ConSORT: Context- and Flow-Sensitive Ownership Refinement Types for Imperative Programs

John Toman[1], Ren Siqi[1], Kohei Suenaga[1][0000−0002−7466−8789], Atsushi Igarashi[1][0000−0002−5143−9764], and Naoki Kobayashi[2]

[1] Kyoto University, Japan,
{jtoman,shiki,ksuenaga,igarashi}@fos.kuis.kyoto-u.ac.jp
[2] The University of Tokyo, Japan, koba@is.s.u-tokyo.ac.jp

Abstract. We present ConSORT, a type system for safety verification in the presence of mutability and aliasing. Mutability requires *strong updates* to model changing invariants during program execution, but aliasing between pointers makes it difficult to determine which invariants must be updated in response to mutation. Our type system addresses this difficulty with a novel combination of refinement types and fractional ownership types. Fractional ownership types provide flow-sensitive and precise aliasing information for reference variables. ConSORT interprets this ownership information to soundly handle strong updates of potentially aliased references. We have proved ConSORT sound and implemented a prototype, fully automated inference tool. We evaluated our tool and found it verifies non-trivial programs including data structure implementations.

Keywords: refinement types, mutable references, aliasing, strong updates, fractional ownerships, program verification, type systems

1 Introduction

Driven by the increasing power of automated theorem provers and recent high-profile software failures, fully automated program verification has seen a surge of interest in recent years [4, 9, 14, 28, 34, 58]. In particular, *refinement types* [8, 20, 23, 57], which refine base types with logical predicates, have been shown to be a practical approach for program verification that are amenable to (sometimes full) automation [40, 53, 54, 55]. Despite promising advances [25, 29, 39], the sound and precise application of refinement types (and program verification in general) in settings with mutability and aliasing (e.g., Java, Ruby, etc.) remains difficult.

One of the major challenges is how to precisely and soundly support *strong updates* for the invariants on memory cells. In a setting with mutability, a single invariant may not necessarily hold throughout the lifetime of a memory cell; while the program mutates the memory the invariant may change or evolve. To model these changes, a program verifier must support different, incompatible invariants which hold at different points during program execution. Further, precise program verification requires supporting different invariants on distinct pieces of memory.

```
1 mk(n) { mkref n }

3 let p = mk(3) in
4 let q = mk(5) in
5 p := *p + 1;
6 q := *q + 1;
7 assert(*p = 4);
```

Fig. 1. Example demonstrating the difficulty of effecting strong updates in the presence of aliasing. The function `mk` is bound in the program from lines 3 to 7; its body is given within the braces.

```
1  loop(a, b) {
2    let aold = *a in
3    b := *b + 1;
4    a := *a + 1;
5    assert(*a = aold + 1);
6    if ⋆ then
7      loop(b, mkref ⋆)
8    else
9      loop(b,a)
10 }
11 loop(mkref ⋆, mkref ⋆)
```

Fig. 2. Example with non-trivial aliasing behavior.

One solution is to use refinement types on the static program names (i.e., variables) which point to a memory location. This approach can model evolving invariants while tracking distinct invariants for each memory cell. For example, consider the (contrived) example in Figure 1. This program is written in an ML-like language with mutable references; references are updated with `:=` and allocated with **mkref**. Variable `p` can initially be given the type $\{\nu : \mathbf{int} \mid \nu = 3\}\,\mathbf{ref}$, indicating it is a reference to the integer 3. Similarly, `q` can be given the type $\{\nu : \mathbf{int} \mid \nu = 5\}\,\mathbf{ref}$. We can model the mutation of `p`'s memory on line 5 by strongly updating `p`'s type to $\{\nu : \mathbf{int} \mid \nu = 4\}\,\mathbf{ref}$.

Unfortunately, the precise application of this technique is confounded by the existence of unrestricted aliasing. In general, updating the type of just the mutated reference is insufficient: due to aliasing, other variables may point to the mutated memory and their refinements must be updated as well. However, in the presence of conditional, *may* aliasing, it is impossible to strongly update the refinements on all possible aliases; given the static uncertainty about whether a variable points to the mutated memory, that variable's refinement may only be *weakly updated*. For example, suppose we used a simple alias analysis that imprecisely (but soundly) concluded all references allocated at the same program point *might* alias. Variables `p` and `q` share the allocation site on line 1, so on line 5 we would have to weakly update `q`'s type to $\{\nu : \mathbf{int} \mid \nu = 4 \vee \nu = 5\}$, indicating it may hold either 4 *or* 5. Under this same imprecise aliasing assumption, we would also have to weakly update `p`'s type on line 6, preventing the verification of the example program.

