{v}; \cdot} \quad \frac{}{\sigma \vdash c \downarrow c; \cdot}$$

$\boxed{\sigma \vdash \sigma \downarrow M}$ Translating FX stores to Fine token stores

$$\frac{\forall \ell'. \sigma(\ell') = \{\overline{f=v}\}^\oplus \Rightarrow \ell \notin FV(\bar{v}) \quad \sigma_0 \vdash \ell \downarrow \hat{v}; L \quad \sigma_0 \vdash (\sigma \setminus L) \downarrow M}{\sigma_0 \vdash (\sigma, \{\ell \mapsto \{\overline{f=v}\}^\oplus\}) \downarrow M, (\text{tok}_\ell, \text{MkToken } \hat{v})} \text{ M-1} \\
\frac{}{\sigma_0 \vdash (\sigma, \{\ell \mapsto \{\overline{f=v}\}^\ominus\}) \downarrow M} \text{ M-2} \quad \frac{\sigma_0 \vdash \sigma_2, \sigma_1 \downarrow M}{\sigma_0 \vdash \sigma_1, \sigma_2 \downarrow M} \text{ M-3} \quad \frac{}{\sigma_0 \vdash \cdot \downarrow \cdot} \text{ M-4}$$

$\boxed{\sigma \vdash \hat{e}_1 \downarrow \hat{e}_2; M}$ Closing Fine expressions

$$\frac{\sigma \vdash \sigma \downarrow M \quad \hat{e}_2 = \rho \hat{e}_1 \quad \rho = \{[\ell/s] \mid \ell \in FV(\hat{e}_1) \wedge \sigma \vdash \ell \downarrow \hat{v}; L\}}{\sigma \vdash \hat{e}_1 \downarrow \hat{e}_2; M}$$

$\boxed{\sigma; \Gamma; \mathbf{Y} \vdash e \Longrightarrow C \rightsquigarrow \hat{e}; M}$ Top-level judgment: Translating FX expressions to Fine programs with token stores

$$\text{T-Trans} \frac{\Gamma; \mathbf{Y} \vdash e \Longrightarrow C \rightsquigarrow \hat{e}_1 \quad \sigma \vdash \hat{e}_1 \downarrow \hat{e}; M}{\sigma; \Gamma; \mathbf{Y} \vdash e \Longrightarrow C \rightsquigarrow \hat{e}; M}$$

Figure 6. Translating FX runtime configurations to Fine

Given an FX signature S , we derive a Fine signature \hat{S} , using the following judgment: $td, S \rightsquigarrow \hat{S}_{td}, \hat{S}$ when $td \rightsquigarrow \hat{S}_{td}$ and $S \rightsquigarrow \hat{S}$. The translation of a single type declaration is defined below:

$$\begin{aligned}
(\tau = \{\{f_1:\tau_1; \dots; f_n:\tau_n \mid \Phi\}\}) &\rightsquigarrow \\
& \begin{array}{l}
(\text{prims}, \tau::\star), \quad (\text{prims}, \text{Token}_{\tau}::\tau \Rightarrow \mathbf{A}) \\
(\text{prims}, \text{MkToken}_{\tau} : x:\tau \rightarrow \text{Token}_{\tau} x) \\
(\text{prims}, D_{\tau} : f_1:\tau_1 \rightarrow \dots \rightarrow f_n:\tau_n \rightarrow \tau), \\
mk_{\tau} : x_1:\tau_1 \rightarrow \dots \rightarrow x_n:\tau_n \rightarrow (\text{tokens}(x_1:\tau_1, \dots, x_n:\tau_n) * \{-: \text{unit} \mid \Phi[x_i/\text{self}.f_i]\}) \rightarrow (x:\tau * \text{Token } x * \{-: \text{unit} \mid \bigwedge_i x.f_i = x_i\}), \\
upd_{f_k} : x:\tau \rightarrow y':\tau_k \rightarrow (\text{tokens}(x:\tau, y':\tau_k) * \text{unit}) \rightarrow (y:\tau_k * x':\tau * \text{tokens}(y:\tau_k, x':\tau) * \{-: \text{unit} \mid y = x.f_k \wedge y' = x'.f_k\}) \\
destr_{\tau} : x:\tau \rightarrow (\text{tokens}(x:\tau) * \text{unit}) \rightarrow (f_1:\tau_1 * \dots * f_k:\tau_k * \text{tokens}(f_1:\tau_1, \dots, f_n:\tau_n) * \{-: \text{unit} \mid \bigwedge f_i = x.f_i\})
\end{array}
\end{aligned}$$

Lemma 4 (Preservation for Fine). *Given (M, \hat{e}_1) and $\hat{S}; \hat{\Gamma}$, such that*

- (1) $\hat{S}; \hat{\Gamma} \models M$
- (2) $\hat{S}; \hat{\Gamma}; X \vdash \hat{e}_1 : \hat{\tau}$, for $X \subseteq \text{dom } M$
- (3) $(M, \hat{e}_1) \xrightarrow[p]{\text{fine}} (M', \hat{e}_2)$

Then, there exists $\hat{\Gamma}', X'$ such that $\hat{S}; \hat{\Gamma}'; X' \vdash \hat{e}_2 : \hat{\tau}$ and $\hat{S}; \hat{\Gamma}' \models M'$.

Furthermore, for $\Delta_X = (\text{dom}(M) \cup \text{dom}(M')) \setminus (\text{dom}(M) \cap \text{dom}(M'))$, if $\text{dom}(M') \supseteq \text{dom}(M)$ then $X' = X \cup \Delta_X$; otherwise $X' = X \setminus \Delta_X$.

Proof. Extension of Fine's soundness proof to account for the three additional base cases, R-New, R-Upd, and R-Destr. □

Lemma 5 (Progress for Fine). *Given (M, \hat{e}_1) and $\hat{S}; \hat{\Gamma}$, such that*

- (1) $\hat{S}; \hat{\Gamma} \models M$
- (2) $\hat{S}; \hat{\Gamma}; X \vdash \hat{e}_1 : \hat{\tau}$, for $X \subseteq \text{dom } M$
- (3) $\hat{e}_1 \neq \hat{v}$

Then, there exists M', \hat{e}_2 such that $(M, \hat{e}_1) \xrightarrow[p]{\text{fine}} (M', \hat{e}_2)$.

Proof. Extension of Fine's soundness proof to account for the three additional base cases, R-New, R-Upd, and R-Destr. □

Equivalence of Fine configurations

| | |
|---|---|
| $(M_1, \hat{e}_1) \cong (M_2, \hat{e}_2)$ | $\frac{M_1 \cong M_2 \quad M_1 \vdash \hat{e}_1 \cong \hat{e}_2}{(M_1, \hat{e}_1) \cong (M_2, \hat{e}_2)} \text{ C-Cfg}$ |
| $M_1 \cong M_2$ | $\frac{(M_2, M_1) \cong M}{(M_1, M_2) \cong M} \cdot \cong \cdot \quad \frac{M_1 \cong M_2}{(M_1, (\text{tok}_l, \hat{v})) \cong (M_2, (\text{tok}_l, \hat{v}))}$ |
| $M \vdash \hat{e}_1 \cong \hat{e}_2$ | $\frac{\text{tok}_l \notin \text{dom}(M)}{M \vdash \hat{u}_1^{\ell:\tau} \cong \hat{u}_2^{\ell:\tau}} \text{ C-Dead}$ |
| $M = M', (\text{tok}_l, \text{MkToken } (\hat{v}_1, \dots, \hat{v}_n)^{\ell:\tau})$ | $\frac{\hat{u}_i = (\hat{v}'_1, \dots, \hat{v}'_n) \quad \exists k, \hat{v}_k \neq \hat{v}'_k}{M \vdash \hat{u}_1^{\ell:\tau} \cong \hat{u}_2^{\ell:\tau}} \text{ C-Stale}$ |
| $M \vdash \hat{e} \cong \hat{e}$ | $\text{C-id} \quad \frac{M \vdash \hat{e}_i \cong \hat{e}'_i \quad \forall i}{M \text{ let } \bar{x} = \hat{e}_1 \text{ in } \hat{e}_2 \cong \text{let } \bar{x} = \hat{e}'_1 \text{ in } \hat{e}'_2} \text{ C-Let}$ |
| $M \vdash \hat{v}_i \cong \hat{v}'_i \quad \forall i$ | $\text{C-Data} \quad \frac{M \vdash \hat{v}_i \cong \hat{v}'_i \quad \forall i}{M \vdash D \hat{v}_1 \dots \hat{v}_n \cong D \hat{v}'_1 \dots \hat{v}'_n} \text{ C-Data} \quad \frac{M \vdash \hat{v}_i \cong \hat{v}'_i \quad \forall i}{M \vdash (\hat{v}_1, \hat{v}_2) \cong (\hat{v}'_1, \hat{v}'_2)} \text{ C-Pair}$ |
| $M \vdash \hat{e}_1 \cong \hat{e}_2$ | $\text{C-Lam} \quad \frac{M \vdash \hat{e}_1 \cong \hat{e}_2}{M \vdash \lambda x:\hat{\tau}. \hat{e}_1 \cong \lambda x:\hat{\tau}. \hat{e}_2} \text{ C-Lam} \quad \frac{M \vdash \hat{v}_i \cong \hat{v}'_i \quad \forall i}{M \vdash (\hat{v}_1 \hat{v}_2) \cong (\hat{v}'_1 \hat{v}'_2)} \text{ C-App}$ |

