Rely-Guarantee References for Refinement Types Over Aliased Mutable Data (Extended Version [*])

Colin S. Gordon, Michael D. Ernst, and Dan Grossman

University of Washington

{csgordon,mernst,djg}@cs.washington.edu

Abstract

Reasoning about side effects and aliasing is the heart of verifying imperative programs. Unrestricted side effects through one reference can invalidate assumptions about an alias. We present a new type system approach to reasoning about safe assumptions in the presence of aliasing and side effects, unifying ideas from reference immutability type systems and rely-guarantee program logics. Our approach, *rely-guarantee references*, treats multiple references to shared objects similarly to multiple threads in rely-guarantee program logics. We propose statically associating rely and guarantee conditions with individual references to shared objects. Multiple aliases to a given object may coexist only if the guarantee condition of each alias implies the rely condition for all other aliases. We demonstrate that existing reference immutability type systems are special cases of rely-guarantee references.

In addition to allowing precise control over state modification, rely-guarantee references allow types to depend on mutable data while still permitting flexible aliasing. Dependent types whose denotation is stable over the actions of the rely and guarantee conditions for a reference and its data will not be invalidated by any action through any alias. We demonstrate this with refinement (subset) types that may depend on mutable data. As a special case, we derive the first reference immutability type system with dependent types over immutable data.

We show soundness for our approach and describe experience using rely-guarantee references in a dependently-typed monadic DSL in CoQ.

1. Introduction

A common way to reason about side effects in imperative languages is to restrict (disable) mutating some state in some code sections. This is seen most clearly in reference immutability [27, 50, 55, 56], but also in ownership [17] and region-based type systems [6]. The common approach is to attach permission/ownership/region information to references, where certain operations (mainly writes to the heap) through references with certain permissions are prohibited.

The program logic literature includes work ensuring that actions by one section of code do not interfere destructively with the assumptions of another section of code. This appears most often in the form of concurrent program logics, where the goal is to prevent destructive interference between threads. This reaches at least as far back as Owicki and Gries's technique [43], which checks thread non-interference by ensuring that no action would invalidate any intermediate assumption of another thread. Jones abstracted cross-thread interactions to a *rely* relation bounding interference by other threads, and a *guarantee* relation bounding actions of the current thread [31]. Each thread's local proof then requires all local actions to fall within its guarantee, and that all of its intermediate assertions are stable with respect to (that is, not invalidated by) any possible action permitted by the rely. Parallel composition of threads is then safe if each thread's guarantee implies each other thread's rely.

Our central idea is to treat aliases to objects similarly to threads of control in rely-guarantee program logics. Each reference's type carries a rely and a guarantee, bounding actions on an object through other references (rely) and bounding actions through the reference itself (guarantee). We call these augmented reference types *rely-guarantee references*. The type system maintains the invariant that the guarantee of any reference implies the rely of any alias. The type system checks these constraints when a program duplicates an alias. This raises the issue that some references cannot soundly coexist: no two references to the same object can each guarantee nothing (the reference permits arbitrary actions) and rely on restricted behavior through aliases. This presents us with a logical account of aliasing: some references may not be aliased without weakening the rely or guarantee of the source, and a reference with an empty rely necessarily has no aliases.

Rely-guarantee references generalize *reference immutability* [27, 50, 55, 56] to finer-grained control over interference through aliases. The traditional reference immutability qualifiers correspond to simple rely and guarantee conditions. For ref τ[R,G] as a reference to data of type τ with rely R and guarantee G:

- readable τ ≡ ref τ[any interference, no writes]
- writable τ ≡ ref τ[any interference, any writes]
- immutable τ ≡ ref τ[no interference, no writes]

Rely-guarantee references let us reason about some refinements of referents. Let a *stable* predicate over a reference be one that is preserved by its rely. Then a stable predicate cannot be invalidated by actions through an alias, and any new predicate that is stable and ensured by a guarantee-permitted action (on an object satisfying the old predicate) is true after the action, providing a form of strong update on arbitrarily aliased mutable data. (An action allowed by the guarantee that preserves the current predicate is a special case.)

1.1 Contributions

Refinement Types Over Mutable Data Rely-guarantee references permit refinement types [25] that depend on mutable data, without requiring any aliasing restrictions to support strong updates. We leverage the notion of a *stable assertion* from rely-guarantee program logics, allowing any refinements that are not invalidated by actions performed through other references. We prove that our type system is sound.

Generalizing Reference Immutability We generalize reference immutability by combining it with rely-guarantee techniques. This is of independent interest, but also outlines an effort/precision spectrum from unrestricted references to reference immutability to rely-guarantee references.

A Prototype Implementation We prototype an implementation as a shallow monadic embedding in COQ. We have used it to verify the examples in the paper, including implementing reference immutability as a special case. We briefly discuss our experience implementing a language as a COQ DSL and the manual proof burden for our technique versus purely functional versions. The implementation is available at
`https://github.com/csgordon/rgref/`.

We believe rely-guarantee references make a compelling argument that rely-guarantee reasoning is a promising way to statically reason about aliasing. Further, any technique traditionally used to reason about thread interference can be adapted to modularly reason about effects in the presence of aliasing (we present rely-guarantee references as a type system, but our ideas could be implemented in other ways, such as a program logic). Ultimately we believe the proper way to support unknown aliases in program verification is by treating aliases as different threads of control.

2. Rely-Guarantee References

A rely-guarantee reference is a reference to a heap structure of a given type, as in ML's ref τ, with three additional type components:

- A refinement predicate P over the τ and a heap h that can enforce local properties and/or data-structure well-formedness.

- A guarantee relation G over pairs of τs and heaps, restricting the effects to the referent (and state heap-reachable from that referent) that may be performed through this reference or those produced by dereferencing it.

- A rely relation R specifying the actions permitted by (the guarantees of) other aliases to the referent.

We use the form ref$\{\tau \mid P\}[R,G]$ for a rely-guarantee reference. Predicates and relations are defined not only over the τ a reference refers to, but also over heaps, to refine data reachable from the immediate referent. For a rely-guarantee reference type to be well-formed, the predicate P must be *stable* with respect to the rely R: for all values and heaps for which the predicate holds, if the rely R allows another value and heap to be produced by actions on another alias, then the predicate holds for the new value and heap as well: $P\,v\,h \wedge R\,v\,v'\,h\,h' \implies P\,v'\,h'$. This ensures that actions through aliases do not invalidate the refinement, and that all actions that may invalidate the refinement are local, so reasoning about such changes allows strong updates to the refinement. These issues are formally treated in Section 4.

A simple example of rely-guarantee references is a monotonically increasing counter, which we can represent as a value of type

$$\text{ref}\{\texttt{nat} \mid \textsf{any}\}[\textsf{increasing},\textsf{increasing}]$$

where any is the trivial (always true) refinement, and increasing (Section 3.1) is a relation on natural numbers and heaps that requires the second nat to be greater than or equal to the first. Given a variable x with the type above, $x \leftarrow\ !x+1$ type-checks (! is ML's dereference operator). By contrast, incorrect code that decrements the counter cannot satisfy the guarantee relation increasing.

A read-only alias to an increasing counter can be expressed as:

$$\text{ref}\{\texttt{nat} \mid \textsf{any}\}[\textsf{increasing},\approx]$$

where $\approx$ is a relation permitting no change.

We might wish to know more about a counter value, for example that it is greater than 0 so it is safe as a divisor to compute an average. Any write to the counter via any reference will increase its value, and may therefore conclude the result is greater than 0.[1] Furthermore, it is safe to continue assuming the value is greater than 0 because the reference's rely ensures no alias can decrease the value. We say $\lambda x : \texttt{nat}.\lambda h : \textsf{heap}.x > 0$ is *stable* with respect to the rely increasing. When a write establishes a new stable predicate over the data, strong updates to the reference's predicate (changing the predicate in the type) are sound. (Similarly, when a write invalidates a reference's predicate, a strong update is required, to a new predicate stable over the rely.)

Many verification techniques for imperative programs struggle to verify examples of this kind. Reference immutability and fractional permissions [7, 33] can only allow or outright prevent mutation, not control it. Separation logic cannot concisely specify the counter's intended semantics, only code's behavior. Rely-guarantee and related systems can express the semantics among threads [22, 31, 51], but only coarsely [52] among different program sections. Most program logics can constrain the actions of a function on an argument, but the specification must deal with aliasing, either by giving linear semantics to knowledge of the counter (as in separation logic), or by explicitly treating aliasing (as in more traditional Hoare logics [28]).

With rely-guarantee references, functions are written without concern for aliasing among their arguments. A function cannot be called with unsafe aliasing among arguments: since each alias's guarantee must imply each other's rely, each function explicitly accounts for its possible actions. If two arguments of the same type have conflicting rely/guarantee conditions, they cannot be aliased.

2.1 Subtleties of Rely-Guarantee References

While the intuition behind rely-guarantee references is straightforward, this section overviews some more subtle features of our system that avoid problems.

Non-duplicable References A reference may be freely duplicated if its guarantee implies its own rely, as with the monotonically increasing counter. But consider a reference

$$y : \text{ref}\{\texttt{nat} \mid \textsf{any}\}[\textsf{decreasing},\textsf{increasing}]$$

Making an alias to y where the the alias has the same type as y violates soundness, because the guarantee of the duplicate does not imply the rely of y! Instead, aliasing y requires splitting it into two aliases with weaker rely/guarantee conditions. We support such splitting via a novel substructural resource semantics (Section 4).

Reference to References We need a reference's guarantee to restrict all actions performed using that reference, which must include actions performed via references acquired by dereferencing the first reference. Otherwise, reading a reference out of the heap and writing through it could violate the original reference's guarantee, violating the "capability to perform effects in the guarantee" intuition, and potentially invalidating a predicate. So reading from the heap must somehow transform the type of the referent to restrict resulting references. Reference immutability systems can give a simple binary function on permissions [27], to capture the transitive meaning of qualifiers. For example, dereferencing a `readable` reference to a `writable` reference returns a `readable` reference (assuming a *deep* interpretation of reference immutability, where permissions apply transitively). By contrast, our type system combines *arbitrary* relations (Section 3.2). Furthermore, if one reference points to another, how should the rely of the "inner" reference be related to the outer

---

[1] Because the type nat of natural numbers contains no negative numbers.

one? It is unsound if it permits more interference than the outer rely, so our type system prevents this.[2]

Footprint How much of the heap may a rely-guarantee reference's predicate or conditions mention? It is not productive or sound to let a reference constrain unrelated heap data: letting a reference arbitrarily constrain the heap could lead to allocating a new heap cell whose rely is not implied by existing references. The type system restricts the expressiveness of these predicates to ensure sound and tractable reasoning: predicates and relations may depend only on the heap reachable from the reference.

Cycles Many useful data structures contain cycles, so we wish to reason effectively about them. The solution turns out to be simple (propositions describing cycles require finite proofs, and recursion based on heap structure is not permitted in predicates), but was not immediately obvious to us.

3. Examples

We present examples using rely-guarantee references to verify programs. The examples are small, but highlight distinct capabilities of rely-guarantee references. Rather than writing examples in our core language RGREF (Section 4), we present them using a slight simplification of our shallow embedding in COQ [15]. The embedding is largely in the style of YNOT [12, 40], using axioms for heap interactions.

Reading COQ ***Source*** COQ's language for defining functions and types is based strongly on ML, though many keywords are different: `Definition` and `Fixpoint` for non- and recursive definitions, `Inductive` for defining inductive variant datatypes by specifying constructors. `Parameter` declares assumptions, external functions, or abstract elements in a module signature. Functions and parameterized type definitions can put some arguments in braces rather than parentheses; these arguments are implicit, and inferred when possible from later arguments. Another notable syntactic change from ML is that = is an operator for *propositional* equality, not a boolean decision procedure for structural equality. Therefore, := is often used where ML would use = in definitions. The set of types is much richer than ML, not only due to dependent types but because there are universes (types of types): Prop is the type of propositional types (erasable during extraction, such as proof terms with conjunction, implications, etc.), and Set is the type of normal (computationally relevant) data types. COQ also includes a notation feature that allows users to extend the grammar with additional parsing rules, allowing programs to use syntax closer to mathematical definitions (such as ref$\{T \mid P\}[R, G]$). Our notation uses ML's dereference operator (!) and uses $r \leftarrow e$ for writing e to the location referenced by r. We introduce further notations as they arise.

The PROGRAM extension [48] (used via definitions prefixed with `Program`) allows the omission of explicit proof terms in programs. Omitted terms are either solved automatically via a (customizable) proof search tactic, or set aside for subsequent manual interactive solving, improving readability.

3.1 Monotonic Counter

Consider again our running example of a monotonically increasing counter. Generally, rely and guarantee conditions must be defined over pre- and post-heaps as well as values, to describe the interference they tolerate on reachable substructures. For a simple counter, there is no other reachable data, so the pre- and post-heaps may be ignored. Thus the relation for increasing over time is defined as:

```
Definition increasing (n n':nat) (h h':heap) : Prop :=
   n' ≥ n.
```

Code to allocate a counter is straightforward:

```
Program Definition mkCounter (_:unit)
  : ref{nat|any}[increasing,increasing] :=
   alloc 0.
```

The allocation function `mkCounter` generates well-formedness proof obligations for the resulting type:

- that any is stable with respect to increasing

- that any and increasing are precise: they access only the (empty) heap segment reachable from the natural number they apply to

- that any is true of 0

In our prototype implementation (Section 5), most of these obligations are proven automatically by lightly-guided automatic proof search. Type errors for actions that fall outside the guarantee (or ill-formed rely/guarantee relations, or predicates that are not precise, etc.) manifest as unsolvable proof obligations.

Using a monotonic counter is also straightforward:

```
Program Definition example (_:unit) :=
let x = mkCounter () in x ← !x + 1;
```

An assignment typechecks only if the change implied by the write is permitted by the reference's guarantee relation, for any pre-heap and pre-value satisfying the reference's refinement. In this case, the assignment generates a proof obligation of the form

$$\forall x, h. \, \text{any} \, (!x) \, h \implies \text{increasing} \, (!x) \, (!x+1) \, h \, h[x \mapsto h[x]+1].$$

which is easily solved, with little effort beyond what is required to verify a pure-functional increment function (see Section 5.3). Each read also generates a proof obligation that the guarantee `increasing` is "reflexive": it allows a reference to be used without modifying the heap ($\forall n, h. \, \text{increasing} \, n \, n \, h \, h$). By contrast, an empty guarantee relation would disallow using a reference.

The monotonically increasing counter was proposed by Pilkiewicz and Pottier [46] as a challenging goal for program verification. Unlike their solution and another in fictional separation logic [30], we can state the monotonicity property plainly and require no abstraction to prevent unchecked interference. On the other hand, their solutions verify that the increment occurs, while ours ensures that increment is the only permitted action.

3.2 Monotonic List

We can define a monotonically growing (prepend-only) list, either using a mutable reference to a pure-functional list, or using mutable nodes. The former approach is similar to the monotonic counter, so to show the power of rely-guarantee references for recursive data structures, Figure 1 shows the latter.[3]

We first define `hpred` and `hrel`, type-level functions that allow shorter type declarations. We use them throughout this paper. Next we define a linked list structure, with restricted references to the tail. `list_imm` constrains the tail to be immutable. For the reader unfamiliar with COQ, `list_imm` is a GADT [54] constructing a proposition on different constructions of lists. The first constructor declares that an empty list must remain empty, regardless of heaps. The second constructor accepts a `nat`, a tail, two heaps, and a proof that `list_imm` holds over the tail of the list in those heaps, returning a declaration that a cons cell must remain constant. The immutability requirement is not essential to this example (we could, for example,

---

[2] This is actually a design decision that simplifies checking stability. An alternative design could check a predicate for stability over any change permitted by any reference reachable from the predicate's target referent.