Given the precision loss associated with weak updates, it is critical that verification techniques built upon refinement types use precise aliasing information and avoid spuriously applied weak updates. Although it is relatively simple to conclude that `p` and `q` do not alias in Figure 1, consider the example in Figure 2. (In this example, $\star$ represents non-deterministic values.) Verifying this program requires proving `a` and `b` never alias at the writes on lines 3 and 4. In fact, `a` and `b` *may* point to the same memory location, but only in different invocations of `loop`; this pattern may confound even sophisticated symbolic alias analyses.

Additionally, `a` and `b` share an allocation site on line 7, so an approach based on the simple alias analysis described above will also fail on this example. This obligation *can* be discharged with existing techniques [46, 47], but requires an expensive, on-demand, interprocedural, flow-sensitive alias analysis.

This paper presents ConSORT (CONtext Sensitive Ownership Refinement Types), a type system for the automated verification of program safety for imperative languages with mutability and aliasing. ConSORT is built upon the novel combination of refinement types and fractional ownership types [48, 49]. Fractional ownership types extend pointer types with a rational number in the range $[0, 1]$ called an *ownership*. These ownerships encapsulate the permission of the reference; only references with ownership 1 may be used for mutation. Fractional ownership types also obey the following key invariant: any references with a mutable alias must have ownership 0. Thus, any reference with non-zero ownership *cannot* be an alias of a reference with ownership 1. In other words, ownerships encode precise aliasing information in the form of *must-not* aliasing relationships.

To understand the benefit of this approach, let us return to Figure 1. As `mk` returns a freshly allocated reference with no aliases, its type indicates it returns a reference with ownership 1. Thus our type system can initially give `p` and `q` types $\{\nu : \mathbf{int} \mid \nu = 3\}\,\mathbf{ref}^1$ and $\{\nu : \mathbf{int} \mid \nu = 5\}\,\mathbf{ref}^1$ respectively. The ownership 1 on the reference type constructor **ref** indicates both pointers hold "exclusive" ownership of the pointed to reference cell; from the invariant of fractional ownership types `p` and `q` must not alias. The types of both references can be strongly updated *without* requiring spurious weak updates. As a result, at the assertion statement on line 7, `p` has type $\{\nu : \mathbf{int} \mid \nu = 4\}\,\mathbf{ref}^1$ expressing the required invariant.

Our type system can also verify the example in Figure 2 *without* expensive side analyses. As `a` and `b` are both mutated, they must both have ownership 1; i.e., they cannot alias. This pre-condition is satisfied by all invocations of `loop`; on line 7, `b` has ownership 1 (from the argument type), and the newly allocated reference must also have ownership 1. Similarly, both arguments on line 9 have ownership 1 (from the assumed ownership on the argument types).

Ownerships behave linearly; they cannot be duplicated, only *split* when aliases are created. This linear behavior preserves the critical ownership invariant. For example, if we replace line 9 in Figure 2 with `loop(b,b)`, the program becomes ill-typed; there is no way to divide `b`'s ownership of 1 to into *two* ownerships of 1.

Ownerships also obviate updating the refinements of all aliases at mutation. ConSORT ensures that only the trivial refinement $\top$ is used within the type of all mutably-aliased references; i.e., references with 0 ownership. When memory is mutated through a reference with ownership 1, ConSORT simply updates the refinement of the mutated reference variable. From the soundness of ownership types and the above condition enforced by ConSORT, the types of any aliases must only use the refinement $\top$, and already soundly describe all possible contents.[3]

ConSORT is also *context-sensitive*, and can use different summaries of function behavior at different points in the program. For example, consider the variant

[3] This assumption holds only if updates do not change simple types, a condition our type-system enforces.

```
1 get(p) { *p }

3 let p = mkref 3 in
4 let q = mkref 5 in
5 p := get(p) + 1;
6 q := get(q) + 1;
7 assert(*p = 4);
8 assert(*q = 6);
```

Fig. 3. Context-sensitivity example

of Figure 1 shown in Figure 3. The function `get` returns the contents of its argument, and is called on lines 5 and 6. To precisely verify this program, on line 5 `get` must be typed as a function that takes a reference to 3 and returns 3. Similarly, on line 6 `get` must be typed as a function that takes a reference to 5 and returns 5. Our type system can give `get` a function type that distinguishes between these two calling contexts and selects the appropriate summary of `get`'s behavior.