Lemma 6 (Structural equivalence). *For all (M_1, \hat{e}_1) and (M_2, \hat{e}_2) such that $(M_1, \hat{e}_1) \cong (M_2, \hat{e}_2)$.*

There $M_1 = M_2$, and exists \hat{e} and $\sigma_1 = (x_1, u'_1), \dots, (x_n, u'_n)$ and $\sigma_2 = (x_1, v'_1), \dots, (x_n, v'_n)$

Such that $\hat{e}_1 = \sigma_1 \hat{e}$ and $\hat{e}_2 = \sigma_2 \hat{e}$

And vice versa.

Proof. Easy induction on the structure of the equivalence relation. □

Lemma 7 (Equivalent Fine configurations simulate each other). *Given and FX signature S and its Fine counterpart \hat{S}_{base} such that $S \rightsquigarrow \hat{S}_{base}$. Then, for all $(M_1, \hat{e}_1), (M_2, \hat{e}_2)$ such that*

- A1 $(M_1, \hat{e}_1) \cong (M_2, \hat{e}_2)$
- A2 $\hat{S}_{base}; \hat{\Gamma}(M_1); \text{dom } M_1 \vdash \hat{e}_1 : \tau$
- A3 $\hat{S}_{base}; \hat{\Gamma}(M_1); \text{dom } M_1 \vdash \hat{e}_2 : \tau$
- A4 $(M_1, \hat{e}_1) \xrightarrow{\text{fine}}_p (M'_1, \hat{e}'_1)$

Then, $(M_2, \hat{e}_2) \xrightarrow{\text{fine}}_p (M'_2, \hat{e}'_2)$ And $(M'_1, \hat{e}'_1) \cong (M'_2, \hat{e}'_2)$

Proof. We start by appealing to Lemma 6, and deriving \hat{e}, σ_1 , and σ_2 such that $\hat{e}_1 = \sigma_1 \hat{e}$ and $\hat{e}_2 = \sigma_2 \hat{e}$.

From well-typedness and an inversion of (T-Loc) we find that for each $u^{l:\tau}$ in the range of σ_1, σ_2 ; we have $S \ni (\text{prims}, \tau)$.

We proceed by induction on the structure of the reduction relation (A4).

We consider two main sub-cases.

Sub-case $p = \text{pub}$:

From the metatheory of Fine, specifically from the soundness of Fine's module system, we have a value abstraction result.

Namely, for $\hat{S}; \hat{\Gamma}(M), x : (\text{prims}, \hat{\tau}_x); X \vdash e : \tau$

and for any two u_1 and u_2 such that $\hat{S} \vdash u_1 : \hat{\tau}_x$ and $\hat{S} \vdash u_2 : \hat{\tau}_x$.

If $(M, e[u_1/x]) \xrightarrow{\text{fine}}_{\text{pub}} (M', e')$.

Then, there exists e'' such that $e' = e''[u_1/x]$ and $(M, e[u_2/x]) \xrightarrow{\text{fine}}_{\text{pub}} (M', e''[u_2/x])$.

From this result, and the inverse direction of Lemma 6, we get our conclusion.

Sub-case $p = \text{prims}$:

Sub-sub-case R-New:

We have $e_1 = \text{mk}_{\hat{\tau}} \hat{u}^{l_1}, \dots, \hat{u}^{l_n}(l_1, \dots, l_n, ())$ $e_2 = \text{mk}_{\hat{\tau}} \hat{v}^{l_1}, \dots, \hat{v}^{l_n}(l_1, \dots, l_n, ())$

Since we have $M_1 = M_2$ we have $M_1(l_i) = M_2(l_i) = \text{MkToken } \hat{w}_i : \text{Token } \hat{w}_i$.

From well-typedness of e_1 and e_2 we have $l_i : \text{Token } \hat{u}_i$ and $l_i : \text{Token } \hat{v}_i$.

Together, from convertibility of types, we have $u_i = v_i = l_i$.

Thus, $e_1 = e_2$ and we have the conclusion.

Sub-sub-case R-Upd and R-Destr:

Analogous to (R-New)—the presence of the tokens and $M_1 = M_2$ guarantees that $e_1 = e_2$.

Sub-sub-case R-Ctx: Simple, from the induction hypothesis. □

A store and expression respect ownership constraints on object graph

$\text{RespectsOwnership } \sigma \hat{e} \quad \text{RespectsOwnership } \sigma \hat{e} \iff \text{OwnershipOK } \sigma \{l \mid \text{tok}_l \in \text{FV}(\hat{e})\}$

$\text{OwnershipOK } \sigma L$, for $L = l_1, \dots, l_n$

$\frac{\text{OwnershipOK } (\sigma_2, \sigma_1) L}{\text{OwnershipOK } (\sigma_1, \sigma_2) L} \text{ RO-Perm} \quad \frac{}{\text{OwnershipOK } \bullet L} \text{ RO-Nil}$

$\frac{\text{FreeLocs}(\bar{v}) \cap \text{dom}(\sigma) = \emptyset \quad l \notin \text{Reachable } \sigma L \quad \text{OwnershipOK } \sigma L}{\text{OwnershipOK } (\sigma + \{l \mapsto \{\{f = v\}^\oplus\}) L} \text{ RO-Dead}$

$\frac{\text{FreeLocs}(\bar{v}) \cap \text{dom}(\sigma) = \emptyset \quad \text{Parents } \sigma l \cap \text{Reachable } \sigma L = \emptyset \quad \text{OwnershipOK } \sigma L}{\text{OwnershipOK } (\sigma + \{l \mapsto \{\{f = v\}^\oplus\}) (L, l)} \text{ RO-Root}$

$\frac{l \notin L \quad \text{FreeLocs}(\bar{v}) \cap \text{dom}(\sigma) = \emptyset \quad |\text{Parents } \sigma l \cap \text{Reachable } \sigma L| \leq 1 \quad \text{OwnershipOK } \sigma L}{\text{OwnershipOK } (\sigma + \{l \mapsto \{\{f = v\}^\oplus\}) L} \text{ RO-Nested}$

$\ell \in \text{Reachable } \sigma L$

$\frac{\ell \in L}{\ell \in \text{Reachable } \sigma L} \quad \frac{(\text{Parents } \sigma \ell) \cap (\text{Reachable } \sigma L) \neq \emptyset}{\ell \in \text{Reachable } (\sigma + \{\ell \mapsto _ \}^l) L}$

We write $\text{Parents } \sigma \ell$ for $\{\ell' \mid \ell' \in \text{dom}(\sigma) \wedge \ell \in \text{FreeLocs}(\sigma(\ell'))\}$, and $\text{Ancestors } \sigma \ell$ for the fixed point of this relation. We write $\text{Live } \sigma$ for $\{\ell \in \text{dom}(\sigma) \mid \sigma(\ell) = \{\{f = v\}^\oplus\}\}$

Lemma 8 (Parents). *Given σ, L such that $\text{OwnershipOK } \sigma L$.*

Then $\forall l \in \text{dom } \sigma. |L \cap (l \cup \text{Ancestors } \sigma l)| \leq 1$

Proof. Straightforward, by induction on the definition of $\text{OwnershipOK } \sigma L$. □

Lemma 9 (Canonical derivations for RO).

If $\text{OwnershipOK } \sigma (L_1, L_2)$.

Then, there exists a permutation $(\sigma_0, \sigma_1, \sigma_2)$ of σ ,

Such that, $\text{OwnershipOK } (\sigma_0, \sigma_1, \sigma_2) (L_1, L_2)$.