[3] The most natural design uses mutual inductive types where one indexes the other, which we assume here. COQ does not support this, so we use an encoding, discussed further in Section 5 and Appendix A.

```
Definition hpred (A:Set) := A -> heap -> Prop.
Definition hrel (A:Set) := A -> A -> heap -> heap -> Prop.
Inductive list : Set :=
  | nil : list
  | cons : forall (n:nat),
                  ref{list|any}[list_imm,list_imm] -> list
with (* list tails are immutable (_imm) *)
  list_imm : list -> list -> heap -> heap -> Prop :=
  | imm_nil : forall h h', list_imm nil nil h h'
  | imm_cons: forall n tl h h',
        list_imm h[tl] h'[tl] h h' ->
          list_imm (cons n tl) (cons n tl) h h'.
(* Convenient allocation functions *)
Program Definition Nil {P:hpred list}
  : ref{list|P}[list_imm,list_imm] := alloc nil.
Program Definition Cons {P P':hpred list}
    (n:nat)(tl:ref{list|P}[list_imm,list_imm])
  : ref{list|P'∩(λ l h=>l=cons n tl)}[list_imm,list_imm]
  := alloc (cons n tl).
(* A prepend-only list container *)
Record list_container (P:hpred list) :=
  mkList { head : ref{list|P}[list_imm,list_imm] }.
Inductive prepend : hrel (list_container P) :=
  | prepended : forall c c' h h' n,
      h'[head c']=cons n (head c) -> prepend c c' h h'
  | prepend_nop : forall c h h', prepend c c h h'.
Program Definition newList (P:hpred list)
  : ref{list_container|any}[prepend,prepend] :=
    let x = Nil in alloc (mkList P x).
Program Definition doPrepend {P:hpred list}(n:nat)
  (l:ref{list_container|any}[prepend,prepend]) : unit :=
    let x = Cons n (head l) in
        l ← mkList P x.
```

Figure 1. RGREF code for a prepend-only linked list.

permit the numbers to change but require the length to increase), but is included for completeness. We then define convenient helper functions for heap-allocating nil or cons cells.

We enforce the prepend-only behavior through reference to a `list_container` structure, which holds a reference to a list parameterized by some predicate. The `prepend` relation on `list_container`s allows prepending and no-ops (required for reading the reference). Finally we have helper functions to allocate a new list and to prepend the list with a new cell. `prepend` is essentially the specification of what `doPrepend` is permitted to do with the list.

Note the predicate conjunction (∩) in the return type of `Cons`. This, along with the predicate conversion rule, is how flow-sensitive assumptions can be handled (notice that the equality is stable with respect to `list_imm`). This is important in `doPrepend`, where information from the result of one write (inside `Cons`) must be carried into another (the assignment through `l`), because it is otherwise unavailable in the expression stored.

Not shown in Figure 1 are implicit obligations such as $\forall h. P$ nil h in `Nil`. Other such obligations include:

- $\forall tl, h. P\, tl\, h \implies P'\,(\text{cons } n\, tl)\, h$ in `Cons`

- That `prepend` permits the write in `doPrepend` (under the trivial assumption that `any` holds of the initial list container and heap). This obligation requires a richer type for the allocation result, because `mkList` must know `x` is a cons cell whose tail is the old list. This information is not available locally (within the write statement itself). Other systems propagate hypotheses separately, but we only need to track variables: the required equality is present in `x`'s predicate because of `Cons`'s return type.

- The stability and precision properties that must hold for various predicates and `list_imm`.

	▷	immutable	=	immutable
immutable	▷	_	=	immutable
readable	▷	_	=	readable
writable	▷	q	=	q

Figure 2. Combining reference immutability permissions, from [27]. Using a p reference to read a $q\ T$ field produces a $(p \triangleright q)\ T$.

- Propagations of these obligations to indirect polymorphic callers, such as `newList` and `prepend`.

Also omitted in Figure 1 are obligations related to *folding* and *containment*. Folding is the restriction of read result types to ensure that for any reference r with guarantee G, references produced via reads of r do not allow actions exceeding those permitted by G on r's referent. This ensures actions via a reference read from inside a data structure cannot invalidate predicates over the whole structure. Containment is a check that the rely R for a reference r captures all interference allowed by the interference summaries of references reachable from r. This ensures that any predicate preserved by R is also preserved by actions on aliases to internal structures.

Both operations require projecting a given relation component-wise onto a datatype's members. For our prepend-only list, projecting `prepend` is trivial (it does not constrain the list cells' values), and the result of projecting `list_imm` is logically equivalent to `list_imm` itself.

3.3 Reference Immutability as a Special Case

Reference immutability [27, 50, 55, 56] adds permissions (type qualifiers) to references to permit or disallow side effects through a particular reference. Multiple aliases at different permissions may coexist if compatible: for example, there may be write-permitting and read-only aliases to an object. We can define the permissions of reference immutability like this:

```
Definition havoc {A:Set} : hrel A :=
    fun x => fun x' => fun h => fun h' => True.
Definition readable (T:Set) := ref{T|any}[havoc,≈].
Definition writable (T:Set) := ref{T|any}[havoc,havoc].
Definition immutable (T:Set) := ref{T|any}[≈,≈].
```

Our definitions encode the standard semantics for reference immutability qualifiers: only `immutable` assumes limited interference via other aliases, and `readable` and `immutable` disallow mutation through a reference. Restrictions on aliasing among reference immutability permissions are reflected in the rely and guarantee relations: no heap cell may have writable and immutable aliases simultaneously, as the guarantee of the writable reference (`havoc`) is not a subrelation of the immutable rely (≈). The "containment" requirement (Section 3.2) for rely conditions on nested datatypes is satisfied by the rely for readable and writable, and for immutable the rely prevents it from (transitively) referencing mutable data.

Reading through one of these references requires considering how the rely and guarantee affect the read's result. In reference immutability, result types are adapted using a simple binary function on permissions (Figure 2). Our rely-guarantee reference type system must combine arbitrary relations (folding), using the type of the referent. Intuitively, folding is projection of the reference's guarantee onto the referent type. Any projection of `havoc` and ≈ correspond with the simplified version in Figure 2. Projecting `havoc` is equivalent to reading through a `writable` reference, which simply produces the inner type. Projecting ≈ is equivalent to the weakening that occurs when reading through `readable` or `immutable` references.

We can also give a reference immutability system with limited dependent types by a small extension:

```
Definition refined (T:Set) (P:hpred T) := ref{T|P}[≈,≈].
```

At first glance this is weaker than proposed systems that let mutable data's type depend on arbitrary immutable data, because we require any reference predicate to access only heap state reachable from its referent. At the cost of some space the referent could maintain its own extra reference to relevant immutable data. Careful code extraction work can improve the space overhead in executables.

Another benefit of implementing reference immutability via rely-guarantee references is interoperability between reference immutability and richer rely-guarantee references. For example, a function accepting a `readable` reference to a natural number can be passed a read-only monotonically-increasing counter from Section 3.1. This offers a natural path for gradually adding stronger verification guarantees to code using reference immutability (which itself is a gradual refinement of unrestricted references [27]).

3.4 RCC/Java with Reference Immutability

The core of RCC/Java [24] is also implementable as a small library using our COQ DSL, and we present a translation of an early version [23]. The key idea in these type systems and related systems is to parameterize the type of a reference by the identity of a particular lock. The type system tracks the set of held locks and permits reads and writes through a reference only when the reference's lock parameter is statically known to be held.

A COQ module wraps standard acquire and release primitives and exposes a new reference type that quantifies over a lock. The RCC reference type can be abstracted with a module signature, but can be concretely represented by a RGREF ref type:

```
(* Signature *)
Parameter rccref : forall (A:Set),
   hpred A -> hrel A -> hrel A -> lock -> Set.
...
(* Implementation *)
Definition rccref A P R G (l:lock) := ref A P R G.
```

The module then exposes its own read and write primitives, and external ways of proving goals like guarantee satisfaction that do not expose the internal `rccref` representation. This mostly consists of re-exporting existing axioms using new names. Then an explicit lock witness (since the type system is not specialized to track lock witnesses) can be abstracted using:

```
Parameter lockwitness : lock -> Set.
Parameter locked : forall {l:lock}, hpred (lockwitness l).
Parameter unlocked : forall {l:lock}, hpred (lockwitness l).
```

The acquire and release operations must respectively produce and consume a witness that a lock is held, that permits release. Using a binary operator `-->` on predicates that produces a relation allowing changes from states where the first predicate holds to states where the second holds (a limited encoding of protocols), a witness may have type ref{lockwitness | locked}[empty, locked-->unlocked]. Using an empty rely implies uniqueness, and requiring such a witness to release the lock prevents splitting the witness, which would requiring weakening the rely of both resulting references. The read and write for `rccref`s would need to also require some witness that the lock was held:

```
Program Definition rcc_read ...{l:lock}
   (w:ref{lockwitness l | locked}[empty,locked-->unlocked])
   (r:rccref A P R G l)... := (!r,w).(* return the witness *)
```

This encoding of RCC/Java using dependent types is not novel, but note the rely, guarantee, and predicate of the underlying reference are exposed, yielding the first combination of race-free lock acquisition and reference immutability we are aware of, in addition to exposing the full power of rely-guarantee references over lock-guarded data.

Expressions	e	$::=$	x	Variable
		$\mid$	$e\ e$	Application
		$\mid$	$!e$	Dereference
		$\mid$	$e[e]$	Heap Select
		$\mid$	alloc e	Allocation
		$\mid$	$(\lambda x : \tau . e)$	Pure Function
		$\mid$	$(\lambda_{\mathcal{M}}(x : \tau) . e)$	Procedure
		$\mid$	τ	Types
		$\mid$	$x \leftarrow e$	Store
		$\mid$	$\text{swap}(x, e)$	Atomic swap
Types	τ	$::=$	$=$	Propositional equality
		$\mid$	Prop	Propositions
		$\mid$	$\text{ref}\{\tau \mid e\}[e, e]$	Reference
		$\mid$	Type	Type of Prop
		$\mid$	$\Pi x : \tau \rightarrow \tau'$	Dep. Product (pure)
		$\mid$	$\tau \xrightarrow{\mathcal{M}} \tau'$	Impure Function
		$\mid$	heap	Heaps
		$\mid$	e	Expressions

Figure 3. Syntax, omitting booleans (b : bool), unary natural numbers (n : nat), unit, pairs, propositional True and False, and standard recursors. The expression/type division is presentational; the language is a single syntactic category of PTS-style [2] pseudoterms.

4. A Type System for Rely-Guarantee References

Figure 3 gives the syntax for a core language RGREF with rely-guarantee references. The expressions combine features from the ML-family (e.g., references) and dependently typed languages (e.g., dependent product), specifically from the Calculus of Constructions [2, 16]. We include a few basic datatypes (natural numbers, booleans, pairs, unit), a type for propositional equality, and their standard recursors [29]. We also distinguish effectful functions through the term former $\lambda_{\mathcal{M}}$ and the effectful non-dependent function type former $\tau \xrightarrow{\mathcal{M}} \tau'$ ($\mathcal{M}$ for mutation).

The language supports reasoning about heaps: not only is there a standard form of dereference, but there is a term for dereferencing a reference in a particular heap, used to specify predicates and rely/guarantee relations using the propositions-as-types principle. The language is designed to use propositions-as-types to specify predicates and relations, and to use the pure sublanguage as a computational language amenable to rich reasoning, but to use external means for discharging obligations like a write satisfying a guarantee. The presence of current-heap-dereference makes the pure term language itself unsound as a logic, internally offering assurances similar to Cayenne [1]. In general, the language for predicates and relations can be distinct from the term language, and the term language does not require advanced types; our design is motivated by our implementation as a COQ DSL (Section 5).

Figure 4 presents the primary typing rules for the core language. There are two key judgments: $\Gamma \vdash e : \tau$ for pure terms (useful for proofs), and $\Gamma \vdash e : \tau \Rightarrow \Gamma'$ for impure and substructural computation. Those pure rules omitted here (recursors for the assumed inductive types, typing the primitive types in Prop, etc.) are standard for a pure type system.

The imperative typing judgment $\Gamma \vdash e : \tau \Rightarrow \Gamma'$ is flow-sensitive to allow reasoning about when references are duplicated. Crucially, it allows reasoning about when a reference must not be duplicated because its guarantee does not imply its own rely. For this reason, we have two impure variable rules: one consumes the variable (V-∅), and the other uses an auxiliary relation $\Gamma \vdash \tau \prec \tau \divideontimes \tau$ to split a type (V-∗). Primitive types (nat, bool, unit) freely split into two copies of themselves. We require that any variable captured by a closure

has a self-splitting type (Π-I and FUN), and thus functions may be duplicated freely. We require that only values of self-splitting types are captured by dependent type constructors (Π-F and the not-shown propositional equality rule), so types are also self-splitting. Variables read in pure computation must also be self-splitting (V).

References (and structures that may contain them, like pairs) are the only types with non-trivial splitting. Reference types split into reference types that may coexist (each guarantee implies the other rely, both relies are no stronger than the original rely, stable predicates, etc.), and pairs split into pairs of the component-wise split results. For example, the problematic reference from Section 2.1 has non-trivial splitting behavior, mediated by REF-✳:

$$y : \mathsf{ref}\{\mathtt{nat} \mid \mathsf{any}\}[\mathsf{decreasing}, \mathsf{increasing}]$$

Splitting this reference requires consuming the original and producing two "weaker" references: each guarantee may permit at most what the original guarantee allowed, and each rely must assume at least as much interference as the original. For example,

$$\varepsilon \vdash \mathsf{ref}\{\mathtt{nat} \mid \mathsf{any}\}[\mathsf{decreasing}, \mathsf{increasing}]$$
$$\prec \mathsf{ref}\{\mathtt{nat} \mid \mathsf{any}\}[\mathsf{havoc}, \mathsf{increasing}]$$
$$✳ \mathsf{ref}\{\mathtt{nat} \mid \mathsf{any}\}[\mathsf{havoc}, \mathsf{increasing}]$$

The natural use of simply duplicating a reference whose guarantee implies its own rely (as in the monotonic counter) is a degenerate case of the very general rule REF-✳.

The conversion relation $\Gamma \vdash \tau \rightsquigarrow \tau$ is a directed call-by-value β-conversion (so for example β-reduction is not used with arguments containing dereferences) plus reducing abstractions whose bound variable is free in the result, and what amounts to subtyping by converting predicates and relations to weaker versions: P-⇒ weakens the predicate; R-⊂ assumes more interference may occur; and G-⊂ sacrifices some permissions; P and R changes may affect stability.

Mutation The most interesting rules are those for mutation, particularly for writing to the heap (WRITE). This rule requires (beyond basic type safety) that the effects fall within the guarantee, assuming the reference's predicate holds in the current heap.[4] It also allows the option of a strong update to the reference's predicate, if the change establishes some new stable predicate. For example,

$$x : \mathsf{ref}\{\mathtt{nat} \mid \lambda x. \lambda h. x = 3\}[\mathsf{empty}, \mathsf{havoc}] \vdash$$
$$x \leftarrow 4 : \mathtt{unit}$$
$$\Rightarrow x : \mathsf{ref}\{\mathtt{nat} \mid \lambda x. \lambda h. x = 4\}[\mathsf{empty}, \mathsf{havoc}]$$

Thus RGREF naturally supports strong updates on unique references as a degenerate case. The atomic swap operation (which permits modifying substructural fields) leverages the heap-write rule's premises. Allocation simply requires a well-formed type as a result and establishing the predicate over the value in any heap. The imperative part of the language also includes non-dependent function types, application, and the use of pure expressions.

Dereference uses the *relation folding* function $[R, G] \gg \tau$ (Figure 5) to reason about the rely and guarantee in result types. It has no effect for non-reference types. For types containing references, the result type is rewritten by intersecting the projection of the guarantee onto each component with the stated guarantee for the component itself. This can cause some precision loss. The effect of folding on the rely has no impact: because the rely for any well-formed reference type has to contain / admit the effects allowed by the rely of any reachable reference type, the intersection on the rely component would produce a relation semantically equivalent to the syntactically present rely on the inner reference (the same type of relation projection is used to check rely containment as is used in

folding the guarantee). In general, relation folding and checking containment are straightforward for types whose members are always reflected in a type index (e.g., pairs, references). Folding for richer types, such as full inductive types [45] is left to future work. The Deref rule also checks that the source type is self-splitting. This ensures that the (possibly weaker) guarantee of the result implies the original location's rely, and the original value's guarantee(s) will imply the (unaltered) rely relations of the result.

The rule for typing reference types themselves (WF-REF) imposes several additional requirements on predicates and relations. First, the predicate P must be stable with respect to R. Second, the relations and predicates must be *precise*: all depend only on heap state reachable from the referent. This prevents code from rendering the system unsound by allocating a new cell whose rely condition requires the heap to be invariant: that rely would be undesirable if enforced as it prevents all mutation and allocation, but unsound if ignored because all predicates are stable over such a rely. Finally, the rely must be *closed* (contained): any changes permitted by relies of references reachable from the referent are also permitted by the checked rely. This ensures that checking stable P R is sufficient to ensure P is not violated (otherwise P could depend on other references reachable from the referent, whose rely relations might permit P to be invalidated). Figure 5 defines these notions precisely. WF-REF also requires as a side condition that P, R, and G are free of dereference expressions ($!e$), since implicitly heap-dependent predicates are not sensible.