We have formalized ConSORT as a type system for a small imperative calculus and proved the system is sound: i.e., a well-typed program never encounters assertion failures during execution. We have implemented a prototype type inference tool targeting this imperative language and found it can automatically verify several non-trivial programs, including sorted lists and an array list data structure.

The rest of this paper is organized as follows. Section 2 defines the imperative language targeted by ConSORT and its semantics. Section 3 defines our type system and states our soundness theorem. Section 4 sketches our implementation's inference algorithm and its current limitations. Section 5 describes an evaluation of our prototype, Section 6 outlines related work, and Section 7 concludes.

2 Target Language

This section describes a simple imperative language with mutable references and first-order, recursive functions.

2.1 Syntax

We assume a set of *variables*, ranged over by $x, y, z, \ldots$, a set of *function names*, ranged over by f, and a set of labels ranged over by $\ell_1, \ell_2, \ldots$. The grammar of the language is as follows.

$$
\begin{array}{rl}
d & ::= f \mapsto (x_1, \ldots, x_n)e \\
e & ::= x \mid \mathbf{let}\, x = y\, \mathbf{in}\, e \mid \mathbf{let}\, x = n\, \mathbf{in}\, e \mid \mathbf{ifz}\, x\, \mathbf{then}\, e_1\, \mathbf{else}\, e_2 \\
 & \quad \mid \mathbf{let}\, x = \mathbf{mkref}\, y\, \mathbf{in}\, e \mid \mathbf{let}\, x = *y\, \mathbf{in}\, e \mid \mathbf{let}\, x = f^{\ell}(y_1, \ldots, y_n)\, \mathbf{in}\, e \\
 & \quad \mid x := y\,;e \mid \mathbf{alias}(x = y)\,;e \mid \mathbf{alias}(x = *y)\,;e \mid \mathbf{assert}(\varphi)\,;e \mid e_1\,;e_2 \\
P & ::= \langle \{d_1, \ldots, d_n\}, e \rangle
\end{array}
$$

φ stands for a formula in propositional first-order logic over variables, integers and contexts; we discuss these formulas later in Section 3.1.

Variables are introduced by function parameters or let bindings. Like ML, the variable bindings introduced by let expressions and parameters are immutable. A mutable variable declaration like `int x = 1;` in C is achieved in our language with:

$$\mathbf{let}\, y = 1\, \mathbf{in}(\mathbf{let}\, x = \mathbf{mkref}\, y\, \mathbf{in} \ldots)$$

As a convenience, we assume all variable names introduced with let bindings and function parameters are distinct.

Unlike ML (and like C or Java) we do not allow general expressions on the right hand side of let bindings. The simplest right hand forms are a variable y or an integer literal n. **mkref** y creates a reference cell with value y, and $*y$ accesses the contents of reference y. For simplicity, we do not include an explicit null value; an extension to null will be discussed in Section 4. Function calls must occur on the right hand side of a variable binding and take the form $f^{\ell}(x_1, \ldots, x_n)$, where $x_1, \ldots, x_n$ are distinct variables and ℓ is a (unique) label. These labels are used to make our type system context-sensitive as discussed in Section 3.3.

The single base case for expressions is a single variable. If the variable expression is executed in a tail position of a function, then the value of that variable is the return value of the function, otherwise the value is ignored.

The only intraprocedural control-flow operations in our language are if statements. **ifz** checks whether the condition variable x equals zero and chooses the corresponding branch. Loops can be implemented with recursive functions and we do not include them explicitly in our formalism.

Our grammar requires that side-effecting, result-free statements, $\mathbf{assert}(\varphi)$ $\mathbf{alias}(x = y)$, $\mathbf{alias}(x = *y)$ and assignment $x := y$ are followed by a continuation expression. We impose this requirement for technical reasons to ease our formal presentation; this requirement does not reduce expressiveness as dummy continuations can be inserted as needed. The $\mathbf{assert}(\varphi)\,;e$ form executes e if the predicate φ holds in the current state and aborts the program otherwise. $\mathbf{alias}(x = y)\,;e$ asserts a must-aliasing relationship between x and y (resp. x and $*y$) and then executes e. **alias** statements are effectively *annotations* that our type system exploits to gain added precision. $x := y\,;e$ updates the contents of the memory cell pointed to by x with the contents of y. In addition to the above continuations, our language supports general sequencing with $e_1\,;e_2$.