Where,

(1) $\text{OwnershipOK } (\sigma_0, \sigma_1) L_1$ and $\text{OwnershipOK } \sigma_2 L_2$

(2) $\text{dom}(\sigma_0) = \text{dom}(\sigma) \setminus \text{Live } (\sigma)$

(3) $\text{dom}(\sigma_1) = \text{Reachable } \sigma L_1 \cap \text{Live } (\sigma)$

(4) $\text{dom}(\sigma_2) = \text{Reachable } \sigma L_2 \cap \text{Live } (\sigma)$

Proof. By induction on the definition of $\text{OwnershipOK } \sigma L$, using the fact that the side conditions of the inductive rules are monotonic in σ . □

Lemma 10 (Extensibility of RO).

If $\text{OwnershipOK } \sigma_1 L_1$

and $\text{OwnershipOK } \sigma_2 L_2$.

where, $L_1 \cap L_2 = \emptyset$,

$L'_1 = L_1 \setminus \text{Reachable } \sigma_2 L_2$,

$\forall l \in \text{dom}(\sigma_1). \forall l' \in \text{dom}(\sigma_2). l \in \text{FreeLocs}(\sigma_2(l')) \Rightarrow l' \in \text{Reachable } \sigma_2 L_2$,

and $\forall l \in \text{dom}(\sigma_2). \forall l' \in \text{dom}(\sigma_1). l \in \text{FreeLocs}(\sigma_1(l')) \Rightarrow \sigma_1(l') = \{\{f = v\}\}^\ominus$.

Then, $\text{OwnershipOK } (\sigma_1, \sigma_2) (L'_1, L_2)$.

Proof. By induction on the derivation of $\text{OwnershipOK } \sigma_1 L_1$, using the fact that no dead locations in σ_2 are reachable from L_1, L_2 . □

Lemma 11 (Soundness of small footprints on single steps).

If $\Gamma_\sigma \vdash e : C$, where $C \cong \{X\} - \{-\}$ and $(\sigma, e) \xrightarrow{\tau} (\sigma', e')$.

Then, there exists $\sigma_0, \sigma_1, \sigma'_1$ such that σ is a permutation of σ_0, σ_1 ; and

σ' is a permutation of σ_0, σ'_1 ; and

$\text{dom}(\sigma_1) = X$ and $\text{FreeLocs}(\sigma_2) = \text{dom } \sigma_1 \cup L$, where L fresh.

Proof. By induction on the derivation of the transition. □

Lemma 12 (Input variables correspond to tokens). *If $\Gamma_\sigma \vdash e \Rightarrow C \rightsquigarrow \hat{e}$*

and $C \cong \{X\} - \{-\}$

and $\hat{S}_{base}; \hat{\Gamma}_\sigma; \text{dom}(\hat{\Gamma}_\sigma) \vdash \hat{e} : \hat{\tau}$.

Then $\text{FreeTokens}(\hat{e}) = X$.

Proof. Induction over the structure of the token-threading judgment. □

Lemma 13 (Memo-ized derivations). *If $\Gamma; \hat{Y} \models e \Rightarrow C \rightsquigarrow \hat{e}$. Then, $\exists \mathbf{Y}. \Gamma; \mathbf{Y} \vdash e \Rightarrow C \rightsquigarrow \hat{e}$*

Proof. Induction on the rewriting judgment. □

Lemma 14 (Input variables correspond to tokens). *If $\Gamma_\sigma \vdash e \Rightarrow C \rightsquigarrow \hat{e}$ and $C \cong \{l, X\} r : \tau \{-\}$. Then $\text{tok}_l \in \text{FV}(\hat{e})$.*

Proof. Induction over the structure of the token-threading judgment. □

Lemma 15 (Type translation). *If $\Gamma_\sigma; \mathbf{Y} \vdash e \Rightarrow C \rightsquigarrow \hat{e}$. Then, exists $\hat{C}, \Gamma_\sigma \vdash C \Rightarrow \hat{C}$.*

Proof. Induction over the structure of the token-threading judgment. □

Lemma 16 (Well-typed translations preserve types). *If $\Gamma_\sigma \vdash e \Rightarrow C \rightsquigarrow \hat{e}$ and $\hat{\Gamma}_\sigma \vdash \hat{e} : \hat{\tau}$ and $\Gamma_\sigma \vdash C \rightsquigarrow \hat{C}$. Then $\hat{\Gamma}_\sigma \vdash \hat{\tau} <: \hat{C}$.*

Proof. Induction over the structure of the token-threading judgment. □

Lemma 17 (Substitution for type translation). *Given $\Gamma, x; \tau \vdash C \rightsquigarrow \hat{C}$, and $\Gamma \vdash v : \tau$. Then, $\Gamma \vdash C[v/x] \rightsquigarrow \hat{C}[\hat{v}/x]$.*

Proof. By the same property for the formulas and type that constitute C , which is proven by structural induction. □

Lemma 18 (Substitution). *Given all of the following:*

(1) $\Gamma, x; \tau; \mathbf{Y} \vdash e \Rightarrow C \rightsquigarrow \hat{e}$

- (2) $\Gamma; Y' \vdash v \implies C' \rightsquigarrow \hat{v}$, where $C' \cong \{ _ \} x: \tau \{ Y' \}$
- (3) $\hat{\Gamma}, x: \hat{\tau}, Y': \bar{\tau}, \text{toks}(x: \tau, Y'): \bar{\tau}' \vdash \hat{e} : \hat{C}$
- (4) $\hat{\Gamma} \vdash \hat{v} : \hat{C}'$

Then,

- (A) There exists $\hat{v}_x, \hat{v}_{Y'}, \hat{v}_{\text{tokens}}$ such that $\hat{v} = (\hat{v}_x, \hat{v}_{Y'}, \hat{v}_{\text{tokens}}, ())$, and
- (B) $\Gamma; \mathbf{Y} \vdash e[v/x] \implies C[v/x] \rightsquigarrow \hat{e}[\hat{v}_x/x][\hat{v}_{Y'}/Y'][\hat{v}_{\text{tokens}}/\text{toks}(x: \tau, Y')]$

Proof. (A) by induction on (2). (B) by induction on (1). □

Lemma 19 (Effects of FX small steps are small). *If $(\sigma_1, e_1) \xrightarrow{\text{fx}} (\sigma_2, e_2)$. Then, one of the following holds true:*

- (1) $\sigma_1 = \sigma_2$
- (2) \exists fresh $l. \sigma_2 = \sigma_1 + (\{l \mapsto \{\overline{\{f = v\}}\}^\oplus\})$ for some $\{\overline{\{f = v\}}\}$.
- (3) \exists unique $l. \sigma_1(l) = \{\overline{\{f = v\}}\}^\oplus \wedge \sigma_2(l) = \{\overline{\{f = v\}}\}^\ominus$ and $\forall l' \neq l. \sigma_1(l') = \sigma_2(l')$
- (4) \exists unique $l. \sigma_1(l) = \{\overline{\{f = v\}}\}^\dagger \wedge \sigma_2(l) = \{\overline{\{v\}}\}^\dagger$ and $\forall l' \neq l. \sigma_1(l') = \sigma_2(l')$

Proof. By induction on the derivation of the step. □

Lemma 20 (Equivalence of dead locations). *For all $\hat{e}, \hat{v}_1^{\hat{\tau}}, \hat{v}_2^{\hat{\tau}}, M_1, M_2$; if $l \notin \text{dom}(M_1)$ and $l \notin \text{dom}(M_2)$, and $M_1 \cong M_2$, then $M_1 \vdash \hat{e}[\hat{v}_1^{\hat{\tau}}] \cong \hat{e}[\hat{v}_2^{\hat{\tau}}]$.*

Proof. Easy, from application of (C-Dead) □

Lemma 21 (Equivalence of stale locations). *For all $\hat{e}, \hat{v}_1^{\hat{\tau}}, \hat{v}_2^{\hat{\tau}}, M_1, M_2$; if $\hat{v}_1 = (u_1, \dots, u_n)$ and $M_1(l) = (w_1, \dots, w_n)$ where $u_k \neq w_k$; and $M_1 \cong M_2$. Then $M_1 \vdash \hat{e}[\hat{v}_1^{\hat{\tau}}] \cong \hat{e}[\hat{v}_2^{\hat{\tau}}]$. Similarly for M_2, \hat{v}_2 .*

Proof. Easy, from application of (C-Stale) □

Theorem 22 (Fine programs simulate FX programs). *For all FX programs e and well-typed stores σ_1 , if all of the following are true*

- (T0) $\sigma_1; \Gamma \sigma_1; \mathbf{Y} \vdash e_1 \implies C_1 \rightsquigarrow \hat{e}_1; M_1$ where $C_1 \cong \{X_1\} r: \tau \{Y\}$
- (T1) $\Gamma \sigma_1 \vdash C_1 \rightsquigarrow \hat{C}_1$
- (F0) $\hat{S}_{\text{base}}; \hat{\Gamma} \sigma_1; \text{dom } M \vdash \hat{e}_1 : \hat{C}_1$
- (R0) *OwnershipOK* $\sigma_1 X_1$
- (E0) $(e_1, \sigma_1) \xrightarrow{\text{fx}} (e_2, \sigma_2)$