We foresee no technical difficulty in building the system directly atop stronger calculi such as the Calculus of Inductive Constructions [4] (CIC) beyond richer treatment of folding for the full spectrum of inductive types (Section 5.2). There are a few essential qualities required for soundness. First, effectful terms and abstractions are encapsulated in a separate judgment (which corresponds to a monadic treatment of effects in translation to a pure system). Any term in the pure fragment must be reflexively splittable, to avoid introducing resource semantics into the pure sublanguage. Captured variables have reflexively splittable types ($\Gamma \vdash \tau \prec \tau ✳ \tau$). In principle this could be weakened, but we prefer the simplicity of allowing function terms to be arbitrarily duplicated. We build upon CC for simpler presentation.

4.1 Soundness Sketch

Soundness follows a preservation-like structure. Evaluation must preserve a couple invariants beyond standard heap soundness:

- For each reference $r : \mathsf{ref}\{T \mid P\}[R, G]$ in the stack, heap, or expression under reduction, there exists a proof of $P\ (h[r])\ h$ for the current heap h.

- For each pair of references $p : \mathsf{ref}\{T \mid P\}[R, G]$ and $q : \mathsf{ref}\{T \mid P'\}[R', G']$ in the stack, heap, or expression under reduction, if p and q alias (point to the same heap cell) then $\emptyset \subset G' \subseteq R$ and $\emptyset \subset G \subseteq R'$.

Initially there are no references, so these hold trivially. On allocation, the predicate for the new object is true in all heaps by inversion on the allocation typing, so the result is immediate. On any action through a reference, the type system ensures the action falls within the guarantee. The action either preserves the predicate or produces a new (stable) one, and the proof for that reference's new refinement is easy to construct from the typing derivation results. For aliases, the used reference's guarantee must imply the rely of any alias, and the alias's predicate is stable over its rely, and therefore preserved by the action within the used reference's guarantee. For other references, the action will fall within its rely (by containment) and thus a similar use of a stability proof suffices, or the changed cell is not reachable from the reference in question, in which case the predicate is preserved by precision.

---

[4] The reflexivity goals DEREF generates could also assume the predicate (i.e., reflexive on states satisfying the predicate), but we haven't needed this.

To make this proof easier, the dynamic semantics for RGREF have a few quirks:

- Variable binding occurs by stack usage in impure code (making it easier to prove substructural behavior), and by substitution in pure code (justified by the requirement that any variable used in pure code is self-splitting).

- Variables captured by either a pure or impure closure are required to be self-splitting, but it would be unsound to allow those captured uses to read the stack location if the closure were executed after impure code split the stack reference with different permissions. To prevent this, closures are only values when they have no free variables, and there are reduction rules that perform the reads for free variables.

- References are represented as "fat pointers" of the underlying pointer, rely, guarantee, and predicate

- Casts/subsumption and substructural consumption/splitting of variables are explicit — including explicit mention of result predicates and relations — so the semantics can appropriately modify underlying reference values. The translation from the source language presented here to the more explicit language is straightforward by induction on the typing derivation.

- We assume the heap behaves according to the axiomatization used in rely/guarantee conditions and predicates.

- Folding has a runtime representation. Reads produce the heap value wrapped in a deferred folding construct which lazily pushes the guarantee restriction through a data structure as components are evaluated. This is necessary for dependent type constructors like propositional equality, where types appear in values of the type.

- The semantics occasionally reduce multiple dereferences of the same location simultaneously, if they are constrained by types to be definitionally equal. We have yet to conceive of desirable *computational* code that observes this behavior; see Sections C and F.

5. Implementation

To understand RGREF's effect on data structure design and the effort required for verification, we have implemented RGREF as a shallow embedding in COQ, and used it to implement Section 3's examples. This includes implementing reference immutability, meaning our RGREF implementation can be used to write programs using reference immutability, and to gradually refine parts of those programs to use more fine-grained rely and guarantee conditions and predicates. Overall, we found that RGREF required careful choice of type refinements, but did not affect algorithm design and had reasonable proof burden (commensurate with the complexity of the code verified).

The implementation is done largely in the style of YNOT [12, 40], with axioms for heap interaction, and using COQ's notation facilities to elaborate source terms to COQ terms with proof holes, which are then elaborated and semi-automatically solved by Sozeau's PROGRAM extension [48]. Each structure typically requires its own slightly customized PROGRAM tactic for effectively solving most goals, but we find that following the tactic development style Chlipala recommends [11] tends to work well, as each module typically handles its own family of predicates and relations. Proofs involving heaps are carried out using a small set of axioms reflecting invariants maintained by the semantics. The most useful axiom is $\mathsf{heap\_lookup} : \forall h, A, P, R, G. \forall r : \mathsf{ref}\{A \mid P\}[R, G]. P\,(h[r])\,h$ which means that in any heap, the type system ensures there is a proof of the refinement for every valid reference. The implementation also relaxes some restrictions present in the formal language, such as allowing values to be projected into predicates (as in Section

3.2); predicates, rely and guarantee relations must simply abstain from dereference.

We made compromises to fit into COQ. Notably, COQ lacks support for mutual inductive types where one indexes the other (e.g., a datatype simultaneously defined with an inductive predicate on that type). Our implementation adapts a standard encoding [9] of induction-recursion [21] outlined in Appendix A to support examples like the list in Section 3.2. Any use of this encoding somewhat complicates generated proof obligations and data structure designs. Thus our current implementation is best-suited to "functional-first" designs that make only light use of references, as is common in OCaml, Scala, and F# code. We stress that this is a limitation of our implementation by embedding in COQ, not a fundamental limitation of rely-guarantee references.

Our implementation focuses on self-splitting types; not all primitives for handling substructural data monadically have been implemented (or necessary) yet.

To use COQ's rich support for inductive types, we require trusted user-provided definitions of relation folding, immediate reachability (without heap access) of references from a pure datatype, and relation containment. These are provided as typeclass instances. The definitions are fairly mechanical, and can be synthesized automatically for simple types if we extend COQ's support for inductive datatypes.

We also move some proof obligations such as stability, precision, and containment from type formation to allocation. This allows the definition of functions over ill-formed types, but such functions are unusable: only well-formed types may actually be constructed. This avoids some redundant proof obligations.

Our implementation is publicly available at
https://github.com/csgordon/rgref/.

5.1 Proving Obligations with Dependent Types

RGREF contains as a sublanguage the full Calculus of Constructions (CC). Specifically, it contains a full Pure Type System (PTS) with sorts $\mathcal{S} = \{\mathsf{Prop}, \mathsf{Type}\}$, axioms $\mathcal{A} = \{\mathsf{Prop} : \mathsf{Type}\}$, and product formation rules $\mathcal{R} = \{(s_1, s_2) \mid s_1, s_2 \in \mathcal{S}\}$ as formulated by Barendregt [2]. This sublanguage is part of the pure ($\Gamma \vdash e : \tau$) subset of RGREF. Thus the language is amenable to embedding directly into CC with a few extensions (natural numbers, etc.) and RGREF-specific axioms.

Theorem 1 (CC Embedding). *There are mutually recursive transformation functions $\mathcal{P}[\![-]\!]$ and $I[\![-]\!]$ from a version of RGREF types, terms, and contexts with explicit splitting, conversion, and full type annotations to CC types, terms and contexts such that*

- *If $\Gamma \vdash_{\mathrm{RGREF}} e : \tau$ then $\mathcal{P}[\![\Gamma]\!] \vdash_{\mathrm{CC}} \mathcal{P}[\![e]\!] : \mathcal{P}[\![\tau]\!]$, and*
- *If $\Gamma \vdash_{\mathrm{RGREF}} e : \tau \Rightarrow \Gamma'$ then $I[\![\Gamma]\!] \vdash_{\mathrm{CC}} I[\![e]\!]^{\Gamma} : I[\![\tau]\!]^{\Gamma}$*

The key difference between the two cases is the resulting types: translations of the impure judgments generally produce monadic types.

Proof. By simultaneous induction on the RGREF derivations. See Appendix B. □

This provides us with an approach to proving predicate and guarantee obligations in a way well-integrated with the source language, justifying the use of proof terms in RGREF's implementation. This also allows straightforward translation of proof goals from our typing derivations into COQ, where we can use tactic-based theorem proving to solve proof obligations.

The only subtleties arise from the fact that our embedding treats dereference as an uninterpreted function, allowing two potential inconsistencies. First, we permit recursion through the store, so applying a function read (via dereference) from the heap could

result in an infinite loop; by treating dereference as an uninterpreted function in the embedding, this potential recursion is lost. The prototype may accept proofs about non-terminating terms. Second, there is a potential to equate dereference expressions that will be evaluated in different heaps (e.g., when a returned pure closure dereferences some reference). Our implementation only uses full equational reasoning for cases where all dereferences occur in the same heap, and otherwise abstracts terms with only their type.

5.2 Data Structure Design

We were able to use natural data structures for the examples in Section 3, but the types require careful consideration for propagating information through data flow. For example, the return type of `Cons` in Figure 1 must carry the additional refinement that the reference points to a cons cell whose tail is the `tl` argument, or the obligation to prove that the write prepends a cell in `doPrepend` is unprovable.

From Simple Types to Inductive Types For the types whose splitting, folding, and containment we have examined formally (first-order data types, pairs, and references), the structure of the types is simple enough to provide a straightforward structural projection for each type. Much imperative code (e.g. in C, Java, etc.) heap-allocates similarly-simple structures. We have not worked out the theory for full inductive types. For types whose constructor arguments are not reflected as type indices, splitting and the like depend on the values passed to constructors, complicating the definitions for splitting because they then depend not only on type indices, but the actual value being potentially-split. This is an issue even though we expect to only require support for small inductive types. For self-referential datatypes, such as the list, we have only used guarantees for which folding is idempotent. In general, folding a restricted guarantee when dereferencing an datatype defined with concrete relations on "recursive" references is not expressible in the current system; if the guarantee on the recursive member changes, the result may not match a constructor of the type! Supporting this would require some sort of datatype-generic support, or a hybrid dereference-and-pattern-match to avoid directly representing not-quite-typed read results. We leave full support for inductive types to future work.

5.3 Proof Burden

RGREF imposes proof obligations for precision, folding, containment, and guarantee satisfaction. For verifying the examples in Sections 3.1, 3.2, and 3.3, the proof burden is not substantially different from verifying the analogous pure-functional version. This section will call out which parts were particularly straightforward, as well as the few challenging aspects.

Precision obligations are typically easy to prove when they are true, as are the reflexive relation goals generated when references are used for reading: most are discharged by a simple induction, use of constructors for the relevant relation, and/or first order reasoning (e.g., COQ's `firstorder` tactic). When the goals are not true (e.g., a relation or predicate is not precise), the goal is simply not provable, and it is up to the developer to recognize this. In this respect, RGREF is similar to verifying a functional program in COQ: a developer can waste time on unsolvable (false) proof goals.

In cases where relation folding is a no-op (e.g., $[R, G] \gg \tau = \tau$ as in the monotonic counter) folding goals are a simple matter of reduction and basic equality. In cases where relation folding is not a no-op (e.g., for references to references where the outer reference's guarantee bounds effects on other reachable objects), the folding obligations' complexity depends on the relations involved.

The most difficult proof obligations generated are those checking that heap writes satisfy the guarantee.[5] This is partly because these

goals sometimes require reasoning about reads from an updated heap, in particular proving non-aliasing between references to different base types. Some goals are also complicated by non-identity folding results in types. In the formal system, we abstract away the mechanism for checking guarantee satisfaction through a denotational semantics for writes. Our implementation uses COQ's notation facilities for a sort of "punning," to duplicate expressions into two contexts with different semantics for dereference.

The normal program's use of the dereference expression chooses the appropriate relation folding type class instance, while the duplicated version used to check guarantee satisfaction is placed inside a context where a no-op fold instance overrides all others. This way the guarantee and predicates, which are specified as predicates over a type A, can be applied directly to $!x$ at type A in the proofs instead of at a weakened type. The disadvantage of this approach is that the expression duplication also duplicates proof goals. Many of the smaller goals are automatically discharged when using COQ's PROGRAM extension, but because the generated goals are formed in slightly different contexts, the solved lemmas' proof terms have different arities, and are therefore not interchangeable in equalities. We encountered this twice, and solved it by using the proof irrelevance axiom on applications of the equivalent lemmas.

Because stability, reflexivity, and satisfaction results tend to be reused within a module and by clients, it should be considered proper practice for modules exporting a given API to also export most goals proven internally about properties of rely and guarantee relations, and predicates, as lemmas registered in a module-specific hint database. This is best practice for verifying purely functional programs in COQ as well; in general most of the useful habits in verifying functional programs can be reused in our implementation.

6. Future Work: Extensions and Adaptations

The type system we present in this paper has a few technical limitations beyond the implementation limits described in Section 5.

The system we present here is modeled on the *deep* interpretation of reference immutability (the standard interpretation), where the permission of one reference constrains the permissions on read results through that reference (this can be seen best in the rely-guarantee folding in the dereference type rule). This is contrasted against the *shallow* interpretation of reference immutability, where a permission affects only the actions on a particular heap cell (so a writable reference could be obtained by reading through an immutable cell). Our formalism does admit uses of shallow relations and predicates: these propositions simply ignore the heaps, so containment and folding are trivial.

Our type system could be converted to a fully shallow model by:

- removing the rely-guarantee folding,

- removing the restrictions that a rely condition must account for the rely conditions of any reference reachable from it, and

- adding tighter footprint restrictions on rely-guarantee conditions and refinement predicates.

For the last point, we mean specifically that the rely, guarantee, and predicate should be insensitive to heap contents outside the heap cell they apply to. The simplest way to accomplish this is to remove the heap arguments from these predicates, making them closed over only the value itself. Note that any rely/guarantee and predicate matching these restrictions is already valid in our system, and folding and containment are semantically no-ops because component-wise projection of these more restricted relations do not impose any new restrictions.

One weakness of our current system is that every individual effectful action must fall within a reference's guarantee. Some oper-

---

[5] Not all are difficult; Section 3.1's guarantee obligation is discharged by automatic proof search with arithmetic hints: `auto with arith`.

ations (for example, splay tree rotation) are difficult or impossible to write this way while satisfying guarantees that preserve interesting refinements on aliases (such as set membership), limiting the properties the system can verify. Other systems based on various forms of object invariant have a notion of *focusing* or *opening* an object for a series of operations that together preserve an invariant, but where intermediate states violate the invariant [3, 5, 47]. Type systems with focus often make the typing judgment modal, with one "unfocused" mode and one where a particular object is focused. We could add such support to our system. This would grant us additional flexibility, and also allow us to subsume much of the expressive power of Militão et al.'s recently proposed rely-guarantee inspired typestate system (see Section 7).

Another limitation is that effectful function types do not summarize strong updates to predicates of references provided as arguments. Such summaries could be useful, and should not be too difficult to add (particularly as the CC translation already uses a monad that preserves more information than the current source language types).

A promising direction for future work is to further explore the resource-like semantics of splitting references by rely and guarantee, and allowing *recovery* [27] of "stronger" references from the results of splitting, either based on combining provably equal references or by controlling scope. This would help control the gradual loss of precision in the current system when non-duplicable (non-self-splitting) references are repeatedly split to less precise relations.

7. Related Work

The most closely related work falls into three categories: restricting mutation on a per-reference basis, techniques for reasoning about interference among threads (which can be adapted to interference among aliases), and dependent types for imperative languages.

Alias-based Mutation Control   Many techniques exist to control side effects by restricting actions through particular references. Notable examples include reference immutability [27, 50, 55, 56], and the owner-as-modifier interpretation of ownership and universe types [17]. Rely-guarantee references generalize reference immutability permissions (Section 3.3), allowing precise control over what modification is permitted through a given alias, not simply a choice between arbitrary mutation and local immutability. The fact that reference immutability is a special case of rely-guarantee references suggests a natural transition path from reference immutability to stronger verification guarantees. Developers could employ a "pay-as-you-go" model for verification, where a code base first transitions to using reference immutability (which need not be onerous [27]), and then gradually enrich the types for some parts of the system where more assurance is desired.