A program is a pair $\langle D, e\rangle$, where $D = \{d_1, \ldots, d_n\}$ is a set of first-order, mutually recursive function definitions, and e is the program entry point. A function definition d maps the function name to a tuple of argument names $x_1, \ldots, x_n$ that are bound within the function body e.

Paper Syntax In the remainder of the paper, we will write programs that are technically illegal according to our grammar, but can be easily "de-sugared" into an equivalent, valid program. For example, we will write

```
let x = mkref 4 in assert(*x = 4)
```

as syntactic sugar for:

```
let f = 4 in let x = mkref f in
let tmp = *x in assert(tmp = 4); let dummy = 0 in dummy
```

2.2 Operational Semantics

We now introduce the operational semantics for our language. We assume a finite domain of heap addresses **Addr**: we denote an arbitrary address with a.

$$\frac{}{\left\langle H, R, F : \vec{F}, x \right\rangle \longrightarrow_D \left\langle H, R, \vec{F}, F[x] \right\rangle} \text{(R-Var)} \qquad \frac{}{\left\langle H, R, F : \vec{F}, E[x\,;e] \right\rangle \longrightarrow_D \left\langle H, R, \vec{F}, E[e] \right\rangle} \text{(R-Seq)}$$

$$\frac{x' \notin dom(R)}{\left\langle H, R, \vec{F}, E[\mathbf{let}\ x = y\ \mathbf{in}\ e] \right\rangle \longrightarrow_D \left\langle H, R\{x' \mapsto R(y)\}, \vec{F}, E[[x'/x]e] \right\rangle} \text{(R-Let)}$$

$$\frac{x' \notin dom(R)}{\left\langle H, R, \vec{F}, E[\mathbf{let}\ x = n\ \mathbf{in}\ e] \right\rangle \longrightarrow_D \left\langle H, R\{x' \mapsto n\}, \vec{F}, E[[x'/x]e] \right\rangle} \text{(R-LetInt)}$$

$$\frac{R(x) = 0}{\left\langle H, R, \vec{F}, E[\mathbf{ifz}\ x\ \mathbf{then}\ e_1\ \mathbf{else}\ e_2] \right\rangle \longrightarrow_D \left\langle H, R, \vec{F}, E[e_1] \right\rangle} \text{(R-IfTrue)}$$

$$\frac{R(x) \neq 0}{\left\langle H, R, \vec{F}, E[\mathbf{ifz}\ x\ \mathbf{then}\ e_1\ \mathbf{else}\ e_2] \right\rangle \longrightarrow_D \left\langle H, R, \vec{F}, E[e_2] \right\rangle} \text{(R-IfFalse)}$$

$$\frac{a \notin dom(H) \qquad x' \notin dom(R)}{\left\langle H, R, \vec{F}, E[\mathbf{let}\ x = \mathbf{mkref}\ y\ \mathbf{in}\ e] \right\rangle \longrightarrow_D \left\langle H\{a \mapsto R(y)\}, R\{x' \mapsto a\}, \vec{F}, E[[x'/x]e] \right\rangle} \text{(R-MkRef)}$$

$$\frac{R(y) = a \qquad H(a) = v \qquad x' \notin dom(R)}{\left\langle H, R, \vec{F}, E[\mathbf{let}\ x = *y\ \mathbf{in}\ e] \right\rangle \longrightarrow_D \left\langle H, R\{x' \mapsto v\}, \vec{F}, E[[x'/x]e] \right\rangle} \text{(R-Deref)}$$

Fig. 4. Transition Rules (1).

A runtime state is represented by a configuration $\left\langle H, R, \vec{F}, e \right\rangle$, which consists of a heap, register file, stack, and currently reducing expression respectively. The register file maps variables to runtime values v, which are either integers n or addresses a. The heap maps a finite subset of addresses to runtime values. The runtime stack represents pending function calls as a sequence of return contexts, which we describe below. While the final configuration component is an expression, the rewriting rules are defined in terms of $E[e]$, which is an evaluation context E and redex e, as is standard. The grammar for evaluation contexts is defined by: $E ::= E'\,; e \mid []$.