Then, there exists M_2, \hat{e}_2 such that

- (1) $(M_1, \hat{e}_1) \xrightarrow[\text{P}]{\text{fine}^*} (M_2, \hat{e}_2)$
- (2) $\exists \mathbf{Y}'$ such that $\sigma_2; \Gamma \sigma_2; \mathbf{Y}' \vdash e_2 \implies C_2 \rightsquigarrow \hat{e}_2; M_2$, where $C_2 \cong \{X_2\} r: \tau \{Y\}$
- (3) $M_2 \vdash \hat{e}_2 \cong \hat{e}_2'$
- (4) $\hat{S}_{\text{base}}; \hat{\Gamma} \sigma_2; \text{dom } M_2 \vdash \hat{e}_2' : \hat{C}_2$ where $\Gamma \sigma_2 \vdash C_2 \rightsquigarrow \hat{C}_2$
- (5) For all σ, σ_1' such that $\sigma_1 = \sigma, \sigma_1'$ and *OwnershipOK* $\sigma_1' X_1$; there exists σ_2' such that $\sigma_2 = \sigma, \sigma_2'$ and *OwnershipOK* $\sigma_2' X_2$.

Proof. By induction on the structure of (E0), the FX reduction relation.

Case Ctx: $e_1 = \mathbf{let } x = e_{a1} \mathbf{in } e_b$

By inversion on (E0), we have $(e_{a1}, \sigma_1) \xrightarrow{\text{fx}} (e_{a2}, \sigma_2)$

and $(e_1, \sigma_1) \xrightarrow{\text{fx}} (\mathbf{let } x = e_{a2} \mathbf{in } e_b, \sigma_2)$

By discrimination, we have (T0) is an application of T-Let-Top

By inversion of (T0), we get

(T1) $\Gamma \sigma_1; \mathbf{Y}.1 \vdash e_{a1} \implies C_{a1} \rightsquigarrow \hat{e}_{a1\text{-open}}$ where $C_{a1} \cong \{X_{a1}\} x: \tau \{Y_a\}$

(T2) $\Gamma \sigma_1; x: \tau; \mathbf{Y}.2 \vdash e_b \implies C_b \rightsquigarrow \hat{e}_{b\text{-open}}$

And $\hat{e}_1 = (\mathbf{let } x, Y_a, \text{toks}(x: \tau, Y_a) = \hat{e}_{a1} \mathbf{in } \hat{e}_b)$

Where, $\hat{e}_{a1} = \rho_1 \hat{e}_{a1\text{-open}}$ and $\hat{e}_b = \rho_1 \hat{e}_{b\text{-open}}$, for $\rho_1 = \{[\hat{v}] \mid M_1(l) = \text{MkToken } \hat{v}\}$

Inverting (F0), we get $\hat{S}_{\text{base}}; \hat{\Gamma} \sigma_1; X_1 \vdash \hat{e}_a : \hat{C}_a'$ where $X_1, X_2 = \text{dom } M_1$.

And from Lemma 16, we get $\hat{C}_a = \hat{C}_a'$.

We apply the induction hypothesis to get:

(A1-finetto-A2) $(M_1; \hat{e}_{a1}) \xrightarrow[\text{P}]{\text{fine}^*} (M_2; \hat{e}_{a2})$

(A2-transto-A2) $\exists \mathbf{Y}'' . \sigma_2; \Gamma \sigma_2; \mathbf{Y}'' \vdash e_{a2} \implies C_{a2} \rightsquigarrow \hat{e}_{a2}'; M_2$

(A1-tokens) $C_{a2} \cong \{X_{a2}\} x: \tau \{Y_a\}$

(A2-sim-A2') $M_2 \vdash \hat{e}_{a2} \cong \hat{e}_{a2}'$

(A2-RO) For all $\sigma, \sigma_1', \sigma_2'$ such that $\sigma_1 = \sigma, \sigma_1'$ and *OwnershipOK* $\sigma_1' X_{a1}$; there exists σ_2' such that $\sigma_2 = \sigma, \sigma_2'$ and

OwnershipOK $\sigma_2' X_{a2}$.

Now, for the goals:

Goal (1): From an application of R-Ctx, using (A1-fineto-A2) we get

$$(M_1; \hat{e}_1) \xrightarrow[\text{P}]{\text{fine}^*} (M_2; \text{let } x, Y_a, \text{toks}(x, Y_a) = \hat{e}_{a2} \text{ in } \hat{e}_b)$$

Goal (2): Next, we derive

$$\sigma_2; \Gamma_{\sigma_2}; (1: \mathbf{Y}'', 2: \mathbf{Y}.2) \vdash \text{let } x = e_{a2} \text{ in } e_b \Rightarrow C' \rightsquigarrow \text{let } x, Y_a, \text{toks}(x: \tau, Y_a) = \hat{e}_{a2} \text{ in } \hat{e}'_b; M_2$$

We do this by applying (T-Trans),

re-using the first premise of (A1-transto-A2) to translate σ_2 to M_2 .

Then, we apply (T-Let-Top) in the second premise,

re-using the second premise (A1-transto-A2) in the first premise of (T-LetTop).

For the second premise of (T-Let-Top),

we first observe that $\text{dom}(\sigma_1) \subseteq \text{dom}(\sigma_2)$ and hence $\Gamma_{\sigma_2} \geq \Gamma_{\sigma_1}$.

Thus, appealing to weakening for the translation judgment,

we apply (T2) for the second premise of (T-Let-Top)

And $ov(C') = ov(C)$, by construction.

Note, \hat{e}_b differs from \hat{e}'_b in some stale locations.

This is the content of the next sub-goal.

Goal (3): Our goal is to show $(M_2, \hat{e}_2) \cong (M_2, \hat{e}'_2)$.

We apply (C-Cfg) and establish $M_2 \cong M_2$ in its first premise

by inverting (A2-Sim-A2') and using its first premise.

For the second premise, we apply (C-Let)

use the second premise of (A2-sim-A2') in its first premise.

What's left to show is $M_2 \vdash \hat{e}_b \cong \hat{e}'_b$.

Recall, $\hat{e}_b = \rho_1 \hat{e}_{b-open}$, for

$$\rho_1 = \{[l/\hat{v}] \mid \sigma_1 \vdash l \downarrow \hat{v} \wedge (l \in \text{dom}(M_1) \Rightarrow M_1(l) = \text{MkToken } \hat{v})\}$$

And, $\hat{e}'_b = \rho_2 \hat{e}_{b-open}$, for

$$\rho_2 = \{[l/\hat{v}] \mid \sigma_2 \vdash l \downarrow \hat{v} \wedge (l \in \text{dom}(M_2) \Rightarrow M_2(l) = \text{MkToken } \hat{v})\}$$

Where, $\sigma_1 \vdash \sigma_1 \downarrow M_1$ and $\sigma_2 \vdash \sigma_2 \downarrow M_2$.

Appealing to Lemma 19, we have four sub-cases.

Sub-case (FX step is pure): We have $\sigma_1 = \sigma_2$ and we're done.

Sub-case (FX allocates): We have \exists fresh $l. \sigma_2 = \sigma_1 + (\{l \mapsto \overline{\{f = v\}}\})^\oplus$

In this case, since l is fresh, we have $l \notin \text{FreeVars}(\hat{e}_{b-open})$

And, so, $\forall l' \neq l. \sigma_1 \vdash l' \downarrow \hat{v} \Rightarrow \sigma_2 \vdash l' \downarrow \hat{v}$.

Thus, we have $\hat{e}_b = \hat{e}'_b$.

Sub-case (FX destructs): We have

$$\exists \text{unique } l. \sigma_1(l) = \overline{\{f = v\}}^\oplus \wedge \sigma_2(l) = \overline{\{f = v\}}^\ominus$$

and $\forall l' \neq l. \sigma_1(l) = \sigma_2(l)$

In this case, we have $\forall l. \sigma_1 \vdash l \downarrow \hat{v} \Rightarrow \sigma_2 \vdash l \downarrow \hat{v}$.

Thus, once again, we have $\hat{e}_b = \hat{e}'_b$.

Sub-case (FX updates): We have

$$\exists \text{unique } l. \sigma_1(l) = \overline{\{v\}}^l \wedge \sigma_2(l) = \overline{\{v\}}^l$$

and $\forall l' \neq l. \sigma_1(l) = \sigma_2(l)$

In this case, we consider the minimal set of locations Δ where

$l' \in \Delta \Rightarrow \sigma_1 \vdash l' \downarrow \hat{v}_1$ and $\sigma_2 \vdash l' \downarrow \hat{v}_2$ for $\hat{v}_1 \neq \hat{v}_2$.