Typestate approaches typically use reference immutability like access permissions to control sharing of objects in a certain typestate, which is a weak form of refinement [5, 34, 37, 41]. Nistor and Aldrich describe a program-logic style type system [41] using abstract predicates [44] and connectives inspired by separation logic to specify *object propositions*, essentially an enriched typestate much closer to a full predicate logic; we believe object propositions and RGREF have similar expressivity in terms of what predicates they can verify. They pair access permissions with refinements to let aliases share coinciding views of an object's properties, and handle non-atomic updates. However, propositions on aliased objects cannot change over time and all aliases must have identical capabilities (simply preserving the proposition if the object is aliased), while RGREF allows asymmetric permissions and some strong updates to refinements even with aliasing.

Militão and Aldrich present a system that splits objects into substates that each carry their own typestate, and ways to merge particular substates into a typestate of the parent object [34]. Aliases to

an object may exist and allow modification provided each reference is to a different substate, rather than to the full substate.

They extend this to a notion of rely-guarantee [35], where a rely is a single typestate all other aliases to an object assume an object to be in, and a guarantee is a typestate a given alias is expected to produce before other aliases become usable. They ensure that temporarily-ill-formed objects are not used through aliases by adopting a focus-unfocus model, where only one object's fields are accessible at one time and a focused (unpacked) reference must be restored to the guarantee view before unfocusing (re-packing). Focusing on a reference with a rely-guarantee typestate requires a dynamic check for whether the object is in the rely state, or the guarantee state, as it is statically unknown whether the last access was through another reference (leaving the object in the rely) or the focused reference (leaving the object in the guarantee). This gives a programming model more similar to using fractional permissions to make references agree on a sum typestate (which their language has for reasoning about the typestate at the start of a focus block). It also severely constrains operations that access multiple objects. They also include a form of rely-guarantee "narrowing" to allow a more general rely and weaker guarantee typestate (which correspond to super-typestates and sub-typestates respectively). They include a notion of refinement typestate on rely-guarantee typestate, but their refinement is a typestate that is convertible to the guarantee typestate, and appears to add the restriction that the refinement must also be the guarantee of all other aliases (based on examples; the written description of their refinement is quite vague).

Their system has much of the same flavor as our system (both apply rely-guarantee concepts to references) but with less involved specification and correspondingly weaker verification options compared out our approach. Typestate is known for having a fairly good annotation/expressivity ratio. The syntactic overhead of our approach is clear. The semantic gain of our approach is primarily due to our rely and guarantee being arbitrary (computable) binary relations on states (rather than views, which are effectively predicates over a single state). This allows us to express relationships over infinite state spaces, rather than simply membership within a particular partition of the state space. This is why classical rely-guarantee, and derivatives such as RGsep [51] and Local Rely-Guarantee [22] use binary conditions. A simple example that is inexpressible in Militão and Aldrich's system is a condition such as that one reference only grants the holder permission to increase a counter. Rely-guarantee views can only express finitely many states, and enforce transitions between those. One of their examples [36] is a counter where one reference may not bring the counter to 0. Their views include a zero view, and a positive view. The closest they could come to specifying a reference for a monotonically increasing counter is a rely-guarantee view reference whose guarantee was either the view corresponding to a particular interval (increasing to further values would require another reference with a different guarantee typestate) or a view for exceeding a particular lower bound, at which point the reference would grant permission to increase *or decrease* the counter above the lower bound.

Reasoning About Thread Interference   Generally, any technique that can reason about interference among threads can be adapted to reason about interference among references. The most closely related system for reasoning about thread interference is rely-guarantee reasoning [31] described in Section 1.

The original rely-guarantee approach focused on global relations and assertions, hampering modularity. Several adaptations exist to treat interference over disjoint state separately. Vafeiadis and Parkinson integrated rely-guarantee reasoning with separation logic [51], allowing separation of state with linear resource semantics from shared state with interference. Feng later generalized this to add separating conjunction of rely and guarantee conditions [22].

The conditions split into separate relations over separated pieces of shared state. Rely-guarantee references are heavily inspired by these approaches. However, RGREF allows substantial overlap among heap segments.

Dodds et al. adapted standard (non-separating) rely-guarantee reasoning to give resource semantics to rely and guarantee relations as assumptions in a context [20]. This allows the interference on shared state to change over time as permission to modify disjoint parts in particular ways is split, rejoined, and split again differently. This style of strong changes to the rely and guarantee over time could be adapted in a rely-guarantee reference system to allow the natural rely-guarantee reference generalization of the *recovery* technique of Gordon et al. [27], which allows recovering unique (or immutable) references from writable (or readable) references in a flow-sensitive type system given some constraints on the input context to a block of code.

Wickerson et al. [52] apply a modularized rely-guarantee logic to treat (non-)interference in the degenerate case of sequential access to the UNIXv7 memory manager. Related systems [18, 22, 51] could be applied similarly, but to our knowledge haven't. Most have only first-order treatment of interference. Only Concurrent Abstract Predicates [18] can (with some effort) store capabilities into the heap, while RGREF naturally supports this since mutation capabilities are tied to data. Our design closely follows a technique already shown successful in large-scale uses (reference-immutability [27]).

Dependent Types for Imperative Code Many others have worked on integrating dependent types into imperative programming languages. Most take the approach of using refinement types [25, 49] that restrict modification to mutable data, but the refinements themselves may depend only on immutable data. Examples include DML [53], ATS [10], and X10's constrained types [42]. The refinement language is often also restricted to some theory that can be effectively decided by an SMT solver, as in Liquid Types [47]. RGREF allows refinements to depend on mutable heap data, and does not artificially restrict the properties that can be verified (at the cost of requiring manual guidance for proofs).

A notable approach to dependent types in imperative code is Hoare Type Theory (HTT) [38, 39] and its implementation YNOT [12, 40]. HTT uses a monadic Hoare Triple to encode effectful computation. It allows using effectful code in specifications: it decides equality of effectful specification terms by using canonical forms where traditional dependent type systems use β-conversion. This approach could be adapted for rely-guarantee references as well. YNOT implements the core ideas of HTT as a domain specific language embedded in COQ. It supports traditional Hoare logic specifications and, by embedding, separation logic specifications as well. A later version [12] builds a family of targeted proof search tactics that can automatically discharge most separation logic proof goals generated while typechecking YNOT programs. We modeled our implementation after YNOT, and are currently building a library of combinators and transformers in hopes of supporting similarly robust automatic proof discharge.

HTT (and separation logic in general) support proving functional correctness rather than the somewhat weaker safety properties verified by rely-guarantee references (and most other stable-assertion-based approaches [22, 31, 51]). But this comes at the cost of specifications explicitly specifying aliasing constraints through choice of standard or separating conjunction. Separation logic specifies the behavior of code, not the restrictions on data transformation. Rely-guarantee reference types specify the possible evolution of data in the description of data itself. This means that assumptions and permission to modify state follow data-flow, rather than the control-flow-centric passing of assertions in most program logics.

7.1 Further Related Work

7.1.1 Fractional Permissions

Boyland's fractional permissions [8] ensure non-interference among aliases to a single object, by assigning a real- or rational-valued access permission to each alias allowing updates only for a full permission, read access for nonzero permissions, and allowing the natural splitting and merging of permissions that comes with real/rational arithmetic. This has since been generalized to separation algebras [19] (arbitrary cancellative partial commutative monoids). Both have been used for temporary sharing without interference for state and assertions [7], and are used in metatheory to define the meanings of qualifiers in typestate [37] and reference immutability [27]. However, by design, they ensure complete non-interference, not cooperation and robustness to interference.

7.1.2 VCC

VCC [13, 14] is a verification effort for concurrent C programs, using methodology derived directly from Spec# [3]. The main idea is to use invariants for shared state, which must be preserved by any action. Objects are packed and unpacked as in Spec#, where invariants must be restored upon re-packing. They encode a sort of rely-guarantee-style reasoning using *claim objects*, whose existence ensures no other thread can have the shared object open. A claim is essentially an object whose invariant depends on other objects. Because multi-object invariants are only permitted when admissible (any action preserving the invariant of one object must not violate the invariant of the other, similar to assertion stability in rely-guarantee logics or rely containment and predicate stability in our system), a claim serves as a witness that no other part of the program is actively using the claimed objects.

8. Conclusion

We have introduced *rely-guarantee references*, an adaptation of rely-guarantee program logics to reasoning about interference among aliases to shared objects. The technique generalizes reference immutability, connecting two previously-separate lines of research and addressing a fundamental problem in verifying imperative programs. We have shown the technique's usefulness by verifying correctness for several small examples (which are difficult to specify or verify with other approaches) in a prototype implementation. Our experience suggests that at least for small examples, the proof burden is reasonable. Rely-guarantee references demonstrate that aliasing in program verification can be addressed by adapting ideas from reasoning about thread interference.

Acknowledgments

References

[1] L. Augustsson. Cayenne — A Language with Dependent Types. In *ICFP*, 1998.

[2] H. Barendregt. *Lambda Calculi with Types*. 1991.

[3] M. Barnett, M. Fähndrich, K. R. M. Leino, P. Müller, W. Schulte, and H. Venter. Specification and Verification: The Spec# Experience. *Commun. ACM*, 54(6):81–91, June 2011.

[4] Y. Bertot and P. Castéran. *Interactive Theorem Proving and Program Development; Coq'Art: The Calculus of Inductive Constructions*. Springer Verlag, 2004.

[5] K. Bierhoff and J. Aldrich. Modular Typestate Checking of Aliased Objects. In *OOPSLA*, 2007.

[6] R. L. Bocchino, Jr., V. S. Adve, D. Dig, S. V. Adve, S. Heumann, R. Komuravelli, J. Overbey, P. Simmons, H. Sung, and M. Vakilian. A Type and Effect System for Deterministic Parallel Java. In *OOPSLA*, 2009.

[7] R. Bornat, C. Calcagno, P. O'Hearn, and M. Parkinson. Permission Accounting in Separation Logic. In *POPL*, 2005.

[8] J. Boyland. Checking Interference with Fractional Permissions. In *SAS*, 2003.

[9] V. Capretta. A Polymorphic Representation of Induction-Recursion. Retrieved 9/12/12. URL: http://www.cs.ru.nl/˜venanzio /publications/induction_recursion.pdf, March 2004.

[10] C. Chen and H. Xi. Combining Programming with Theorem Proving. In *ICFP*, 2005.

[11] A. Chlipala. *Certified Programming with Dependent Types*. http://adam.chlipala.net/cpdt/.

[12] A. Chlipala, G. Malecha, G. Morrisett, A. Shinnar, and R. Wisnesky. Effective Interactive Proofs for Higher-order Imperative Programs. In *ICFP*, 2009.

[13] E. Cohen, M. Dahlweid, M. Hillebrand, D. Leinenbach, M. Moskal, T. Santen, W. Schulte, and S. Tobies. Vcc: A practical system for verifying concurrent c. In *TPHOL*. 2009.

[14] E. Cohen, M. Moskal, W. Schulte, and S. Tobies. Local verification of global invariants in concurrent programs. In *CAV*. 2010.

[15] Coq Development Team. The COQ Proof Assistant Reference Manual: Version 8.4, 2012.

[16] T. Coquand and G. Huet. The Calculus of Constructions. *Information and Computation*, 76, 1988.

[17] W. Dietl, S. Drossopoulou, and P. Müller. Generic Universe Types. In *ECOOP*, 2007.

[18] T. Dinsdale-Young, M. Dodds, P. Gardner, M. Parkinson, and V. Vafeiadis. Concurrent Abstract Predicates. In *ECOOP*, 2010.

[19] R. Dockins, A. Hobor, and A. Appel. A fresh look at separation algebras and share accounting. In *ESOP*. 2009.

[20] M. Dodds, X. Feng, M. Parkinson, and V. Vafeiadis. Deny-Guarantee Reasoning. In *ESOP*. 2009.

[21] P. Dybjer. Inductive Families. *Formal Aspects of Computing*, 6:440–465, 1994.

[22] X. Feng. Local Rely-Guarantee Reasoning. In *POPL*, 2009.

[23] C. Flanagan and M. Abadi. Types for Safe Locking. In *ESOP*, 1999.

[24] C. Flanagan and S. N. Freund. Type-Based Race Detection for Java. In *PLDI*, 2000.

[25] T. Freeman and F. Pfenning. Refinement types for ml. In *PLDI*, 1991.

[26] H. Geuvers. The Church-Rosser Property for βη-reduction in Typed λ-calculi. In *LICS*, 1992.

[27] C. S. Gordon, M. J. Parkinson, J. Parsons, A. Bromfield, and J. Duffy. Uniqueness and Reference Immutability for Safe Parallelism. In *OOPSLA*, 2012.

[28] C. A. R. Hoare. An Axiomatic Basis for Computer Programming. *Commun. ACM*, 12(10):576–580, Oct. 1969.

[29] M. Hofmann. *Syntax and Semantics of Dependent Types, in Semantics and Logics of Computation*, chapter 3. 1997.

[30] J. B. Jensen and L. Birkedal. Fictional Separation Logic. In *ESOP*, 2012.

[31] C. B. Jones. Tentative Steps Toward a Development Method for Interfering Programs. *ACM TOPLAS*, 5(4):596–619, Oct. 1983.

[32] Kiselyov, Oleg and Lämmel, Ralf and Schupke, Keean. Strongly Typed Heterogeneous Collections. In *Haskell '04*, 2004.

[33] K. R. Leino and P. Müller. A Basis for Verifying Multi-threaded Programs. In *ESOP*, 2009.

[34] F. Militão, J. Aldrich, and L. Caires. Aliasing Control with View-based Typestate. In *FTfJP*, 2010.

[35] F. Militão, J. Aldrich, and L. Caires. Rely-Guarantee View Typestate. Retrieved 8/24/12, July 2012. URL http://www.cs.cmu.edu/ ˜foliveir/papers/rgviews.pdf.

[36] F. Militão, J. Aldrich, and L. Caires. Rely-guarantee view typestate (extended version). Retrieved 8/24/12, July 2012. URL http://www. cs.cmu.edu/~foliveir/papers/rgviews-TR.pdf.

[37] K. Naden, R. Bocchino, J. Aldrich, and K. Bierhoff. A Type System for Borrowing Permissions. In *POPL*, 2012.

[38] A. Nanevski, G. Morrisett, and L. Birkedal. Polymorphism and Separation in Hoare Type Theory. In *ICFP*, 2006.

[39] A. Nanevski, A. Ahmed, G. Morrisett, and L. Birkedal. Abstract Predicates and Mutable ADTs in Hoare Type Theory. In *ESOP*. 2007.

[40] A. Nanevski, G. Morrisett, A. Shinnar, P. Govereau, and L. Birkedal. Ynot: Dependent Types for Imperative Programs. In *ICFP*, 2008.

[41] L. Nistor and J. Aldrich. Verifying Object-Oriented Code Using Object Propositions. In *IWACO*, 2011.

[42] N. Nystrom, V. Saraswat, J. Palsberg, and C. Grothoff. Constrained Types for Object-Oriented Languages. In *OOPSLA*, 2008.

[43] S. Owicki and D. Gries. An Axiomatic Proof Technique for Parallel Programs I. *Acta Informatica*, pages 319–340, 1976.

[44] M. Parkinson and G. Bierman. Separation Logic and Abstraction. In *POPL*, 2005.

[45] C. Paulin-Mohring. Inductive Definitions in the System Coq: Rules and Properties. In *Typed Lambda Calculi and Applications*, 1993.

[46] A. Pilkiewicz and F. Pottier. The Essence of Monotonic State. In *TLDI*, 2011.

[47] P. M. Rondon, M. Kawaguchi, and R. Jhala. Low-Level Liquid Types. In *POPL*, 2010.

[48] M. Sozeau. Program-ing Finger Trees in Coq. In *ICFP*, 2007.

[49] N. Swamy, J. Chen, C. Fournet, P.-Y. Strub, K. Bhargavan, and J. Yang. Secure Distributed Programming with Value-dependent Types. In *ICFP*, 2011.

[50] M. S. Tschantz and M. D. Ernst. Javari: Adding Reference Immutability to Java. In *OOPSLA*, 2005.

[51] V. Vafeiadis and M. Parkinson. A Marriage of Rely/Guarantee and Separation Logic. In *CONCUR*. 2007.

[52] J. Wickerson, M. Dodds, and M. Parkinson. Explicit Stabilisation for Modular Rely-Guarantee Reasoning. In *ESOP*, 2010.