Our operational semantics is given in Figures 4 and 5. We write $dom(H)$ to indicate the domain of a function and $H\{a \mapsto v\}$ where $a \notin dom(H)$ to denote a map which takes all values in $dom(H)$ to their values in H and which additionally takes a to v. We will write $H\{a \hookleftarrow v\}$ where $a \in dom(H)$ to denote a map equivalent to H except that a takes value v. We use similar notation for $dom(R)$ and $R\{x \mapsto v\}$. We also write $\emptyset$ for the empty register file and heap. The step relation $\longrightarrow_D$ is parameterized by a set of function definitions D; a program $\langle D, e\rangle$ is executed by stepping the initial configuration $\langle \emptyset, \emptyset, \cdot, e\rangle$ according to $\longrightarrow_D$. The semantics is mostly standard; we highlight some important points below.

Return contexts F take the form $E[\mathbf{let}\ y = []^\ell\ \mathbf{in}\ e]$. A return context represents a pending function call with label ℓ, and indicates that the return value of the callee should be bound to y in e within the larger execution context E. The call stack $\vec{F}$ is a sequence of these contexts, with the first such return context representing the most recent function call. The stack grows at function calls as described by rule R-Call. For a call $E[\mathbf{let}\ x = f^\ell(y_1, \ldots, y_n)\ \mathbf{in}\ e]$ where f is defined as $(x_1, \ldots, x_n)e'$, the return context $E[\mathbf{let}\ y = []^\ell\ \mathbf{in}\ e]$ is prepended onto

$$\frac{f \mapsto (x_1, .., x_n)e \in D}{\left\langle H, R, \vec{F}, E[\mathbf{let}\, x = f^{\ell}(y_1, \ldots, y_n)\, \mathbf{in}\, e']\right\rangle \longrightarrow_D \left\langle H, R, E[\mathbf{let}\, x = []^{\ell}\, \mathbf{in}\, e'] : \vec{F}, [y_1/x_1] \cdots [y_n/x_n]e\right\rangle} \text{(R-Call)}$$

$$\frac{R(x) = a \qquad a \in dom(H)}{\left\langle H, R, \vec{F}, E[x := y\,; e]\right\rangle \longrightarrow_D \left\langle H\{a \hookleftarrow R(y)\}, R, \vec{F}, E[e]\right\rangle} \text{(R-Assign)}$$

$$\frac{R(x) = R(y)}{\left\langle H, R, \vec{F}, E[\mathbf{alias}(x = y)\,; e]\right\rangle \longrightarrow_D \left\langle H, R, \vec{F}, E[e]\right\rangle} \text{(R-Alias)}$$

$$\frac{R(y) = a \qquad H(a) = R(x)}{\left\langle H, R, \vec{F}, E[\mathbf{alias}(x = *y)\,; e]\right\rangle \longrightarrow_D \left\langle H, R, \vec{F}, E[e]\right\rangle} \text{(R-AliasPtr)}$$

$$\frac{R(x) \neq R(y)}{\left\langle H, R, \vec{F}, E[\mathbf{alias}(x = y)\,; e]\right\rangle \longrightarrow_D \mathbf{AliasFail}} \text{(R-AliasFail)}$$

$$\frac{R(x) \neq H(R(y))}{\left\langle H, R, \vec{F}, E[\mathbf{alias}(x = *y)\,; e]\right\rangle \longrightarrow_D \mathbf{AliasFail}} \text{(R-AliasPtrFail)}$$

$$\frac{\models [R]\,\varphi}{\left\langle H, R, \vec{F}, E[\mathbf{assert}(\varphi)\,; e]\right\rangle \longrightarrow_D \left\langle H, R, \vec{F}, E[e]\right\rangle} \text{(R-Assert)}$$

$$\frac{\not\models [R]\,\varphi}{\left\langle H, R, \vec{F}, E[\mathbf{assert}(\varphi)\,; e]\right\rangle \longrightarrow_D \mathbf{AssertFail}} \text{(R-AssertFail)}$$

Fig. 5. Transition Rules (2).

the stack of the input configuration. The substitution of formal arguments for parameters in e', denoted by $[y_1/x_1] \cdots [y_n/x_n]e'$, becomes the currently reducing expression in the output configuration. Function returns are handled by R-Var. Our semantics return values by name; when the currently executing function fully reduces to a single variable x, x is substituted into the return context on the top of the stack, denoted by $E[\mathbf{let}\, y = []^{\ell}\, \mathbf{in}\, e][x]$.