It is easy to show that $\Delta = \text{Ancestors } \sigma_1 l$.

For the distinguished location l , since $M_1(\text{tok}_l) \neq M_2(\text{tok}_l)$,

we use Lemma 21 (equivalence of stale locations) to show that

$M_2 \vdash \hat{e}_{b-open}[\hat{v}_1/l] \cong \hat{e}_{b-open}[\hat{v}_2/l]$, where $M_i(\text{tok}_l) = \text{MkToken } \hat{v}_i$

For every other location $l' \neq l \wedge l' \in \Delta$, we will show $\text{tok}_{l'} \notin \text{dom}(M_2)$.

As a consequence, using Lemma 20, we can derive $M_2 \vdash \hat{e}_b \cong \hat{e}'_b$.

First, we have $(e_{a1}, \sigma_1) \xrightarrow{\text{fx}} (e_{a2}, \sigma_2)$.

Applying Lemmas 11 and 14, we have $l \in X_{a1}$ and $\text{tok}_l \in \text{FV}(\hat{e}_{a1})$.

Since we have *OwnershipOK* $\sigma_1 X_1$, where $X_1 = \text{FreeTokens}(\hat{e}_1)$, we have:

$\forall l_i \neq l \in (\text{Ancestors } \sigma_1 l = \Delta). l_i \in \text{Reachable } \sigma_1 X_1$

And, $\sigma_1(l_i) = \overline{\{f_i = v_i\}}^\ominus$, and hence $\text{tok}_{l_i} \notin \text{dom}(M_1)$.

Since, by assumption, $\forall l_i \neq l. \sigma_1(l_i) = \sigma_2(l_i)$, we have $\text{tok}_{l_i} \notin \text{dom}(M_2)$, as required.

Goal (4): Easy from soundness of Fine and the type-translation lemma.

Goal (5):

We start by considering an auxiliary goal, *OwnershipOK* $\sigma_2 X_2$.

Sub-case Cut *OwnershipOK* $\sigma_2 X_2$:

First, from *OwnershipOK* $\sigma_1 X_1$ we conclude *OwnershipOK* $\sigma_1 (X_b, X_{a1})$,

Where, by Lemma 14, $X_{a1} = \text{FreeTokens}(\hat{e}_{a1}) = \text{iv}(C_{a1})$ and $X_b = \text{FreeTokens}(\hat{e}_b)$.
 By Lemma 9 (canonical forms for RO), we have *OwnershipOK* $(\sigma, \sigma') (X_b, X_{a1})$,
 where *OwnershipOK* $\sigma' X_{a1}$ and $\text{dom}(\sigma') = \text{Reachable } \sigma_1 X_{a1} \cap \text{Live}(\sigma_1)$
 Next, by Lemma 11 (small footprints) and (T1), we have that $l \notin X_{a1} \Rightarrow \sigma_1(l) = \sigma_2(l)$.
 Since we have *OwnershipOK* $\sigma' X_{a1}$, we clearly have $\text{dom } \sigma \cap X_{a1} = \emptyset$ and hence $\sigma_2 = (\sigma, \sigma'')$,
 where (FL-Sigma-2) $\text{FreeLocs}(\sigma'') = \text{dom } \sigma' \cup L$, for L fresh.
 From (A2-RO), we get that *OwnershipOK* $\sigma'' X_{a2}$, and our goal is now to show *OwnershipOK* $(\sigma, \sigma'') (X_b, X_{a2})$.
 We apply Lemma 10 (extensibility for RO) to *OwnershipOK* σX_b and *OwnershipOK* $\sigma'' X_{a2}$
 From soundness of affine typing, we have that $X_{a2} \cap X_b = \emptyset$, since these correspond to tokens.
 We are left with the sub-goal $\forall l \in \text{dom}(\sigma''). \exists l'. l \in \text{FreeLocs}(\sigma(l')) \Rightarrow l' \notin \text{Live}(\sigma)$
 However this follows from (FL-Sigma-2) and *OwnershipOK* $\sigma, \sigma' X_b, X_{a1}$, and $\text{dom } \sigma' = \text{Reachable } \sigma_1 X_{a1}$,
 Which gives us $l \in \text{dom } \sigma' \Rightarrow l \notin \text{FreeLocs}(\text{Live } \sigma)$, as needed.

We return now to the main goal:

For all $\sigma, \sigma'_1, \sigma'_2$ such that $\sigma_1 = \sigma, \sigma'_1$ and *OwnershipOK* $\sigma'_1 X_1$; there exists σ'_2 such that $\sigma_2 = \sigma, \sigma'_2$ and *OwnershipOK* $\sigma'_2 X_2$.

We start by appealing to canonical forms on *OwnershipOK* $\sigma_2 X_2$.

This gives us *OwnershipOK* $(\sigma_0, \sigma'_2) X_2$, where $\text{dom } \sigma_0 \cap \text{Reachable } \sigma_2 X_2 = \emptyset$

And, importantly, *OwnershipOK* $\sigma'_2 X_2$, for $\text{dom}(\sigma'_2) = \text{Reachable } \sigma_2 X_2$.

I.e., σ'_2 is the smallest fragment of σ_2 on which *RespectsOwnership* holds true.

Now, consider some arbitrary partition of $\sigma_1 = \sigma'_0, \sigma'_1$ where *OwnershipOK* $\sigma'_1 X_1$.

By small footprints, Lemma 11, $\sigma_2 = \sigma'_0, \sigma'_2$.

By construction, since σ'_2 is the smallest fragment of σ_2 on which *RespectsOwnership* holds,

$\sigma'_2 = \sigma''_0, \sigma'_2$, where $\sigma''_0 \subseteq \sigma_0$.

Since we have that $\sigma_0 \cap \text{Reachable } \sigma_2 X_2 = \emptyset$ and *OwnershipOK* σ''_0 .

we apply Lemma 10 (extensibility for RO), to get *OwnershipOK* $\sigma''_0 X_2$, as needed.

Case Let: $e = \text{let } x = v_a \text{ in } e_b$

By inversion on (E0), we have $(e, \sigma) \xrightarrow{\text{fx}} (e_b[v_a/x], \sigma)$

By discrimination, we have (T0) is an application of (T-Let-Top)

By inversion of (T0), we get

(T1) $\Gamma\sigma; \mathbf{Y}.1 \vdash v_a \Rightarrow C_a \rightsquigarrow \hat{v}_a$, where $C_a \cong \{-\}x:\tau\{Y_a\}$

(T2) $\Gamma\sigma, x:\tau; \mathbf{Y}.2 \vdash e_b \Rightarrow C_b \rightsquigarrow \hat{e}_b$

And $\hat{e} = \text{let } x, Y_a, \text{toks}(x:\tau, Y_a), - = \hat{v}_a \text{ in } \hat{e}_b$

Inverting (F0), and using Lemma 16, we get

(Fa) $\hat{S}_{base}; \hat{\Gamma}\sigma; X_a \vdash \hat{v}_a : \hat{C}_a$

(Fb) $\hat{S}_{base}; \hat{\Gamma}\sigma, x:\tau, Y_a : \tau_{Y_a}, \text{toks}(x:\tau, Y_a), - : \text{unit}; X_b \vdash \hat{e}_b : \hat{C}_b$

We appeal to the translation of computation types to note that

(Ca-Tuple) $\hat{C}_a = \Sigma(x:\tau, Y_a, \text{toks}(x:\tau, Y_a), \text{unit})$

Appealing to canonical forms in Fine, and well-typedness of \hat{v}_a , we conclude

(Va-Tuple) $\hat{v}_a = (\hat{v}_{ax}, \hat{v}_{Y_a}, \hat{v}_{\text{tokens}}, ())$

Now, for the goals:

Goal (1): We use (Fine-Let) to derive:

$$(M_1; \text{let } x, Y_a, \text{toks}(x:\tau, Y_a), - = \hat{v}_a \text{ in } \hat{e}_b) \xrightarrow{\text{fine}} (M_1; \hat{e}_b[\hat{v}_{ax}/x][\hat{v}_{Y_a}/Y_a][\hat{v}_{\text{tokens}}, \text{toks}(x:\tau, Y_a)])$$

Goal (2): From Lemma 18, with (T2), (T1), (Fb) and (Fa), we derive:

$$\Gamma\sigma; \mathbf{Y}.2 \vdash e_b[v_a/x] \Rightarrow C_b[v_a/x] \rightsquigarrow \hat{e}_b[\hat{v}_{ax}/x][\hat{v}_{Y_a}/Y_a][\hat{v}_{\text{tokens}}, \text{toks}(x:\tau, Y_a)]$$

(Using determinism of canonical forms to identify \hat{v}_{ax} etc.)