[53] H. Xi and F. Pfenning. Dependent Types in Practical Programming. In *POPL*, 1999.

[54] H. Xi, C. Chen, and G. Chen. Guarded Recursive Datatype Constructors. In *POPL*, 2003.

[55] Y. Zibin, A. Potanin, M. Ali, S. Artzi, A. Kiezun, and M. D. Ernst. Object and Reference Immutability Using Java Generics. In *ESEC-FSE*, 2007.

[56] Y. Zibin, A. Potanin, P. Li, M. Ali, and M. D. Ernst. Ownership and Immutability in Generic Java. In *OOPSLA*, 2010.

A. COQ, Positivity, and Inductive Datatypes

As mentioned elsewhere, the most natural way to implement linked data structures with structure-specific predicates and rely/guarantee conditions is to use mutual inductive types, with one (the datatype) indexing the other (the predicate or relation). Unfortunately for our implementation, COQ does not support this; Agda does, but lacks the tactic-based proving language we desire for managing proof burden.

Capretta [9] describes a general encoding of inductive-recursive [21] definitions (mutual definition of an inductive type and a recursive *function* destructing the datatype) in COQ.[6] The core idea is to make constructors that accept recursive members polymorphic over the type of recursive members (making the data type considerably "large" and therefore requiring COQ's -impredicative-set), coupling this with an inductive predicate that forces all recursive members to be instantiated with the correct type, and representing the final datatype as a dependent sum.

In RGREF, rather than using a dependent sum, we leverage the predicate parameter of the reference type. Thus, the monotonically growing list from Section 3.2 is defined in our implementation as:

```
Inductive rgrList' : Set :=
  | rgrl_nil' : rgrList'
  | rgrl_cons' : forall (C:Set) (P:hpred C) (R G:hrel C),
              nat -> ref{C|P}[R,G] -> rgrList'.
Inductive list_imm : hrel rgrList' :=
```

---

[6] Technically, in the original Calculus of Inductive Constructions with Set being impredicative.

```
  | imm_nil : forall h h', list_imm rgrl_nil' rgrl_nil' h h'
  | imm_cons : forall h h' P R G n (tl:ref{rgrList'|P}[R,G]),
            list_imm (h[tl]) (h'[tl]) h h' ->
            list_imm (rgrl_cons' rgrList' P R G n tl)
                      (rgrl_cons' rgrList' P R G n tl) h h'.
...
Inductive plumb : hpred rgrList' -> hpred rgrList' :=
  | plumb_nil : forall P h, plumb P rgrl_nil' h
  | plumb_cons : forall P R G h n (r:ref{rgrList'|P}[R,G]),
              plumb P (h[r]) h ->
              plumb P (rgrl_cons' rgrList' P R G n r) h.
...
Definition list P := ref{rgrList'|plumb P}[list_imm,list_imm].
```

Even if COQ implemented mutual inductive types that could
index each other, there are some cases where this encoding is still
desirable. Specifically, parameterizing a type constructor like the
linked list over an arbitrary rely and guarantee specific to rgrList'
would run afoul of the strict positivity requirement [45] imposed on
logically-sound inductive datatypes. For any type *A*, hpred *A* uses *A*
negatively. So parameterizing a constructor over a rely, guarantee,
or predicate would require adopting a similar encoding. None of our
paper examples use this style of expression, but in general, allowing
such polymorphism in a more principled way would be highly
desirable. Note that types in our COQ DSL *can* be parameterized
over type-polymorphic relations, such as The reference immutability
permissions in Section 3.3.

Another notable aspect of our shallow embedding in COQ is that
no predicate or relation that accesses heap-linked recursive structure
members may be defined as a fixpoint, because such members
require dereferencing within a specific heap, and COQ's termination
checker requires recursive arguments to be a syntactic subterm of
the recursion argument. Thus any "recursive" predicate or relation
over a heap-linked structure must actually be defined as an inductive
datatype in the universe Prop. Note that COQ's Fix combinator
for giving recursive definitions based on a well-founded relation
is essentially using an inductive datatype: the combinator allows
definitions to check by actually performing structural recursion on
the *proof term* that the recursion relation is well-founded [11].

B. Proof of CC Embedding

To justify our use of dependent types for proving, we give here a
shallow embedding of RGREF's core language into an extension
of CC. Two important subtleties are that our axiomatization of
heaps omits the potential for non-well-founded recursion through the
heap, and naïve proof terms might equate dereference expressions
evaluated in different heaps; these are discussed in Sections B.1
and F. Note that because dereference is treated as an uninterpreted
function, the CC reduction semantics (which lack a heap) are not
useful on the result of this translation; it is useful only for proving
properties of the source program.

Here we give a proof of Theorem 1, restated here:

There are mutually recursive transformation functions $\mathcal{P}[\![-]\!]$
and $I[\![-]\!]$ from a version of RGREF types, terms, and
contexts with explicit splitting, conversion, and full type
annotations to CC types, terms and contexts such that

- If $\Gamma \vdash_{\text{RGREF}} e : \tau$ then $\mathcal{P}[\![\Gamma]\!] \vdash_{\text{CC}} \mathcal{P}[\![e]\!] : \mathcal{P}[\![\tau]\!]$, and
- If $\Gamma \vdash_{\text{RGREF}} e : \tau \Rightarrow \Gamma'$ then $I[\![\Gamma]\!] \vdash_{\text{CC}} I[\![e]\!]^{\Gamma} : I[\![\tau]\!]^{\Gamma}$

Proof. The proof proceeds by simultaneous induction on the RGREF
derivations. Throughout the proof, we consider semantic judgments
(stability, containment, predicate and guarantee satisfaction, rela-
tional implication) true if and only if they are intuitionistically
(constructively) true, specifically by construction of an appropri-
ately typed proof term in the pure fragment of RGREF subject to

restrictions described in Section F. Figure 6 defines $\mathcal{P}[\![\cdot]\!]$ and $I[\![\cdot]\!]$
for RGREF types and environments.

First we consider the translation of the pure fragment of RGREF.

- Primitive rules (typing naturals, booleans, unit, pairs, Prop):
  These translate naturally into the corresponding axioms in CC
  with the basic extensions.
- Primitive recursors for naturals, booleans, unit, propositional
  equality; projection on pairs: These also translate into standard
  extensions of CC using the inductive hypothesis.
- V: Translates into the appropriate use of the variable rule (or
  START in PTS terms). Note that the reflexive splitting hypothesis
  requires no special treatment: our embedding only translates
  reflexively splitting values directly into CC, and substructurally
  treated values always end up in the pre- or post-environment of
  the $\mathcal{M}$ monad.
- Π-I: Translates to the analogous use of PRODUCT, using the
  inductive hypothesis for satisfying premises.
- Π-E: Translates similarly to APPLICATION.
- DEREF: Translates into the axiom for dereference, providing the
  CC-equivalent proof terms for the guarantee's reflexivity and the
  type results of relation folding.
- Π-F: Translates to ABSTRACTION similarly to the Π-I and Π-E
  rules.
- CONV: Translates one of two ways, according to the derivation
  of $\Gamma \vdash \tau \rightsquigarrow \tau'$:
  - If the derivation does not use P-⇒, R-⊂, or G-⊂, it translates
    to a natural use of CONVERSION in CC (which is actually
    more permissive than that permitted by CONV).
  - If the derivation does rely on some combination of P-⇒, R-
    ⊂, or G-⊂, then the term is actually an explicit conversion
    term. Translate the term itself into a use the appropriate
    conversion axioms on the source term, possibly combined
    with use of CONVERSION.
  Note that aside from P-⇒, R-⊂, and G-⊂, all of the type
  conversion rules correspond to restricted forms of β-equality
  on pure terms, which satisfies the conversion hypothesis for
  CONVERSION in CC.
- IMPURE: Impure abstractions and impure function types are
  typed according to the impure translation; apply the other
  inductive hypothesis.

The second half, the embedding of the impure fragment, proceeds
similarly, but by translating terms of type τ in RGREF into monadic
terms of type $\mathcal{M}\ p\ v\ p'$, for pseudovariable v of type $\mathcal{V}$, pre- and
post-environments p and p', and an indexed monad

$$\mathcal{M} : \text{list } (\mathcal{V} * \text{Prop}) \to \mathcal{V} \to \text{list } (\mathcal{V} * \text{Prop}) \to \text{Type}$$

that models substructural variable behavior in the imperative frag-
ment. The lists of variable and proposition (type) pairs are used as
type arguments to an HList-like structure that is a map from vari-
ables to values of the type matched in the list. $\mathcal{M}$ is essentially a
state monad, containing a value of its type argument, and a heap-
transformer in the form of a function from heaps to heaps that may
also modify the environment.

The environment structure is intuitively similar to Kiselyov et
al.'s better-known HList (heterogeneous list) structure [32]. In fact,
Kiselyov et al. describe an adaptation of HLists to use as extensible
records, which is roughly the role our environments play: a struc-
tured map from a pseudo-variable for a given (substructural) type to
elements of that type. Inside the monad $\mathcal{M}$, elements of the environ-
ments are accessed only indirectly, through primitives for splitting
or using elements. Substructural arguments to imperative functions
are passed as pseudo-variables that index the pre-environment of
the body's monadic type.

Each rule translates into a use of a simple monadic primitive, in some cases leaving holes for proof terms witnessing proof obligations such as a write satisfying a guarantee; it is up the the implementation to choose appropriate proof terms that avoid the issues discussed in Section F. In some cases the expected type of the translation has fewer elements in the output context than the natural typing for the term, in which case the translation described below should be monadically sequenced with the primitive for dropping substructural context elements.

- FUN: Inductively translate the function body, which will produce a derivation of

$$\Gamma, x : \mathcal{V} \vdash_{\text{CC}} I[\![e]\!]^{\Gamma, x : \mathcal{P}[\![\tau]\!]} : I[\![\tau']\!]^{\Gamma, x : \mathcal{P}[\![\tau]\!]}$$

  Encoding this as a standard λ expression, binding the argument x as a pseudo-variable ($\mathcal{V}$), preceding it with injection of the captured variables, and sequencing after it a drop of the argument from the output gives the result a type matching the $I[\![-]\!]$ metafunction results.
- RUN: By induction the subexpressions translate to expressions with monadic types. The metafunction result expression binds the names for the (substructurally-treated) function and argument, and the m-app primitive builds a computation that consumes both names, applies the function to the argument variable inside the monad, and runs the resulting computation.
- WRITE: The heap_write primitive is reasonably straightforward, producing a computation that stores the result of e's translation into the heap cell referenced by x. The only subtle part is the translation of the proof obligation. The implementation uses runM (for running $\mathcal{M}$ monads) in the proof. runM is effectively the denotational semantics of the expression in the source language, and properties of its results are proven using axioms relating the computation to the behavior of a resource-insensitive heap transformer function. runM has the type

$$\Pi \, \Gamma \, \Gamma' : \text{list}(\mathcal{V} * \text{Prop}) \to \Pi x : \mathcal{V} \to \Pi e : E \, \Gamma \to$$
$$\mathcal{M} \, \Gamma \, x \, \Gamma' \to \text{heap} \to \text{heap}$$

- SWAP: Similar to write.
- ALLOC: Allocation uses a similar translation style to writing, using runM and a few axioms to inspect and reason about the results of the computation in order to establish the initial predicate.
- IDEREF: Translates to an imperative dereference primitive.
- SUB: The pure expression translates to a monadic lifting of the pure value, and the type translates simply to a monadic wrapping of the value in a substructural context.
- Variable reads, recursors, constructors, etc. translate similarly to above.

$\square$

B.1 Recursion Through the Heap

One important aspect the proof above ignores is the potential to recur through the heap, as this term would:

```
Program Example heap_recursion :=
  (* Allocate a unit->nat on the heap *)
  (* Assume predicate any, rely/guarantee havoc *)
  fn <- alloc (fun _:unit => 3);
  (* Close fn2 over fn *)
  fn2<- alloc (fun u:unit => (!fn) u);
  (* Point fn2's function reference to itself! *)
  fn <- fn2;
  return (!fn2) tt.
```

We do not consider this fundamentally problematic; it simply requires our implementation to prevent creating proof terms relying

on such behavior; we restrict equational reasoning when returning closures from the pure sublanguage to the impure context or when invoking a pure function pulled from the heap (see Section F). So proofs that a predicate holds of a new allocation, or that a new predicate holds and a guarantee is satisfied after a write, are valid *when the allocated/stored expression terminates*, and otherwise the program does not terminate at runtime.

We do not consider this a serious issue for three reasons. First, non-termination is standard in impure languages. Most rely-guarantee logics and other program logics are modulo termination, and this proof-by-nontermination issue arises only when the imperative code has explicitly used mutation to introduce indirect recursion. Second, our implementation can ensure that the only "invalid" proofs prove properties of non-terminating program terms, rather than relying on non-termination in the proof itself. Third, disallowing recursion through the heap is straightforward: an implementation could for example

- Restrict the heap to first-order (not containing closures) values (a common restriction in program logics), or
- Stratify the term language further and allow only functions that do not (transitively) dereference a reference to a function type (or a structure containing a function) to be stored in the heap, or
- Restrict dereference in the pure fragment to references to first-order types (which will not contain pure closures) and impure abstractions (which are opaque and not invocable in the pure fragment).

C. Soundness

The proof of soundness follows the sketch in Section 4.1. The dynamic semantics are mostly standard for an ML-like language with references, with a couple small changes:

- Binding: Variable binding in the imperative core is stack-based, not substitution-based, to handle the substructural splitting behavior in the imperative language. Binding in the pure language is substitution based (essentially pre-splitting and inlining read results). [7]
- Variable capture: Variables captured by either a pure or impure closure are required to be self-splitting, but it would be unsound to allow those captured uses to read the stack location if the closure were executed after impure code split the stack reference with different permissions. To prevent this, closures are only values when they have no free variables, and there are reduction rules that perform the reads for free variables.
- Locations: Locations are represented by a tuple of not only the heap "index" (the traditional basic location) but also with the heap storage type, predicate and rely/guarantee relations explicitly tagged as part of the value. The tags are not required for functionality, but their presence simplifies the soundness argument.
- Subsumption: We actually prove soundness for a slight translation of the calculus to one with explicit reference conversion casts placed wherever expressions were typed using the rules P-$\Rightarrow$, R-$\subset$, and G-$\subset$.

---

[7] In an implementation, substitution is fine for both bindings as the relations attached to arguments serve no operational purpose: they are present here only for proving that resource semantics are respected. This would also simplify interaction between the pure and impure languages, allowing self-splitting variable bindings to be captured between sublanguages. In our implementation embedded in COQ, variables of self-splitting types may be captured by either abstraction and used only purely; substructural variables are encapsulated in a monad.

- Substructural behavior: Variable splitting and dropping are explicitly identified in the source, to allow the semantics to modify the stack appropriately.

- Step Granularity: Small-step reduction semantics are used, but the pure computation's semantics include only a single input heap, and no output, since the pure terms cannot modify the heap.

- Folding: Folding has a runtime representation. Reads produce the heap value wrapped in a deferred folding construct which lazily pushes the guarantee restriction through a data structure as components are evaluated. This is necessary for dependent type constructors like propositional equality, where types appear in values of the type.

- It is sometimes necessary to reduce multiple heap dereferences at once if they are in some way connected by equality of dependent types containing that dereference. In a way this is a very unusual variable binding. This differs from naïve semantics only when two dereferences of the same location are related by dependent types, inside different closures, and and one of those closures is returned to the impure context. We have yet to imagine an example of desirable *computational* code that observes this difference.

While considering the role of predicate and guarantee obligations from the typing derivations, note that the type judgments themselves do not require proof terms to prove predicates and guarantee obligations. Those predicates and relations are *specified* using propositions-as-types, but the actual proof method is up to the implementation. Section F discusses the subtleties of using proofs-as-programs to actually prove various obligations.

Execution must preserve two critical invariants beyond standard invariants for well-typed heaps, are:

- For each reference $r : \mathsf{ref}\{T \mid P\}[R, G]$ in the stack, heap, or expression under reduction, there exists a proof of $P\ (h[r])\ h$ for the current heap h.

- For each pair of references $p : \mathsf{ref}\{T \mid P\}[R, G]$ and $q : \mathsf{ref}\{T \mid P'\}[R', G']$ in the stack, heap, or expression under reduction, if p and q alias (point to the same heap cell) then $\emptyset \subset G' \subseteq R$ and $\emptyset \subset G \subseteq R'$

Soundness proceeds as a type preservation proof, by induction on the evaluation step taken.