In the rules R-Assert we write $\models [R]\,\varphi$ to mean that the formula yielded by substituting the concrete values in R for the variables in φ is valid within some chosen logic (see Section 3.1); in R-AssertFail we write $\not\models [R]\,\varphi$ when the formula is *not* valid. The substitution operation $[R]\,\varphi$ is defined inductively as $[\emptyset]\,\varphi = \varphi, [R\{x \mapsto n\}]\,\varphi = [R]\,[n/x]\varphi, [R\{x \mapsto a\}]\,\varphi = [R]\,\varphi$. In the case of an assertion failure, the semantics steps to a distinguished configuration **AssertFail**. The goal of our type system to show that no execution of a well-typed program may reach this configuration. The **alias** form checks whether the two references actually alias; i.e., if the must-alias assertion provided by the programmer is correct. If not, our semantics steps to the distinguished **AliasFail** configuration. Our type system does *not* guarantee that **AliasFail** is unreachable; aliasing assertions are effectively trusted annotations that are assumed to hold.

In order to avoid duplicate variable names in our register file due to recursive functions, we refresh the bound variable x in a let expression to x'. Take expression $\mathbf{let}\, x = y\, \mathbf{in}\, e$ as an example; we substitute a fresh variable x' for x in e, then bind x' to the value of variable y. We assume this refreshing of variables preserves our assumption that all variable bindings introduced with let and function parameters are unique, i.e. x' does not overlap with variable names that occur in the program.

$$
\begin{array}{rll}
\text{Types} & \tau ::= & \{\nu : \mathbf{int} \mid \varphi\} \mid \tau\, \mathbf{ref}^r \\
\text{Ownership} & r \in & [0,1] \\
\text{Refinements} & \varphi ::= & \varphi_1 \vee \varphi_2 \mid \neg\varphi \mid \top \\
& \mid & \phi(\widehat{v}_1, .., \widehat{v}_n) \\
& \mid & \widehat{v}_1 = \widehat{v}_2 \\
& \mid & \mathcal{CP} \\
\text{Refinement Values} & \widehat{v} ::= & x \mid v \mid \nu
\end{array}
\qquad
\begin{array}{rll}
\text{Function Types} & \sigma ::= & \forall \lambda.\, \langle x_1 : \tau_1, \ldots, x_n : \tau_n \rangle \\
& & \to \langle x_1 : \tau'_1, \ldots, x_n : \tau'_n \mid \tau \rangle \\
\text{Context Variables} & \lambda \in & \mathbf{CVar} \\
\text{Concrete Context} & \vec{\ell} ::= & \ell : \vec{\ell} \mid \epsilon \\
\text{Pred. Context} & \mathcal{C} ::= & \ell : \mathcal{C} \mid \lambda \mid \epsilon \\
\text{Context Query} & \mathcal{CP} ::= & \vec{\ell} \subseteq \mathcal{C} \\
\text{Typing Context} & \mathcal{L} ::= & \lambda \mid \vec{\ell}
\end{array}
$$

Fig. 6. Syntax of types, refinements, and contexts

3 Typing

We now introduce a fractional ownership refinement type system that guarantees well-typed programs do not encounter assertion failures.

3.1 Types and Contexts

The syntax of types is given in Figure 6. Our type system has two type constructors: references and integers. $\tau\, \mathbf{ref}^r$ is the type of a (non-null) reference to a value of type τ. r is an ownership which is a rational number in the range $[0, 1]$. An ownership of 0 indicates a reference that cannot be written, and for which there may exist a mutable alias. By contrast, 1 indicates a pointer with exclusive ownership that can be read and written. Reference types with ownership values between these two extremes indicate a pointer that is readable but not writable, and for which no mutable aliases exist. ConSORT ensures that these invariants hold while aliases are created and destroyed during execution.

Integers are refined with a predicate φ. The language of predicates is built using the standard logical connectives of first-order logic, with (in)equality between variables and integers, and atomic predicate symbols ϕ as the basic atoms. We include a special "value" variable ν representing the value being refined by the predicate. For simplicity, we omit the connectives $\varphi_1 \wedge \varphi_2$ and $\varphi_1 \implies \varphi_2$; they can be written as derived forms using the given connectives. We do not fix a particular theory from which ϕ are drawn, provided a sound (but not necessarily complete) decision procedure exists. $\mathcal{CP}$ are context predicates, which are used for context sensitivity as explained below.

Example 1. $\{\nu : \mathbf{int} \mid \nu > 0\}$ is the type of strictly positive integers. The type of immutable references to integers exactly equal to 3 can be expressed by $\{\nu : \mathbf{int} \mid \nu = 3\}\, \mathbf{ref}^{0.5}$.