Goal (3): We need to show:

$$M_1 \vdash \hat{e}_b[\hat{v}_{ax}/x][\hat{v}_{Y_a}/Y_a][\hat{v}_{\text{tokens}}, \text{toks}(x:\tau, Y_a)] \cong \hat{e}_b[\hat{v}_{ax}/x][\hat{v}_{Y_a}/Y_a][\hat{v}_{\text{tokens}}, \text{toks}(x:\tau, Y_a)]$$

Which is trivial from an application of (C-Id).

Goal (4): Easy, from the soundness of Fine.

Goal (5): For any partition of $\sigma_1 = \sigma_0, \sigma'_1$ such that *OwnershipOK* $\sigma'_1 X_1$, we have to show *OwnershipOK* $\sigma'_1 X_2$.

However, $X_2 = \text{FreeTokens}\hat{e}_b[\hat{v}_{ax}/x][\hat{v}_{Y_a}/Y_a][\hat{v}_{\text{tokens}}, \text{toks}(x:\tau, Y_a)]$

and $X_1 = \text{FreeTokens}\hat{e}_b, \hat{v}_{ax}, \hat{v}_{Y_a}, \hat{v}_{\text{tokens}}$.

Thus, we have $X_2 \subseteq X_1$, and *OwnershipOK* $\sigma'_1 X_2$ follows readily.

Case Fun: Much like (Let), using Lemma 18 in the special case where $Y' = \emptyset$

Case New: $e_1 = \text{new}_\tau \{\overline{f = v}\}$

And we have $(e_1, \sigma_1) \xrightarrow{\text{fx}} (l, \sigma_2)$ where $\sigma_2 = (\sigma_1 + \{l \mapsto \{\overline{f = v}\}\})$

By discrimination, we have (T0) is an application of (T-Base), with (T-Y) in the premise, with (T-New) in its premise.

So, we have

$\hat{e}_1 = \text{let } x, \text{tok}_{x,-} = \text{mk}_\tau \hat{v}_1 \cdots \hat{v}_n (\text{tokens}(\cdots), ()) \text{ in } x, \hat{v}_Y, \text{toks}(x, \hat{v}_Y)$

where $\hat{v}_Y = \{\hat{v}_{y_1} \mid y \in Y \wedge \sigma_1 \vdash y \downarrow \hat{v}_{y_1}\}$

Goals:

Goal (1) We aim to use (R-Ctx) with (R-New) in the premise to get

$(M_1; \hat{e}_1) \xrightarrow{\text{fine}} (M_2; \text{let } x, \text{tok}_{x,-} = (\hat{v}, \text{tok}_l, ()) \text{ in } x, \hat{v}_Y, \text{tokens}(x, \hat{v}_Y)) \xrightarrow{\text{fine}} (M_2; (\hat{v}, \hat{v}_Y, \text{tok}_l, \text{toks}(\hat{v}_Y)))$

For the first step, we need to show $M_1 = M_0, \text{TokenVal}(\hat{v}_1), \dots, \text{TokenVal}(\hat{v}_n)$.

For this, we appeal to well-typedness of the Fine program and the base signature \hat{S}_{base} .

We have in \hat{S}_{base} the binding $\text{mk}_\tau : (x_1 : \hat{\tau}_1 * \dots * x_n : \hat{\tau}_n * \text{Token } x_1 * \dots * \text{Token } x_n * \{-:\text{unit} \mid \Phi\}) \rightarrow \hat{\tau}$.

Where $\text{MkToken} : x : \tau_i \rightarrow \text{Token } x$.

Thus, soundness of dependent typing in Fine gives us the premise of (R-New) and we get

$M_2 = M_0, (\text{tok}_l, \hat{v})$, for l fresh and $\hat{v} = \text{MkToken}(l, \hat{v}_1, \dots, \hat{v}_n)^{l:\hat{\tau}}$.

Goal (2)

We derive $\Gamma \sigma_2; Y \vdash l \Longrightarrow \{l, Y\} x : \tau \{Y\} \rightsquigarrow \hat{e}'_2; M_2$ using (T-Base) with (T-Return).

where $\hat{e}'_2 = \rho((l, Y, \text{toks}(r, Y)))$ and $\rho = \{l' / \hat{v}' \mid l' \in (l, Y) \wedge \sigma_2 \vdash l' \downarrow \hat{v}' \wedge (\text{tok}_l \in \text{dom } M_2 \Rightarrow M_2(\text{tok}_l) = \text{MkToken } \hat{v})\}$

And we need to show $\sigma_2 \vdash \sigma_2 \downarrow M_2$,

where $\sigma_2 = \sigma_1, (l \mapsto \{\overline{f = v}\}^\oplus)$ and $\sigma_1 = (l_1 \mapsto v_1^\oplus), \dots, (l_n \mapsto v_n^\oplus)$

and $\sigma_1 \vdash \sigma_1 \downarrow M_1$ and $\sigma_1 \vdash v_i \downarrow \hat{v}_i$.

To do this, we start with an application of (M-1), and obtain $M'_0, (\text{tok}_l, \hat{v})$ where $\sigma_2 \vdash \sigma_0 \downarrow M'_0$.

To show that $M'_0 = M_0$, we appeal to *OwnershipOK* $\sigma_1 (\text{tokens}(l_1, \dots, l_n))$ which ensures that $\{l_1, \dots, l_n\} \cap (\text{FreeLocs}(\sigma_0) \cap \text{Live } \sigma_0) = \emptyset$.

Goal (3) We need to show $M_2 \vdash (\hat{v}, \hat{v}_Y, \text{tok}_l, \text{toks}(\hat{v}_Y)) \cong \rho((l, Y, \text{toks}(r, Y)))$

By construction, we have $\rho(l) = \hat{v}$, but, for $y \in Y$, $\rho(y) = \hat{v}_{y_2}$, where $\sigma_2 \vdash y \downarrow \hat{v}_{y_2}$

We need to show that $\sigma_2 \vdash y \downarrow \hat{v}_{y_2} \iff \sigma_1 \vdash y \downarrow \hat{v}_{y_2}$.

Suffice to show that $Y \subseteq \text{dom}(\sigma_0)$.

This follows from affine typing of Fine and freshness of l , which ensures that $\{\text{tok}_y \mid y \in Y\} \cap \{l_1, \dots, l_n, l\} = \emptyset$.

Goal (4) Well-typedness of the \hat{e}'_2 follows from soundness of Fine.

Goal (5) Our goal is essentially to show *OwnershipOK* $\sigma_2 l, Y$, given *OwnershipOK* $\sigma_0 Y$.

Suffice to show *OwnershipOK* $((l_1 \mapsto v_1^\oplus), \dots, (l_n \mapsto v_n^\oplus), (l \mapsto \{\overline{f = v}\}^\oplus)) l$.

This follows from a single application of (RO-Root) with (RO-Nil).

We reach the goal using an application of Lemma 10.

Case FldUpd: $e_1 = l.f_k := v'_k$

And we have $(e_1, \sigma_1) \xrightarrow{\text{fx}} (v_k, \sigma_2)$ where

$\sigma_1 = (\sigma_0 + \{l \mapsto \{\overline{f_i = v_i}\}_{i=1}^n\})$ and

$\sigma_2 = (\sigma_0 + \{l \mapsto \dots; f_k = v'_k; \dots\})$

By discrimination, we have (T0) is an application of (T-Base), with (T-Y) in the premise, and (T-Upd) in its premise.

So, we have

$\hat{e}_1 = \text{let } x, y, \text{tok}_{x:\tau,y,-} = \text{upd}_{f_k}(\hat{v}, \hat{v}'_k, \text{toks}(\hat{v}, \hat{v}'_k), ()) \text{ in } x, \hat{v}_Y, \text{tok}_{x:\tau,\hat{v}_Y}$

Goals:

Goal (1):

We aim to use (R-Ctx) with (R-Upd) in the premise to get

$(M_1; \hat{e}_1) \xrightarrow{\text{fine}} (M_2; \text{let } x, y, \text{tok}_{x:\tau,y,-} = \hat{v}_k, \hat{v}', \text{toks}(\hat{v}, \hat{v}'_k), () \text{ in } x, \hat{v}_Y, \text{tok}_{x:\tau,\hat{v}_Y}) \xrightarrow{\text{fine}} (M_2; (\hat{v}_k, \hat{v}_Y, \text{tok}_{x:\tau,\hat{v}_Y}))$

For the first step, we need to show $M_1 = M_0, (\text{TokenVal}(\hat{v}), \text{TokenVal}(\hat{v}'_k))$.