Lemma 1 (Pure Preservation). *If $H;\Sigma;\Gamma \vdash e : \tau$ and $H; e \to e'$, then there exists some τ' such that $H;\Sigma;\Gamma \vdash e' : \tau'$ and $H;\Gamma \vdash \tau' \rightsquigarrow \tau$.*

The pure cases are mostly straightforward (recall that subject reduction — progress and preservation — hold for CC with β-conversion [2, 26]), so we focus discussion on the cases for impure rules. The one notable point in the pure cases is reduction of a dereference expression — assume $!r$ — which applies relation folding ($[R, G] \gg \tau$, Figure 5). In this case note that the folding — a form of weakening guarantees of read results — ensures that the guarantees in the result of $!r$ allow no more heap changes than r's guarantee. This is required to preserve the compatible alias guarantee, because some alias of r may have a rely equal to r's guarantee, so if a read result allowed too strong a guarantee, that result might allow actions that would violate the alias's rely, potentially invalidating refinements.

Evaluating $!r$ also produces labeled expressions, $\langle !r \rangle a$, where a is the folded result of the heap lookup, and is convertible to $!r$ itself to aid type preservation. The labeled expression is a witness of a heap dependency. It is the reduction of these heap reads that requires the type before and after be related by $H;\Gamma \vdash \tau' \rightsquigarrow \tau$, which gains an additional runtime rule that a dereference-tagged expression

converts to the dereference expression. This is useful in structural rules like the inductive cases of pure function application, where for example a reduction of the function term could introduce a tagged value into the argument type. For the application to continue to check, appearances of the same (dereferenced) value in the type of the argument position must be convertible to the argument type.

All heap witnesses are removed (reduced to a) by another reduction rule (from the surrounding imperative context) before producing a value that flows back into imperative computation, taking advantage of the fact that replacing $\langle !r \rangle a$ by a preserves typing, up to replacing the tagged expression with a in the type. This is done in two steps by the context rule for embedding pure expressions within impure contexts, under the meta-function $\mathsf{heap\_indep}$. The runtime typing rule for the pure-in-impure embedding requires the pure expression typecheck without using the untagging conversion rule, so $\mathsf{heap\_indep}$ rewrites the pure evaluation result to remove the use of untagging. This involves changing (as few as possible) instances of $!l$ to $\langle !l \rangle a$ in types when previously a conversion was necessary, and at constructors occasionally changing a term-level dereference as well. Thus, this is the reduction of multiple dereferences at once mentioned earlier. Then a substitution is performed using the following lemma:

Lemma 2 (Untagging). *For all references ℓ and tagged expressions $\langle !\ell \rangle a$ such that $H;\Sigma;\varepsilon \vdash \langle !\ell \rangle a : \tau'$, if $H;\Sigma;\Gamma \vdash e : \tau$ without using the untagging conversion, then $H;\Sigma;\Gamma \vdash e[\langle !\ell \rangle a/a] : \tau[\langle !\ell \rangle a/a]$.*

Proof. By inversion on the typing of the tagged value, $H;\Sigma;\varepsilon \vdash a : \tau'$. Proof follows by induction on the typing derivation for e. The CONV case proceeds by induction on the type conversion, where the untagging case holds vacuously. □

Because only closed result types are permitted to flow from pure subterms back into imperative computation (SUB and ΠI-E), this extra "unlabelling" step preserves the type, so all values that persist across pure evaluations are heap-independent.

With these lemmas in hand, impure preservation is reasonably straightforward.

Lemma 3 (Impure Preservation). *If $\vdash S;H;e : \Gamma;\Sigma;\tau \Rightarrow \Gamma'$, and $S;H;e \to S';H';e'$, then there exists a Γ'' and Σ' such that $\vdash S';H';e' : \Gamma'';\Sigma';\tau \Rightarrow \Gamma'$.*

Proof. Note that the initial state (well-typed expression and the empty heap and stack) satisfies all invariants.

- CALL: Reduction of the procedure or argument is sound by induction. In the case where the procedure is actually applied, the only affected parts of the state are the stack (which gains a fresh variable with the argument as its value, with the obvious type), and the expression (now the body with the fresh variable substituted for the source variable). Stack typing follows naturally, expression typing follows from α-renaming and a lemma that replacing well-typed subexpressions preserves typing, and heap typing is unchanged.

- ASSIGN: Stack and expression typing is straightforward. Basic heap typing is straightforward. Establishing that the predicate holds in the new heap is straightforward, using inversion on the heap write type rule (including strong updates to predicates). This rule may create a new alias of some reference, but preserving non-conflicting relies and guarantees among aliases is straightforward. The more subtle part of this case is proving that all other references' predicates still have proofs in the new heap. For direct aliases of the write target, by the agreeable R/G condition, stability of all predicates over their respective relies, and the guarantee satisfaction (by inversion on the typing rule), those predicates are preserved. For references from which the

modified cell is reachable, a similar reasoning applies using the containment requirements of well-formed reference types. For references from which the modified cell is not reachable, preservation is by the precision requirements on predicates.

- SWAP: Similar to the assignment case.
- ALLOCATE: Stack soundness is preserved, expression soundness is straightforward. Basic heap soundness is straightforward, leaving the RGREF-specific heap invariants as remaining proof obligations. By inversion on the typing rule for allocation, there is a proof of the predicate on the allocated value in all heaps, therefore one exists for the new heap. The new object is (initially) unaliased, so all aliases' rely and guarantee imply each other. For previously-existing references, the allocation is not reachable from any existing allocation, so proofs are preserved by the fact that all existing references' predicates are precise (insensitive to changes outside their reachable heap).
- DROP-VAR: Stack and heap typing are straightforward, as is expression typing. The main invariant that could be violated is that aliases' rely and guarantee conditions might conflict; this invariant is preserved because the operation moves an existing reference, it does not create an additional alias.
- SPLIT-VAR: Similar to the DROP-VAR case, except the value is actually split according to the elaborated syntax for splitting variable reads, which is only well-typed if values of the type split according to $\Gamma \vdash \tau \prec \tau' \divideontimes \tau''$, which preserves compatible rely and guarantee conditions.
- PURE: Justified by soundness for the pure sublanguage, plus restrictions on pure/impure interactions (specifically that runtime typing requires that the pure expression must typecheck without using the heap as a source of definitional equality — i.e., without the untagging conversion).

$\square$

D. Operational Semantics

This section describes the operational semantics of the core language from Section 4. Figure 8 gives the most interesting rules, while Figure 9 gives the remaining rules that effectively define contexts and evaluation order. Recursors (not shown) follow the standard reduction rules.

The most unusual part of the semantics is the imperative reduction "evaluation context" rule for reducing pure expressions in an impure context, mentioned earlier. This rule sometimes reduces multiple dereferences of the same location in one step, including inside unevaluated closures.

The semantics reducing multiple dereference expressions is very unusual, and could hypothetically lead to unexpected results. Implementations could impose stronger analyses to prevent writing terms that would observe the change in behavior (thus obviating the unusual semantics). However, we have been unable to conceive of a desirable term that observes this atypical reduction; recall that we use the pure fragment for two different purposes: computation and specification. Observing the multiple-dereference behavior requires using dereference expressions in types.

The simplest example we can think of and a few reduction steps are shown in Figure 7. Assuming ℓ is a reference to a natural number (3 in the current heap), the type of this term is $\text{unit} \to \text{nat}$. The first reduction shown substitutes the reference, and the second performs the first actual dereference. The result of the second reduction (the last line) only type-checks because of the untagging conversion rule, which says $\langle !\ell \rangle 3 \rightsquigarrow !\ell$. If the tag were stripped from the term without further ado and reduction proceeded naïvely, the recursor's result terms would be a proof that $3 = 3$ and a proof that $!\ell = !\ell$. While this is typeable, it would require introducing additional

dependencies into the type (not strict preservation) and we don't know if even that would be possible in general. And while one could conceivably coerce the new term to type-checking using the heap for unrestricted definitional equality, this would be unsound: this problematic term could be reduced inside an impure computation which subsequently stored 4 to ℓ's heap cell. For this reason, our semantics would produce the final term in the figure, deduce that immediately reducing the inner dereference would ensure the term is well-typed in *any* heap, and perform that additional dereference.

Note that to observe the unusual semantics, a term must:

- Dereference a location in a context that introduce the dereference expression into a type (such as the redex location of a reflexivity proof), in a part of the term that is reduced before returning to the impure context
- Return a closure to the surrounding impure context, which also dereferences the same location in a context that injects dereference into a type
- Relate (via types) the "outer" and "inner" terms whose types depend on dereference of the same location.

If a pure embedding's type is any "flat" type (not containing a closure), then the multiple-dereference semantics are irrelevant; all dereference expressions in the pure context will be reduced in the same heap (or thrown away, for example from an unused branch of a recursor). The heap_indep step also imposes restrictions on reasoning about pure terms; see Section F.

E. Static Semantics of Dynamic State

Figure 10 gives the typing judgments for dynamic program states.

F. Equational Reasoning with Dereference

The unusual simultaneous-dereference semantics in the pure sublanguage highlight an important subtlety of equational reasoning (a subtlety that would exist even with alternative restrictions to remove the unusual semantics). Specifically, it is unsound to equate two dereference expressions that may be reduced in different heaps! Further, as mentioned in Section B, the most straightforward axiomatization for properties of heaps would allow the implementation to form proof terms using dereference to accidentally use non-well-founded recursion. Note that our type rules do not actually say how stability, precision, relations, or predicates are proven; this is why.

There are a few approaches to handling the issues with equating dereferences that are evaluated in different heaps (Section B.1 discusses the non-termination issues). The approach we favor is to only permit full equational reasoning for properties of pure terms whose return type contains no closures. In this case, all dereferences of the same location will be reduced with the same heap (before control returns to the imperative fragment), so equating syntactically identical dereferences is sound. In cases where a pure subterm returns a value that may contain unreduced dereference expressions, either the whole term or (a conservative overapproximation of) the closures that may be unreduced when the pure term becomes a value must be abstracted, hiding them from equational reasoning principles.

This is also the motivation for disallowing references in the definition of predicates and relations in the core language. Rather than complicating the type system's core ideas with richer checks to prevent dereferences in predicates and the like, we simply prevent the introduction of references into those subterms. An alternative would be to introduce more distinct syntactic categories so dereference expressions in predicates, rely and guarantee relations simply would not even parse. Our prepend-only list in Section 3.2 does project a location (and another value) from the term language

into a predicate, but this is sound because it is not used in a heap-sensitive way (dereference).

Our prototype does not enforce the required restrictions on equational reasoning, because it is in some sense "too shallow" of an embedded DSL: because we directly leverage Coq's dependent product for RGref's dependent product, we cannot restrict its application without an additional preprocessor or adding a compiled OCaml plugin to restrict interaction between native CIC and RGref-specific terms. We believe this is a reasonable trade-off. The prototype's goal is to evaluate the rely-guarantee reference approach's proof burden. The technically required restrictions on equational reasoning should not be difficult in principle to enforce, and we do not believe they would noticeably affect how code would be written. A production-quality implementation of RGref would of course need to enforce the richer restrictions. It is worth noting that this weakness from interaction between Coq's raw terms and DSL-specific terms is not unique to us; the Ynot [12, 40] implementation of Hoare Type Theory [38, 39] has similar risks when Coq primitives are used in unintended ways with Ynot axioms. This source of unsoundness could be avoided with a deeper embedding, as has been done for some separation logic work, or in a system not based on dependent type theory (which would need to construct many reasoning principles on its own).

$$\boxed{\Gamma \vdash \tau \prec \tau \divideontimes \tau} \qquad \frac{\tau \in \{\texttt{nat}, \texttt{bool}, \texttt{unit}, \mathsf{Prop}, \mathsf{Type}, \texttt{heap}, \_ = \_, \Pi x : \tau \to \tau', \tau \xrightarrow{\mathcal{M}} \tau'\}}{\Gamma \vdash \tau \prec \tau \divideontimes \tau} \qquad \frac{\Gamma \vdash \tau \prec \tau_a \divideontimes \tau_b \quad \Gamma \vdash \sigma \prec \sigma_a \divideontimes \sigma_b}{\Gamma \vdash (\tau, \sigma) \prec (\tau_a, \sigma_a) \divideontimes (\tau_b, \sigma_b)}$$

$$\text{Ref-}\divideontimes \; \frac{\Gamma \vdash \mathsf{ref}\{b \mid \phi''\}[R'', G''] \quad \emptyset \subset \llbracket G' \rrbracket \subseteq \llbracket R'' \rrbracket \quad \begin{array}{c} \Gamma \vdash \mathsf{ref}\{b \mid \phi'\}[R', G'] \\ \emptyset \subset \llbracket G'' \rrbracket \subseteq \llbracket R \rrbracket' \end{array} \quad \llbracket G' \rrbracket \cup \llbracket G'' \rrbracket \subseteq \llbracket G \rrbracket \quad \llbracket R \rrbracket \subseteq \llbracket R' \rrbracket \quad \llbracket R \rrbracket \subseteq \llbracket R'' \rrbracket}{\Gamma \vdash \mathsf{ref}\{b \mid \phi\}[R, G] \prec \mathsf{ref}\{b \mid \phi'\}[R', G'] \divideontimes \mathsf{ref}\{b \mid \phi''\}[R'', G'']}$$

$$\boxed{\Gamma \vdash e : \tau} \qquad \overline{\Gamma \vdash n : \texttt{nat}} \qquad \overline{\Gamma \vdash b : \texttt{bool}} \qquad \overline{\Gamma \vdash \texttt{tt} : \texttt{unit}} \qquad \text{Axiom} \; \frac{}{\Gamma \vdash \mathsf{Prop} : \mathsf{Type}} \qquad \text{V} \; \frac{\Gamma \vdash \tau \prec \tau \divideontimes \tau}{\Gamma, x : \tau \vdash x : \tau}$$

$$\Pi\text{-I} \; \frac{\Gamma \vdash \tau \prec \tau \divideontimes \tau \quad \forall \tau \in \Gamma . \vdash \tau \prec \tau \divideontimes \tau \quad \Gamma \vdash (\Pi x : \tau \to \tau') : \sigma \quad \sigma \in \{\mathsf{Type}, \mathsf{Prop}\} \quad \Gamma, x : \tau \vdash e : \tau'}{\Gamma \vdash (\lambda x : \tau . e) : \Pi x : \tau \to \tau'}$$

$$\Pi\text{-E} \; \frac{\Gamma \vdash e_1 : \Pi x : \tau \to \tau' \quad \Gamma \vdash e_2 : \tau}{\Gamma \vdash e_1 \, e_2 : \tau'[x/e_2]} \qquad \frac{\begin{array}{c} \varepsilon \vdash \tau \quad \varepsilon \vdash \tau' \\ x : \tau \vdash e : \tau' \Rightarrow \Gamma \end{array}}{\Gamma \vdash (\lambda_{\mathcal{M}} x : \tau . e) : \tau \xrightarrow{\mathcal{M}} \tau'} \qquad \text{Deref} \; \frac{\begin{array}{c} \Gamma \vdash e : \mathsf{ref}\{\tau \mid P\}[R, G] \quad \tau' = [R, G] \gg \tau \\ \Gamma \vdash \tau \prec \tau \divideontimes \tau \quad G \text{ reflexive} \end{array}}{\Gamma \vdash \, !e : \tau'}$$

$$\Pi\text{-F} \; \frac{\begin{array}{c} \Gamma \vdash \tau : \gamma \quad \Gamma \vdash \tau \prec \tau \divideontimes \tau \\ \Gamma, x : \tau \vdash \tau' : \sigma \quad \gamma, \sigma \in \{\mathsf{Type}, \mathsf{Prop}\} \end{array}}{\Gamma \vdash \Pi x : \tau \to \tau' : \sigma} \qquad \text{Conv} \; \frac{\Gamma \vdash e : \tau \quad \Gamma \vdash \tau \leadsto \tau'}{\Gamma \vdash e : \tau'}$$