As is standard, we denote a type environment with Γ, which is a finite map from variable names to type τ. We write $\Gamma[x : \tau]$ to denote a type environment Γ such that $\Gamma(x) = \tau$ where $x \in dom(\Gamma)$, $\Gamma, x : \tau$ to indicate the extension of Γ with the type binding $x : \tau$, and $\Gamma[x \hookleftarrow \tau]$ to indicate the type environment Γ with the binding of x updated to τ. We write the empty environment as

•. The treatment of type environments as mappings instead of sequences in a dependent type system is somewhat non-standard. The standard formulation based on ordered sequences of bindings and its corresponding well-formedness condition did not easily admit variables with mutually dependent refinements as introduced by our function types (see below). We therefore use an unordered environment and relax well-formedness to ignore variable binding order.

Function Types, Contexts, and Context Polymorphism Our type system achieves context sensitivity by allowing function types to depend on where and how a function is called, i.e., the *execution context* of the function invocation. Our system represents a *concrete* execution contexts with strings of call site labels (or just "call strings"), defined by $\vec{\ell} ::= \epsilon \mid \ell : \vec{\ell}$. As is standard (e.g., [42, 43]), the string $\ell : \vec{\ell}$ abstracts an execution context where the most recent, active function call occurred at call site ℓ which itself was executed in a context abstracted by $\vec{\ell}$; ϵ is the context under which program execution begins. *Context variables*, drawn from a finite domain **CVar** and ranged over by $\lambda_1, \lambda_2, \ldots$, represent an arbitrary, unknown context.

A function type takes the form $\forall\lambda.\, \langle x_1 : \tau_1, \ldots, x_n : \tau_n \rangle \to \langle x_1 : \tau'_1, \ldots, x_n : \tau'_n \mid \tau \rangle$. The arguments of a function are an n-ary tuple of types τ_i. To model side-effects on arguments, the function type includes the same number of *output types* τ'_i. In addition, function types have a direct return type τ. The argument and output types are given names: refinements within the function type may refer to these names. Function types in our language are context polymorphic, expressed by universal quantification "$\forall\lambda$." over a context variable. Intuitively, this context variable represents the many different execution contexts under which a function may be called.

Argument and return types may depend on this context variable by including *context query predicates* in their refinements. A context query predicate $\mathcal{CP}$ usually takes the form $\vec{\ell} \subseteq \lambda$, and is true iff $\vec{\ell}$ is a prefix of the concrete context represented by λ. Intuitively, a refinement $\vec{\ell} \subseteq \lambda \implies \varphi$ states that φ holds in any concrete execution context with prefix $\vec{\ell}$, and provides no information in any other context. In full generality, a context query predicate may be of the form $\vec{\ell_1} \subseteq \vec{\ell_2}$ or $\vec{\ell} \subseteq \ell_1 \ldots \ell_n : \lambda$; these forms may be immediately simplified to $\top$, $\bot$ or $\vec{\ell'} \subseteq \lambda$.

Example 2. The type $\{\nu : \mathbf{int} \mid (\ell_1 \subseteq \lambda \implies \nu = 3) \wedge (\ell_2 \subseteq \lambda \implies \nu = 5)\}$ represents an integer that is 3 if the most recent active function call site is ℓ_1, 5 if the most recent call site is ℓ_2, and is otherwise unconstrained. This type may be used for the argument of `f` in, e.g., $\texttt{f}^{\ell_1}\texttt{(3) + f}^{\ell_2}\texttt{(5)}$.

As types in our type system may contain context variables, our typing judgment (introduced below) includes a typing context $\mathcal{L}$, which is either a single context variable λ or a concrete context $\vec{\ell}$. This typing context represents the assumptions about the execution context of the term being typed. If the typing context is a context variable λ, then no assumptions are made about the execution context of the term, although types may depend upon λ with context query predicates. Accordingly, function bodies are typed under the context variable universally quantified over in the corresponding function type; i.e., no

assumptions are made about the exact execution context of the function body. As in parametric polymorphism, consistent substitution of a concrete context $\vec{\ell}$ for a context variable λ in a typing derivation yields a valid type derivation under concrete context $\vec{\ell}$.