For this, we appeal to well-typedness of the Fine program and the base signature \hat{S}_{base} .

We have in \hat{S}_{base} the binding $\text{upd}_{f_k} : (x : \hat{\tau} * f : \hat{\tau}_k * \text{Token } x * \text{Token } x * \{-:\text{unit} \mid \Phi\}) \rightarrow \hat{\tau}'$.

Where $\text{MkToken} : x : \tau_i \rightarrow \text{Token } x$.

Thus, soundness of dependent typing in Fine gives us the premise of (R-New) and we get $M_2 = M_0, \text{TokenVal } (\hat{v}'), \text{TokenVal } (\hat{v}_k)$, for $\hat{v}' = \text{MkToken } (j, \hat{v}_1, \dots, \hat{v}'_k, \dots, \hat{v}_n)^{l:\hat{t}}$.

Goal (2):

We derive $\Gamma_{\sigma_2}; Y \vdash v_k \implies \{-\}x:\tau\{Y\} \rightsquigarrow \hat{e}'_2; M_2$ using (T-Base) with (T-Return).

where $\hat{e}'_2 = \rho((\hat{v}_k, Y, \text{toks}(\hat{v}_k, Y)))$ and $\rho = \{[l'/\hat{v}'] \mid l' \in (l, Y) \wedge \sigma_2 \vdash l' \downarrow \hat{v}' \wedge (\text{tok}_l \in \text{dom } M_2 \Rightarrow M_2(\text{tok}_l) = \text{MkToken } \hat{v}')\}$

And we need to show $\sigma_2 \vdash \sigma_2 \downarrow M_2$,

Which follows along the lines of the previous case (New).

Goal (3) We need to show $M_2 \vdash (\hat{v}_k, \hat{v}_Y, \text{tok}_l, \text{toks}(\hat{v}_k, \hat{v}_Y)) \cong \rho((\hat{v}_k, Y, \text{toks}(\hat{v}_k, Y)))$

By construction, if $l \in Y$ then $\rho(l) = \hat{v}'$, but, for $y \neq l \in Y$, $\rho(y) = \hat{v}_{y_2}$, where $\sigma_2 \vdash y \downarrow \hat{v}_{y_2}$

We need to show that $\sigma_2 \vdash y \downarrow \hat{v}_{y_2} \iff \sigma_1 \vdash y \downarrow \hat{v}_{y_2}$.

Suffice to show that $(Y \setminus \{l\}) \subseteq \text{dom}(\sigma_0)$.

This follows from affine typing of Fine, which ensures that $\{\text{tok}_y \mid y \neq l \in Y\} \cap \{l, \hat{v}_k\} = \emptyset$.

Goal (4) Well-typedness follows from soundness of Fine.

Goal (5) *OwnershipOK* σ_2 $\text{tokens}(\hat{v}_k, Y)$, follows from construction.

Case Destr: $e_1 = \text{destruct } \ell \text{ as } \{\{f = \bar{z}\}\} \text{ in } e_b$

And, we have $(e_1, \sigma_1) \xrightarrow{\text{fin}} (e_b [v_i/x_i]_{i=1}^n, \sigma_2)$

where $\sigma_1 = \sigma_0, (l \mapsto \{\{f = v\}\}^l)$

and $\sigma_2 = \sigma_0, (l \mapsto \{\{f = v\}\}^\ominus)$.

By inversion of (T0), we have

$\hat{e}_1 = \text{let } \bar{z}_i, \text{tok}_{\bar{z}_i}, - = \text{in destr}_\tau \hat{v} \text{ tok}_v \hat{e}_b$

where $\sigma_1 \vdash l \downarrow \hat{v}$

Goal (1) We aim to use (R-Ctx) with (R-Destr) followed by (R-Let) to derive:

$$(M_1; \hat{e}_1) \xrightarrow{\text{fin}} (M_2; \text{let } \bar{z}_i, \text{toks} = \hat{v}, \text{tokens}(\hat{v}) \text{ in } \hat{e}_b) \xrightarrow{\text{fin}} (M_2; \hat{e}_b[\hat{v}, \text{tokens}(\hat{v})/\bar{z}_i, \text{toks}])$$

From *OwnershipOK* $\sigma_1 \text{ tok}_v$, we have $M_1 = M_0, (i, \hat{v}_1, \dots, \hat{v}_n)$

And hence, $M_2 = M_0, \text{TokenVal } \hat{v}_1, \dots, \text{TokenVal } \hat{v}_n$.

Goal (2)

We derive $\Gamma_{\sigma_2}; Y \vdash (e_b [v_i/x_i]_{i=1}^n \implies \{-\}x:\tau\{Y\} \rightsquigarrow \hat{e}'_2; M_2$

using (T-Base) and appealing to the substitution lemma for the translation (as in the Let case).

We show that $\sigma_2 \vdash \sigma_2 \downarrow M_2$ by observing that $\sigma_2 = \sigma_0, (l \mapsto \{\{f = v\}\}^\ominus)$.

while $\sigma_1 = \sigma_0, (l \mapsto \{\{f = v\}\}^\oplus)$.

Hence the contents of $\sigma_1(l)$, previously enclosed within $M_1(\text{tok}_l)$, each receive their own location in M_2 , as required.

Goal (3) To show that $M_2 \vdash \hat{e}_b[\hat{v}, \text{tokens}(\hat{v})/\bar{z}_i, \text{toks}] \cong \hat{e}'_2$

we first observe using the substitution lemma, that $\hat{e}'_2 = \hat{e}_b[\hat{v}', \text{tokens}(\hat{v}')/\bar{z}_i, \text{toks}]$

where $\sigma_1 \vdash v_i \downarrow \hat{v}_i$, while $\sigma_2 \vdash v_i \downarrow \hat{v}'_i$.

However, σ_1 and σ_2 differ only on the l -marker on location l and hence $\bar{v} = \bar{v}'$.

Goal (4) Well-typedness follows from soundness of Fine.

Goal (5) *OwnershipOK* σ_2 $\text{tokens}(\hat{v}_1, \dots, \hat{v}_n), Y$ follows from construction,

since $\text{tok}_l \notin Y, \text{tokens}(\hat{v}_1, \dots, \hat{v}_n)$.

□

Definition 23 (Efficiency Preorder(Arun-Kumar and Hennessy 1992)).

A binary relation $\mathcal{R} \subseteq \text{Fine} \times \text{FX}$ is an efficiency prebisimulation up to \mathcal{S} (on the right) if for all $(e_1, \sigma) \mathcal{R} (M, e_2)$,

(1) if $(e_1, \sigma) \xrightarrow{\text{fx}} (e'_1, \sigma')$ then $(M, e_2) \xrightarrow{\text{fine}}^+ (M', e'_2)$ such that $(e'_1, \sigma') \mathcal{R} \mathcal{S} (M', e'_2)$.

(2) if $(M, e_2) \xrightarrow{\text{fine}} (M', e'_2)$ then $(e_1, \sigma) \xrightarrow{\text{fx}} (e'_1, \sigma')$ such that $(e'_1, \sigma') \mathcal{R} \mathcal{S} (M', e'_2)$.

We write $(e_1, \sigma) \lesssim (M, e_2)$ if there is an efficiency prebisimulation \mathcal{R} up to the identity relation such that $(e_1, \sigma) \mathcal{R} (M, e_2)$.

Not that for all relations \mathcal{S} contained in strong bisimilarity, if $(e_1, \sigma) (M, e_2)$ are related by some efficiency prebisimilarity up to \mathcal{S} then $(e_1, \sigma) \lesssim (M, e_2)$.

Theorem 24. If $\cdot; \Gamma; \cdot \vdash e \implies C \rightsquigarrow \hat{e}; \cdot$, where $C \cong \{\emptyset\}r:\tau\{Y\}$, $\cdot \vdash C \rightsquigarrow \hat{C}$ and $\hat{S}_{\text{base}}; \cdot; \emptyset \vdash \hat{e} : \hat{C}$, then $(e, \cdot) \lesssim (M, \hat{e})$.

Proof. Theorem 22 yields the simulation (1) up to \cong , which is a bisimulation by Lemma 6. The reverse simulation follows from reduction $\xrightarrow[p]{\text{fine}}$ being deterministic. \square

Theorem 25. *Given a closed FX program e , such that $\bullet \vdash e \rightsquigarrow \hat{e}$ and $\hat{S}_{base}; \bullet \vdash \hat{e} : \hat{\tau}$. Then for any reduction sequence $(e, \cdot) \xrightarrow{*} (e', \sigma')$, e' is not **error**.*

Proof. Assume that there is a reduction sequence that ends in **error**.