$$\boxed{\Gamma \vdash \tau \leadsto \tau} \qquad \text{P-}\Rightarrow \; \frac{\forall v, h . \, P \, v \, h \to P' \, v \, h \quad \text{stable } P' \, R}{\Gamma \vdash \mathsf{ref}\{\tau \mid P\}[R, G] \leadsto \mathsf{ref}\{\tau \mid P'\}[R, G]} \qquad \text{R-}\subset \; \frac{\llbracket R \rrbracket \subseteq \llbracket R' \rrbracket \quad \text{stable } P \, R'}{\Gamma \vdash \mathsf{ref}\{\tau \mid P\}[R, G] \leadsto \mathsf{ref}\{\tau \mid P\}[R', G]}$$

$$\text{G-}\subset \; \frac{\llbracket G' \rrbracket \subseteq \llbracket G \rrbracket}{\Gamma \vdash \mathsf{ref}\{\tau \mid P\}[R, G] \leadsto \mathsf{ref}\{\tau \mid P\}[R, G']}$$

$$\boxed{\Gamma \vdash e : \tau \Rightarrow \Gamma} \quad \text{V-}\emptyset \; \frac{}{\Gamma, x : \tau \vdash x : \tau \Rightarrow \Gamma} \quad \text{V-}\divideontimes \; \frac{\Gamma \vdash \tau \prec \tau' \divideontimes \tau''}{\Gamma, x : \tau \vdash x : \tau' \Rightarrow \Gamma, x : \tau''} \quad \mathcal{M}\text{-I} \; \frac{\begin{array}{c} \forall \tau \in \Gamma . \vdash \tau \prec \tau \divideontimes \tau \\ \Gamma, x : \tau \vdash e : \tau' \Rightarrow \Gamma' \end{array}}{\Gamma \vdash (\lambda_{\mathcal{M}} (x : \tau) . e) : \tau \xrightarrow{\mathcal{M}} \tau' \Rightarrow \Gamma}$$

$$\mathcal{M}\text{-E} \; \frac{\begin{array}{c} \Gamma \vdash e_1 : \tau \xrightarrow{\mathcal{M}} \tau' \Rightarrow \Gamma_1 \\ \Gamma_1 \vdash e_2 : \tau \Rightarrow \Gamma_2 \end{array}}{\Gamma \vdash e_1 \, e_2 : \tau' \Rightarrow \Gamma_2} \qquad \text{Write} \; \frac{\begin{array}{c} \Gamma \vdash e : \tau \Rightarrow \Gamma', x : \mathsf{ref}\{\tau \mid P\}[R, G] \quad \forall h, h' : heap . \, h' = \llbracket e \rrbracket(h) \to P \, (!x) \, h' \to P' \, \llbracket e \rrbracket \, h'[x \mapsto \llbracket e \rrbracket] \\ \text{stable } P' \, R \quad \forall h, h' : heap . \, h' = \llbracket e \rrbracket(h) \to P \, (!x) \, h' \to G \, (!x) \, e \, h' \, h'[x \mapsto \llbracket e \rrbracket] \end{array}}{\Gamma \vdash x \leftarrow e : \texttt{unit} \Rightarrow \Gamma', x : \mathsf{ref}\{\tau \mid P'\}[R, G]}$$

$$\text{IDeref} \; \frac{\begin{array}{c} \Gamma \vdash e : \mathsf{ref}\{\tau \mid P\}[R, G] \quad \tau' = [R, G] \gg \tau \\ \Gamma \vdash \tau \prec \tau \divideontimes \tau \quad G \text{ reflexive} \end{array}}{\Gamma \vdash \, !e : \tau' \Rightarrow \Gamma'} \qquad \text{Alloc} \; \frac{\Gamma \vdash e : \tau \Rightarrow \Gamma' \quad \emptyset \subset G \quad \text{stable } P \, R \quad \forall h : heap . \, P \, e \, h}{\Gamma \vdash \mathsf{alloc} \, e : \mathsf{ref}\{\tau \mid P\}[R, G] \Rightarrow \Gamma'}$$

$$\text{Swap} \; \frac{\Gamma \vdash x \leftarrow e : \texttt{unit} \Rightarrow \Gamma', x : \mathsf{ref}\{\tau \mid P'\}[R, G]}{\Gamma \vdash \mathsf{swap}(x, e) : \tau \Rightarrow \Gamma', x : \mathsf{ref}\{\tau \mid P'\}[R, G]} \qquad \text{Pure} \; \frac{\varepsilon \vdash e : \tau \quad \varepsilon \vdash \tau : \mathsf{Prop}}{\Gamma \vdash e : \tau \Rightarrow \Gamma}$$

$$\text{I}\Pi\text{-E} \; \frac{\Gamma \vdash e_1 : \Pi x : \tau \to \tau' \Rightarrow \Gamma' \quad \Gamma' \vdash e_2 : \tau \Rightarrow \Gamma'' \quad x \notin FV(\tau')}{\Gamma \vdash e_1 \, e_2 : \tau' \Rightarrow \Gamma''}$$

$$\boxed{\Gamma \vdash \tau : \sigma} \qquad \text{wf-Ref} \; \frac{\begin{array}{c} \varepsilon \vdash R : hrel(\tau) \quad \varepsilon \vdash G : hrel(\tau) \quad \varepsilon \vdash P : hprop(\tau) \quad \text{closed } R \\ \text{precise } R \quad \text{precise } G \quad \text{precise } P \quad \text{stable } P \, R \quad \varepsilon \vdash \tau : \mathsf{Prop} \end{array}}{\Gamma \vdash \mathsf{ref}\{\tau \mid P\}[R, G] : \mathsf{Prop}}$$

with metafunctions $hrel(\tau) \overset{\text{def}}{=} \tau \to \tau \to \texttt{heap} \to \texttt{heap} \to \mathsf{Prop}$ and $hprop(\tau) \overset{\text{def}}{=} \tau \to \texttt{heap} \to \mathsf{Prop}$

Figure 4. Typing. Not shown: standard recursors for naturals, booleans, pairs, identity types [29]. Also not shown: standard well-formed contexts, most of (pure) expression/type conversion ($\Gamma \vdash \tau \leadsto \tau$) (Figure 5).

$$\boxed{\Gamma \vdash \tau \rightsquigarrow \tau \text{ cont.}} \qquad \frac{\Gamma \vdash \tau}{\Gamma \vdash \tau \rightsquigarrow \tau} \qquad \frac{\Gamma \vdash \tau \rightsquigarrow \tau' \quad \Gamma \vdash \tau' \rightsquigarrow \tau''}{\Gamma \vdash \tau \rightsquigarrow \tau''} \qquad \frac{\Gamma \vdash \sigma \rightsquigarrow \tau \quad \Gamma, x : \tau \vdash \tau' \rightsquigarrow \sigma'}{\Gamma \vdash \Pi x : \tau \to \tau' \rightsquigarrow \Pi x : \sigma \to \sigma'} \qquad \frac{\Gamma \vdash \sigma \rightsquigarrow \tau}{\Gamma \vdash (\lambda x : \tau.e) \rightsquigarrow (\lambda x : \sigma.e)}$$

$$\frac{\tau =_{\beta_v} \tau'}{\Gamma \vdash \tau \rightsquigarrow \tau'} \qquad \boxed{\vdash \Gamma} \qquad \overline{\vdash \varepsilon} \qquad \frac{\vdash \Gamma \quad x \notin \Gamma \quad \Gamma \vdash \tau}{\vdash \Gamma, x : \tau}$$

$$\boxed{\Gamma \vdash \tau : \sigma \text{ cont.}} \qquad \frac{\tau \in \{\mathtt{nat}, \mathtt{bool}, \mathtt{unit}, \mathsf{True}, \mathsf{False}\}}{\Gamma \vdash \tau : \mathsf{Prop}} \qquad \frac{\Gamma \vdash \tau : \gamma \quad \Gamma \vdash \sigma : \gamma}{\Gamma \vdash (\tau, \sigma) : \gamma} \qquad \frac{\Gamma \vdash \tau : \mathsf{Prop} \quad \Gamma \vdash \tau' : \mathsf{Prop}}{\Gamma \vdash \tau \xrightarrow{\mathcal{M}} \tau' : \mathsf{Prop}}$$

$$\frac{\Gamma \vdash A : \sigma \quad \Gamma \vdash a : A}{\Gamma \vdash \mathsf{refl}(a) : a = a} \qquad \frac{\Gamma \vdash \tau \prec \tau \circledast \tau}{\Gamma \vdash =_\tau : \tau \to \tau \to \mathsf{Prop}} \qquad \frac{}{\Gamma \vdash \mathsf{truth} : \mathsf{True}}$$

$$\frac{\Gamma \vdash u : \mathtt{unit} \quad \Gamma \vdash d : C[x/\mathsf{tt}] \quad \Gamma, x : \mathtt{unit} \vdash C : \sigma \quad \sigma \in \{\mathsf{Prop}, \mathsf{Type}\}}{\Gamma \vdash \mathcal{R}_{\mathtt{unit}}(u, d) : C[x/u]} \qquad \frac{\Gamma \vdash p : (\tau \ast \sigma)}{\Gamma \vdash \mathsf{fst}\ p : \tau} \qquad \frac{\Gamma \vdash p : (\tau \ast \sigma)}{\Gamma \vdash \mathsf{snd}\ p : \sigma}$$

$$\frac{\Gamma \vdash h : \mathsf{heap} \quad \Gamma \vdash e : \mathsf{ref}\{A \mid P\}[R, G]}{\Gamma \vdash h[e] : A} \qquad \frac{\Gamma \vdash b : \mathsf{False} \quad \Gamma, x : \mathsf{False} \vdash C : \sigma \quad \sigma \in \{\mathsf{Prop}, \mathsf{Type}\}}{\Gamma \vdash \mathcal{R}_{\mathsf{False}}(b) : C[x/b]}$$

$$\frac{\Gamma \vdash n : \mathtt{nat} \quad \Gamma \vdash c : C[x/0] \quad \Gamma \vdash f : \Pi n : \mathtt{nat} \to (C[x/n] \to C[x/Sn]) \quad \Gamma, x : \mathtt{nat} \vdash C : \sigma \quad \sigma \in \{\mathsf{Prop}, \mathsf{Type}\}}{\Gamma \vdash \mathcal{R}_{\mathtt{nat}}(n, c, f) : C[x/n]}$$

$$\frac{\Gamma \vdash b : \mathtt{bool} \quad \Gamma \vdash t : C[x/\mathsf{true}] \quad \Gamma \vdash f : C[x/\mathsf{false}] \quad \Gamma, x : \mathtt{bool} \vdash C : \sigma \quad \sigma \in \{\mathsf{Prop}, \mathsf{Type}\}}{\Gamma \vdash \mathcal{R}_{\mathtt{bool}}(b, t, f) : C[x/b]}$$

$$\frac{\Gamma \vdash b : \mathsf{True} \quad \Gamma \vdash t : C[x/\mathsf{truth}] \quad \Gamma, x : \mathsf{True} \vdash C : \sigma \quad \sigma \in \{\mathsf{Prop}, \mathsf{Type}\}}{\Gamma \vdash \mathcal{R}_{\mathsf{True}}(b, t) : C[x/b]}$$

$$\frac{\Gamma \vdash A : \sigma \quad \Gamma \vdash c : a =_A b \quad \Gamma \vdash a : A \quad \Gamma \vdash b : A \quad \Gamma \vdash d : C[x/a, y/a, z/\mathsf{refl}(a)] \quad \Gamma, x : A, y : A, z : x = y \vdash C : \mathsf{Prop}}{\Gamma \vdash \mathcal{R}_=(c, d) : C[x/a, y/b, z/c]}$$

$$
\begin{aligned}
\mathsf{stable}_\tau\ (P : hprop(\tau))\ (R : hrel(\tau)) \quad &\overset{\mathrm{def}}{=} \quad \forall x : \tau. \forall h, h' : \mathsf{heap}.\, P\, x\, h \to R\, x\, h\, h' \to P\, x\, h' \\
\mathsf{precise}_\tau\ (P : hpred(\tau)) \quad &\overset{\mathrm{def}}{=} \quad \forall x, h, h'.\, (\forall l, \mathsf{ReachFromIn}\ l\ x\ h \to h[l] = h'[l]) \to P\, x\, h \to P\, x\, h' \\
\mathsf{precise}_\tau\ (R : hrel(\tau)) \quad &\overset{\mathrm{def}}{=} \quad \forall x, x', h, h', \overline{h}, \overline{h'}.\, (\forall l, \mathsf{ReachFromIn}\ l\ x\ h \to h[l] = h'[l]) \to \\
&\qquad\quad (\forall l, \mathsf{ReachFromIn}\ l\ x'\ \overline{h} \to \overline{h}[l] = \overline{h'}[l]) \to \\
&\qquad\quad R\, x\, x'\, h\, \overline{h} \to R\, x\, x'\, h'\, \overline{h'}
\end{aligned}
$$

$$
\mathsf{closed}_\tau\ (R : hrel(\tau)) \quad \overset{\mathrm{def}}{=} \quad
\begin{cases}
\mathsf{True} & \text{if } \tau \in \{\mathtt{nat}, \mathtt{bool}, \mathtt{unit}, \Pi x : \tau' \to \tau'', \tau \xrightarrow{\mathcal{M}} \tau', \mathsf{True}, \mathsf{False}, \mathsf{heap}, \_ = \_\} \\
\mathsf{closed}\ R.1 \wedge \mathsf{closed}\ R.2 & \text{if } \tau = (\sigma \ast \sigma') \\
\mathsf{closed}\ R' \wedge \forall l, h, h'.\, R'\ h[l]\ h'[l]\ h\ h' \to R\, l\, l\, h\, h' & \text{if } \tau = \mathsf{ref}\{\tau' \mid P'\}[R', G']
\end{cases}
$$

$$
[R, G] \gg \tau \quad \overset{\mathrm{def}}{=} \quad
\begin{cases}
\tau & \text{if } \tau \in \{\mathtt{nat}, \mathtt{bool}, \mathtt{unit}, \Pi x : \tau' \to \tau'', \tau \xrightarrow{\mathcal{M}} \tau', \mathsf{True}, \mathsf{False}, \mathsf{heap}, \_ = \_\} \\
([R.1, G.1] \gg \sigma \ast [R.2, G.2] \gg \gamma) & \text{if } \tau = (\sigma \ast \gamma) \\
\mathsf{ref}\{\sigma \mid P\}[R',\ \begin{subarray}{l}(\lambda a, a', h, h'.\, G'\, a\, a'\, h\, h' \wedge \\ (\forall l, h[l] = a \to h'[l] = a' \to G\, l\, l\, h\, h'))\end{subarray}\ ] & \text{if } \tau = \mathsf{ref}\{\sigma \mid P\}[R', G']
\end{cases}
$$

where for $R : hrel(\sigma \ast \sigma')$:

$$R.1 \overset{\mathrm{def}}{=} \lambda x, x' : \sigma.\, \lambda h, h' : \mathsf{heap}.\, \forall y : \sigma'.\, R\, (x, y)\, (x', y)\, h\, h' \qquad \text{and} \qquad R.2 \overset{\mathrm{def}}{=} \lambda y, y' : \sigma.\, \lambda h, h' : \mathsf{heap}.\, \forall x : \sigma'.\, R\, (x, y)\, (x, y')\, h\, h'$$

Figure 5. Selected auxiliary judgments and predicates

$\mathcal{P}[\![e]\!]$ = e — for atomic primitives (nat, n, Prop, etc.) (includes treating = as regular term)

$\mathcal{P}[\![e_1\, e_2]\!]$ = $\mathcal{P}[\![e_1]\!]\ \mathcal{P}[\![e_2]\!]$

$\mathcal{P}[\![(e_1, e_2)]\!]$ = $(\mathcal{P}[\![e_1]\!], \mathcal{P}[\![e_2]\!])$

$\mathcal{P}[\![(e_1 * e_2)]\!]$ = $(\mathcal{P}[\![e_1]\!] * \mathcal{P}[\![e_2]\!])$

$\mathcal{P}[\![\Pi x : \tau \to \tau]\!]$ = $\Pi x : \mathcal{P}[\![\tau]\!] \to \mathcal{P}[\![\tau']\!]$

$\mathcal{P}[\![\text{ref}\{A \mid P\}[R, G]]\!]$ = $\text{ref}\{\mathcal{P}[\![A]\!] \mid \mathcal{P}[\![P]\!]\}[\mathcal{P}[\![R]\!], \mathcal{P}[\![G]\!]]$ — Not shown: translating and passing well-formedness proofs

$\mathcal{P}[\![\tau \xrightarrow{\mathcal{M}} \tau']\!]$ = $I[\![\tau \xrightarrow{\mathcal{M}} \tau']\!]^0$

$\mathcal{P}[\![h[e]]\!]$ = heap_lookup $\mathcal{P}[\![\tau]\!]\ \mathcal{P}[\![P]\!]\ \mathcal{P}[\![R]\!]\ \mathcal{P}[\![G]\!]\ h\ \mathcal{P}[\![e]\!]$ — where e is annotated as $\text{ref}\{\tau \mid P\}[R, G]$

$\mathcal{P}[\![!e]\!]$ = deref $\mathcal{P}[\![\tau]\!]\ \mathcal{P}[\![P]\!]\ \mathcal{P}[\![R]\!]\ \mathcal{P}[\![G]\!]\ \mathcal{P}[\![\tau']\!]\ \_\ \mathcal{P}[\![e]\!]$ — where $\tau/P/R/G$ are the ref type, and τ' is the folded result, _ is folding proof

$\mathcal{P}[\![x]\!]$ = x

$\mathcal{P}[\![\lambda x : \tau. e]\!]$ = $\lambda x : \mathcal{P}[\![\tau]\!]. \mathcal{P}[\![e]\!]$

$\mathcal{P}[\![\lambda_{\mathcal{M}} x : \tau. e]\!]$ = $I[\![\lambda_{\mathcal{M}} x : \tau. e]\!]$

$I[\![e]\!]^\Gamma$ = liftAs $\Gamma\, v\, \tau\, e$ — for atomic value primitives (n, b, (), etc.) of atomic value type τ (nat, bool, etc.)