We derive false by induction on the length of the sequence, using Theorem 22 for the induction case.

In the base case, $e = \mathbf{error}$, but there is no \hat{e} such that $\Gamma; \mathbf{Y} \vdash \mathbf{error} \Longrightarrow C \rightsquigarrow \hat{e}$. \square

B. Example FX programs and their translation to Fine

=====

Example program 1 and its translation

FX program:

```
let x = new_t {f=0} in
let s = x.f := 1 in
s
```

Annotated FX program:

```
let x = new_t {f=0} : {}x:t{} in <---- {} r:int {}
let s = x.f := 1 : {x}s:int{x} in <---- {x}r:int{x}
s : {}r:int{}
```

Note, although the second let has a footprint $\{x\}s:int\{x\}$ we translate it by PUSHING in a type with a smaller output set $\{x\}s:int\{\}$ to respect the scoping rules. I.e., PUSHING can decrease the size of the output set.

Fine program:

```
let x, tok_x, _ = mk_t 0 () in
let s,x,tok_x, _ = upd_f x 1 tok_x () in
s, ()
```

=====

Example program 2 and its translation

FX program:

```
let x = new_t {f=0} in
let s = x.f := 1 in
x
```

Annotated FX program:

```
let x = new_t {f=0} : {}x:t{} in <--- {}r:t{}
let s = x.f := 1 : {x}s:int{x} in <--- {x}r:t{}
x : {x}r:t{}
```

Fine program:

```
let x, tok_x, _ = mk_t 0 () in
let s,x,tok_x, _ = upd_f x 1 tok_x () in
x, tok_x, ()
```

=====

Example program 3 and its translation

FX program:

```
let x = new_t {f=0} in
let y = new_t {f=1} in
let _ = x.f := 2 in
let r = y.f := 3 in
r
```

Annotated FX program:

```
let x = new_t {f=0} : {} x:t {} in <-----{}s:int{}
let y = new_t {f=1} : {} y:t {} in <----- {x}s:int{x}
```

```

let _ = x.f := 2 : {x}_:int{x} in <----- {x,y}s:int{x,y}
let r = y.f := 3 : {y}r:int{y} in <----- {y} s:int {y}
r : {} s:int {}

```

PUSH in {x}s:int{} when checking the body of the first let
PUSH {x,y}s:int{} when checking body of 2nd let
PUSH {y}s:int{} for 3rd let
PUSH {}s:int{} for lastlet

Fine program:

```

let x, tok_x = mk_t 0 in
let y, tok_y = mk_t 1 in
let _, x, tok_x = upd_f x 2 tok_x in
let r, y, tok_y = upd_f y 3 tok_y in
r, ()

```

=====
Example program 4 and its translation

FX program:

```

let x = new_t {f=0} in
let y = new_t {f=1} in
let z =
  let _ = x.f := 2 in
  let yf = y.f := 3 in
  yf in
let _ = x.f := 4 in
()

```

Annotated FX program:

```

let x = new_t {f=0} : {} x:t {} in <----- {}_:unit{}
let y = new_t {f=1} : {} y:t {} in <----- {x}_:unit{x}
let z = <----- {x,y}_:unit{x,y}
  let _ = x.f := 2 : {x}_:int{x} in <--- {x,y}r:int{x,y}
  let yf = y.f := 3 : {y}yf:int{y} in <--- {y}r:int{y}
  yf : {} r:int {} in
let _ = x.f := 4 :a {x} _:int {x} in <--- {x}_:unit{x}
() : {} _:unit {}

```

When translating the nested let, we first synth the type {x,y}_:unit{x,y} for the whole nested let. Then PUSH the type {y}r:int{x,y} when checking the second second sub-expr. I.e., PUSHING can also increase the size of the output set.

Fine program:

```

let x, tok_x, _ = mk_t 0 () in
let y, tok_y, _ = mk_t 1 () in
let z, x, y, tok_x, tok_y, _ =
  let _, x, tok_x, _ = upd_f x 2 tok_x () in
  let yf, y, tok_y, _ = upd_f y 3 tok_y () in
  yf, x, y, tok_x, tok_y, () in
let _ = upd_f x 4 tok_x () in
()

```

=====
Example program 5 and its translation

FX program:

```

let x = new_t {f=0} in
let f = \y:T. y.g := x in

```

```

let x1 = new_t {f=1} in
let x2 = new_t {f=2} in
let y1 = new_T {g=x1} in
let y2 = new_T {g=x2} in
let _ = f y1 in
let _ = f y2 in
  ()

```

Annotated FX program:

```

let x = new_t {f=0} : {}x:t{} in <-----{}_:_unit{}
let f = \y:T. y.g := x : {} (y:T -> {y,x} r:t {y}) {} in <--- {x}_:_unit {}
let x1 = new_t {f=1} : {} x1:t {} in <--- {x}_:_unit {}
let x2 = new_t {f=2} : {} x2:t {} in <--- {x1,x}_:_unit {}
let y1 = new_T {g=x1} : {x1} y1:T {} in <--- {x1,x,x2}_:_unit {}
let y2 = new_T {g=x2} : {x2} y2:T {} in <--- {y1,x,x2}_:_unit {y1}
let _ = f y1 : {y1,x}_:_t {y1} in <---- {y1, y2, x}_:_unit {y1, y2}
let _ = f y2 : {y2,x}_:_t {y2} in <--- {y2,x}_:_unit {y2}
  () : {}_:_unit{}

```

The issue here is that the closure `f` captures the object variable `x` and is then invoked twice. This program fails to type check (as intended) due to an affinity failure on the token of `x`. But, notice that no closure conversion of `f` is necessary. And promotion of the pure function closure into a computation type is automatic.

Fine program:

```

let x, tok_x , _ = mk_t 0 () in
let f , _ = (\y:T.\tok_y:Token y.\tok_x:Token x.
  upd_g y x tok_y tok_x ()), () in
let x1, tok_x1 , _ = mk_t 1 () in
let x2, tok_x2 , _ = mk_t 2 () in
let y1, tok_y1 , _ = mk_T x1 tok_x1 () in
let y2, tok_y2 , _ = mk_T x2 tok_x2 () in
let x1', y1, tok_x1', tok_y1 , _ = f y1 tok_y1 tok_x () in
let x2', y2, tok_x2', tok_y2 , _ = f y2 tok_y2 tok_x () in <-- affinity failure
  ()

```

Example program 6

FX program:

```

let x = new_t {f=1} in
let f = \y:int. x.f += y in
  f 1; f 1

```

FX semantics:

```

  After first call to f, x={f=2}
  After second call to f, x={f=3}

```

Annotated FX program:

```

let x = new_t {f=1} : {} x:t {} in
let f = \y:int. x.f += y : {}f:(int -> {x} r:int {x}){} in
let _ = f 1 : {x} r:int {x} in
let _ = f 1 {x} r:int {x} in
  ()

```

where

```

(x.f += y) \def= let tmp = x.f := 0 in let z = x.f := tmp + y in z

```

After translation to fine:

```
let x, tok_x = mk_t 1 in
let f =
  \y:int. \tok_x.
    let tmp, x, tok_x = upd_f x 0 tok_x in
    let z, x, tok_x = upd_f x (tmp+y) tok_x in
      z, x, tok_x in
let _, x, tok_x = f 1 tok_x in
let _, x, tok_x = f 1 tok_x in <---- fails because wrong value index on token
()
```

=====
=====
Example program 7 and its translation.
Affinity in Fine ensures safety

FX program:

```
foo: y:T -> {(s)s(y).g.f=0} unit {(t)...}

let x = new_t {f=0} in
let y = new_T {g=...} in
let _ = y.g := x in
let _ = x.f := 17 in
  foo y
```

Annotated FX program:

```
let x = new_t {f=0} : {} x:t {(s1) s1(x).f=0} in
let y = new_T {g=...} : {} y:T {(s2) s1(x)=s2(x) && .. } in
let _ = y.g := x : {(s2)} _:_ {(s3) s3(y).g=s2(x)} in
let _ = x.f := 17 : {(s3)} _:_ {(s4) s4(x)=17 && s4(y)=s3(y) } in
  foo y : {(s4) s4(y).g.f=0} unit {(s5)...} <----- Pre-condition check succeeds
                                     s4(y).g=s3(y).g=s2(x)=s1(x)={f=0}
```

Translated to Fine:

```
let x, tok_x = mk_t 0 in
let y, tok_y = mk_T ... in
let _, y, tok_y = upd_T y x tok_y tok_x in
let _, x, tok_x = upd_t x 17 tok_x in <-----affinity failure
  foo y tok_y
```