$I[\![e_1\, e_2]\!]^\Gamma$ = $I[\![e_1]\!]^\Gamma \gg= (\lambda x : \mathcal{V}.\ I[\![e_2]\!]^\Gamma \gg= (\lambda y : \mathcal{V}.\ \text{m-app}\ x\ y))$ — for fresh variables x and y (both pure and impure e_i: Π-E and $\Pi\Pi$-E)

$I[\![\lambda_{\mathcal{M}} x : \tau. e]\!]^\Gamma$ = capture Γ ($\lambda env : E\Gamma.$ liftAs $\varepsilon\, v$ ($I[\![\tau]\!] \to I[\![\tau']\!]$) — for fresh variable v (note capture of bound variables)
$(\lambda x : \mathcal{V}.\ \text{inject}\ env \gg= I[\![e]\!]^{\Gamma, x : \mathcal{P}[\![\tau]\!]} \gg= (\text{dropvar}\ x)))$

$I[\![\tau \xrightarrow{\mathcal{M}} \tau']\!]^\Gamma$ = $\mathcal{M}\, \Gamma\, r\, (\Gamma, r : (\Pi x : I[\![\tau']\!]^{x : \mathcal{P}[\![\tau]\!]}))$ — for fresh r (note Γ's values are captured, and not required to invoke the result)
...for establishing P', fresh result variable v

$I[\![x \leftarrow e]\!]^\Gamma$ = heap-write $\Gamma\, v\, x\, I[\![\tau]\!]\, \_\, I[\![e]\!] \ldots$ — guarantee proof $\_ : \forall s : E\, \Gamma \forall h, h'.\ h' = \text{runM}\, s\, I[\![e]\!]\, h \to \mathcal{P}[\![P]\!]\, (!x)\, h \to \mathcal{P}[\![G]\!]\, (!x)\, h \to \mathcal{P}[\![G]\!]\, (!x)\, h'[x]\, h\, h'$

$I[\![\text{swap}(x, e)]\!]^\Gamma$ = swap $\Gamma\, v\, x\, \_\, I[\![e]\!]^\Gamma \ldots$ — for guarantee and predicate update proofs as above, fresh result variable v

$I[\![\text{alloc}_{\tau, P, R, G}\, e]\!]^\Gamma$ = alloc $\Gamma\, v\, \mathcal{P}[\![\tau]\!]\, \mathcal{P}[\![P]\!]\, \mathcal{P}[\![R]\!]\, \mathcal{P}[\![G]\!]\, I[\![e]\!]^\Gamma$ — for fresh variable v, _ : $\forall h.\ \mathcal{P}[\![P]\!]\, v$ (getResult $I[\![e]\!]$) h

$I[\![x]\!]^\Gamma$ = return $\Gamma\, x\, \Gamma$

$I[\![e]\!]^\Gamma$ = $I[\![e]\!]^\Gamma \gg= (\text{derefAs}\ \mathcal{P}[\![\tau]\!]\, \mathcal{P}[\![P]\!]\, \mathcal{P}[\![R]\!]\, \mathcal{P}[\![G]\!]\, \mathcal{P}[\![\tau']\!]\, \_\, v)$ — for fresh v, folding and reflexivity proofs

$I[\![e]\!]^\Gamma$ = liftAs $\Gamma\, v\, I[\![\tau]\!]\, \mathcal{P}[\![e]\!]$ — for fresh v, pure value expression e of type τ

$I[\![e]\!]^\Gamma$ = $\mathcal{M}\, \Gamma\, v\, (\Gamma, v : I[\![e]\!])$ — for fresh v, pure type expression e (which may contain impure types)

$I[\![\text{consume}\, x]\!]^\Gamma$ = var_rename $\Gamma\, x\, v$ — for fresh $\mathcal{V}\, v$

$I[\![\text{split}\, x \ldots]\!]^\Gamma$ = var_split $\Gamma\, x\, v \ldots$ — for fresh $\mathcal{V}\, v, \ldots$

Pure recursors and type constructors (such as =) are translated in the obvious way (recursively translating all types and arguments). Explicit conversions are translated in the natural way as well, recursively translating proof obligations from P$-$ $\Rightarrow$, R$-$ $\subset$, and G$-$ $\subset$.

Figure 6. Translating RGREF types into CC types.

$$(\lambda r : \mathsf{ref}\{\mathtt{nat} \mid \dots\}[\dots,\dots].\,(\lambda pf :!r =!r.\,(\lambda u : \mathtt{unit}.\,(\lambda pf2 :!r =!r.\,!r)\mathcal{R}_{\mathtt{bool}}(\mathsf{true}, pf, \mathsf{refl}(!r)))) \;\mathsf{refl}(!r)) \;\ell$$
$$\rightarrow \quad (\lambda pf :!\ell =!\ell.\,(\lambda u : \mathtt{unit}.\,(\lambda pf2 :!\ell =!\ell.\,!\ell)\mathcal{R}_{\mathtt{bool}}(\mathsf{true}, pf, \mathsf{refl}(!\ell)))) \;\mathsf{refl}(!\ell)$$
$$\rightarrow \quad (\lambda pf :!\ell =!\ell.\,(\lambda u : \mathtt{unit}.\,(\lambda pf2 :!\ell =!\ell.\,!\ell)\mathcal{R}_{\mathtt{bool}}(\mathsf{true}, pf, \mathsf{refl}(!\ell)))) \;\mathsf{refl}(\langle!\ell\rangle 3)$$

Figure 7. The simplest term we can imagine that observes the multiple-dereference reduction

---

$$e ::= \dots \mid \langle!\ell\rangle e \mid \mathsf{fold}\; e\; e\; e$$

$$v ::= \mathsf{true} \mid \mathsf{tt} \mid n \mid b \mid \mathsf{Type} \mid \mathsf{Prop} \mid \Pi x : \tau \to \tau' \mid \lambda x : \tau.\,e \;(\text{where no annotated expressions appear in } e) \mid (\lambda_{\mathcal{M}} x : \tau.\,e) \;(\text{where } FV(e) = \emptyset)$$

$$H : \mathsf{loc} \rightharpoonup v \qquad \Sigma : \mathsf{loc} \rightharpoonup \tau$$

$$\boxed{H;\Sigma;\Gamma \vdash M : N \;(\text{extends } \Gamma \vdash M : N)} \qquad \frac{H;\Sigma;\Gamma \vdash !\ell : \tau' \qquad H(\ell) = v \qquad \varepsilon \vdash v : \tau \qquad [R,G] \gg \tau = \tau' \qquad v' =_\beta (\mathsf{fold}\; R\; G\; v)}{H;\Sigma;\Gamma \vdash \langle!\ell_{\tau,P,R,G}\rangle v' : \tau'}$$

$$\frac{\Sigma(\ell) = \tau}{H;\Sigma;\Gamma \vdash \ell_{\tau,P,R,G} : \mathsf{ref}\{\tau \mid P\}[R,G]}$$

$$\boxed{H;\Gamma \vdash M \rightsquigarrow N \;(\text{extends } \Gamma \vdash M \rightsquigarrow N)} \qquad \frac{H(\ell) = v \qquad v' =_\beta (\mathsf{fold}\; R\; G\; v)}{H;\Gamma \vdash \langle!\ell_{\tau,P,R,G}\rangle v' \rightsquigarrow !\ell_{\tau,P,R,G}}$$

$$\boxed{H;e \to e'} \quad \frac{\beta\lambda}{H;(\lambda x : \tau.\,e)\,v \to e[x/v]} \qquad \frac{H(\ell) = v}{H;!\ell_{\tau,P,R,G} \to \langle!\ell_{\tau,P,R,G}\rangle(\mathsf{fold}\; R\; G\; v)} \qquad \frac{H;e \to e'}{H;\langle!\ell_{\tau,P,R,G}\rangle e \to \langle!\ell_{\tau,P,R,G}\rangle e'}$$

$$\frac{v \in \{\mathtt{nat},\mathtt{bool},\mathtt{unit},(\lambda x.e),(\lambda_{\mathcal{M}} x.e),\mathsf{true},\mathsf{refl}(\_)\}}{H;(\mathsf{fold}\; R\; G\; v) \to v} \qquad \frac{}{H;(\mathsf{fold}\; R\; G\; (v_1,v_2) \to (\mathsf{fold\; havoc}\; G.1\; v_1, \mathsf{fold\; havoc}\; G.2\; v_2)}$$

$$\frac{}{H;(\mathsf{fold}\; R\; G\; (\ell_{\tau,P,R',G'})) \to \ell_{\tau,P,R',(\lambda a,a',h,h'.G'\,a\,a'\,h\,h' \wedge (\forall l.h[l]=a \to h'[l]=a' \to G\,l\,l\,h\,h'))}}$$

$$\boxed{S;H;e \to S';H';e'} \quad \frac{H;e \to e' \qquad \langle!\ell\rangle a \in e' \qquad e'' = \mathsf{heap\_indep}(e')}{S;H;\mathsf{pure}\; e \to S;H;\mathsf{pure}\; e''[\langle!\ell\rangle a/a]} \qquad \frac{H;e \to e' \qquad \forall \ell,a.\,\langle!\ell\rangle a \notin e'}{S;H;\mathsf{pure}\; e \to S;H;\mathsf{pure}\; e'} \qquad \frac{}{S;H;\mathsf{pure}\; v \to S;H;v}$$

$$\frac{y \in FV(e)}{S;H;(\lambda_{\mathcal{M}} x : \tau.\,e) \to S;H;(\lambda_{\mathcal{M}} x : \tau.\,e[y/S(y)]} \qquad \frac{FV(e) = \emptyset \qquad y \;\mathsf{fresh}}{S;H;(\lambda_{\mathcal{M}} x : \tau.\,e)\,v \to S[y \mapsto v];H;e[x/y]} \qquad \frac{S(x) = \ell \qquad \ell \in \mathsf{dom}(H)}{S;H;\; x \leftarrow v \to S;H[\ell \mapsto v];\mathsf{tt}}$$

$$\frac{S(x) = \ell \qquad H(\ell) = v'}{S;H;\mathsf{swap}(x,v) \to S;H[\ell \mapsto v];v'} \qquad \frac{\ell \notin \mathsf{dom}(H)}{S;H;\mathsf{alloc}\; v \to S;H[\ell \mapsto v];\mathsf{tt}} \qquad \frac{}{S;H;\mathsf{consume}\; x \to S \setminus x;H;S(x)}$$

$$\frac{S(x) = v}{S;H;\mathsf{split}\; x\; \tau\; \tau' \to S[x \mapsto v \;\mathsf{as}\; \tau];H;v \;\mathsf{as}\; \tau'} \qquad \frac{}{S;H;\mathsf{pureApp}\; (\lambda x : \tau.\,e)\,v \to S;H;\mathsf{pure}\; ((\lambda x : \tau.\,e)\,v)}$$

$$\frac{H(\ell) = v}{S;H;\,!\ell_{\tau,P,R,G} \to S;H;\mathsf{pure}\; \langle!\ell_{\tau,P,R,G}\rangle(\mathsf{fold}\; R\; G\; v)}$$

Figure 8. Values, typing of runtime values, and main operational semantics for RGREF.

$$\boxed{H; e \to e'} \qquad \frac{H; e_1 \to e_1'}{H; e_1\, e_2 \to e_1'\, e_2} \qquad \frac{H; e_2 \to e_2'}{H; v\, e_2 \to v\, e_2'} \qquad \frac{H; e_1 \to e_1'}{H; (e_1, e_2) \to (e_1', e_2)} \qquad \frac{H; e_2 \to e_2'}{H; (v, e_2) \to (v, e_2')} \qquad \frac{H; e \to e'}{H; !e \to !e'}$$

$$\frac{H; e \to e'}{H; \mathsf{refl}\, e \to \mathsf{refl}\, e'}$$

$$\boxed{S; H; e \to S'; H'; e'} \qquad \frac{S; H; e_1 \to S'; H'; e_1'}{S; H; e_1\, e_2 \to S'; H'; e_1'\, e_2} \qquad \frac{S; H; e_2 \to S'; H'; e_2'}{S; H; v\, e_2 \to S'; H'; v\, e_2'} \qquad \frac{S; H; e \to S'; H'; e'}{S; H; x \leftarrow e \to S'; H'; x \leftarrow e'}$$

$$\frac{S; H; e \to S'; H'; e'}{S; H; \mathsf{swap}(x, e) \to S'; H'; \mathsf{swap}(x, e')} \qquad \frac{S; H; e \to S'; H'; e'}{S; H; \mathsf{alloc}\, e \to S'; H'; \mathsf{alloc}\, e'} \qquad \frac{S; H; e_1 \to S'; H'; e_1'}{S; H; \mathsf{pureApp}\, e_1\, e_2 \to S'; H'; \mathsf{pureApp}\, e_1'\, e_2}$$

$$\frac{S; H; e_2 \to S'; H'; e_2'}{S; H; \mathsf{pureApp}\, v\, e_2 \to S'; H'; \mathsf{pureApp}\, v\, e_2'}$$

Figure 9. Structural / context operational semantics for RGREF.

$$\boxed{\vdash S; H; e : \Gamma; \Sigma; \tau \Rightarrow \Gamma'}$$

$$\frac{\begin{array}{c} \vdash \Sigma \qquad \vdash H : \Sigma \qquad \vdash \Gamma \qquad H; \Sigma \vdash S : \Gamma \qquad H; \Sigma; \Gamma \vdash e : \tau \Rightarrow \Gamma' \qquad \forall \tau, \ell, P, R, G. \ell_{\tau, P, R, G} \in S; H; e \implies [\![P]\!]\, H(\ell)\, H \\ \forall \tau, \ell, P, P', R, R', G, G'. \ell_{\tau, P, R, G} \in S; H; e \wedge \ell_{\tau, P', R', G'} \in S; H; e \wedge \text{the references are in different locations} \Rightarrow (G \implies R' \wedge G' \implies R) \end{array}}{\vdash S; H; e : \Gamma; \Sigma; \tau \Rightarrow \Gamma'}$$

$$\boxed{H; \Sigma \vdash S : \Gamma} \qquad\qquad \overline{H; \Sigma \vdash S : \varepsilon} \qquad\qquad \frac{H; \Sigma \vdash S/x : \Gamma \qquad H; \Sigma; \varepsilon \vdash S(x) : \tau}{H; \Sigma \vdash S : \Gamma, x : \tau}$$

$$\boxed{\vdash \Sigma} \qquad \frac{\forall \ell \in \mathsf{dom}(\Sigma). \emptyset; \emptyset; \varepsilon \vdash \Sigma(\ell) : \mathsf{Prop}}{\vdash \Sigma} \qquad\qquad \boxed{\vdash H : \Sigma} \qquad \frac{\mathsf{dom}(H) = \mathsf{dom}(\Sigma) \qquad \forall \ell \in \mathsf{dom}(H). \emptyset; \Sigma; \varepsilon \vdash H(\ell) : \Sigma(\ell)}{\vdash H : \Sigma}$$

Figure 10. Dynamic state typing.