Remark 1. The context-sensitivity scheme described here corresponds to the standard CFA approach [43] without *a priori* call-string limiting. We chose this scheme because it can be easily encoded with equality over integer variables (see Section 4), but in principle another context-sensitivity strategy could be used instead. The important feature of our type system is the inclusion of predicates over contexts, not the specific choice for these predicates.

Function type environments are denoted with Θ and are finite maps from function names (f) to function types (σ).

Well Formedness We impose two well-formedness conditions on types: *ownership well-formedness* and *refinement well-formedness*. The ownership condition is purely syntactic: τ is ownership well-formed if $\tau = \tau' \, \mathbf{ref}^0$ implies $\tau' = \top_n$ for some n. $\top_i$ is the "maximal" type of a chain of i references, and is defined inductively as $\top_0 = \{\nu : \mathbf{int} \mid \top\}$, $\top_i = \top_{i-1} \, \mathbf{ref}^0$.

The ownership well-formedness condition ensures that aliases introduced via heap writes do not violate the invariant of ownership types *and* that refinements are consistent with updates performed through mutable aliases. Recall our ownership type invariant ensures all aliases of a mutable reference have 0 ownership. Any mutations through that alias will therefore be consistent with the "no information" $\top$ refinement required by this well-formedness condition.

Refinement well-formedness, denoted $\mathcal{L} \mid \Gamma \vdash_{WF} \varphi$, ensures that free program variables in refinement φ are bound in a type environment Γ and have integer type. It also requires that for a typing context $\mathcal{L} = \lambda$, only context query predicates over λ are used (no such predicates may be used if $\mathcal{L} = \vec{\ell}$). Notice this condition forbids refinements that refer to references. Although ownership information can signal when refinements on a mutably-aliased reference must be discarded, our current formulation provides no such information for refinements that *mention* mutably-aliased references. We therefore conservatively reject such refinements at the cost of some expressiveness in our type system.

We write $\mathcal{L} \mid \Gamma \vdash_{WF} \tau$ to indicate a well-formed type where all refinements are well-formed with respect to $\mathcal{L}$ and Γ. We write $\mathcal{L} \vdash_{WF} \Gamma$ for a type environment where all types are well-formed. A function environment is well-formed (written $\vdash_{WF} \Theta$) if, for every σ in Θ, the argument, result, and output types are well-formed with respect to each other and the context variable quantified over in σ. As the formal definition of refinement well-formedness is fairly standard, we omit it for space reasons (the full definition may be found in Appendix B).

3.2 Intraprocedural Type System

We now introduce the type system for the intraprocedural fragment of our language. Accordingly, this section focuses on the interplay of mutability and

$$\frac{}{\Theta \mid \mathcal{L} \mid \Gamma[x : \tau_1 + \tau_2] \vdash x : \tau_1 \Rightarrow \Gamma[x \hookleftarrow \tau_2]} \quad \text{(T-Var)}$$

$$\frac{\Theta \mid \mathcal{L} \mid \Gamma[y \hookleftarrow \tau_1 \wedge_y y =_{\tau_1} x], x : (\tau_2 \wedge_x x =_{\tau_2} y) \vdash e : \tau \Rightarrow \Gamma' \qquad x \notin dom(\Gamma')}{\Theta \mid \mathcal{L} \mid \Gamma[y : \tau_1 + \tau_2] \vdash \mathbf{let}\, x = y\, \mathbf{in}\, e : \tau \Rightarrow \Gamma'} \quad \text{(T-Let)}$$

$$\frac{\Theta \mid \mathcal{L} \mid \Gamma, x : \{\nu : \mathbf{int} \mid \nu = n\} \vdash e : \tau \Rightarrow \Gamma' \qquad x \notin dom(\Gamma')}{\Theta \mid \mathcal{L} \mid \Gamma \vdash \mathbf{let}\, x = n\, \mathbf{in}\, e : \tau \Rightarrow \Gamma'} \quad \text{(T-LetInt)}$$

$$\frac{\begin{array}{c}\Theta \mid \mathcal{L} \mid \Gamma[x \hookleftarrow \{\nu : \mathbf{int} \mid \varphi \wedge \nu = 0\}] \vdash e_1 : \tau \Rightarrow \Gamma' \\ \Theta \mid \mathcal{L} \mid \Gamma[x \hookleftarrow \{\nu : \mathbf{int} \mid \varphi \wedge \nu \neq 0\}] \vdash e_2 : \tau \Rightarrow \Gamma'\end{array}}{\Theta \mid \mathcal{L} \mid \Gamma[x : \{\nu : \mathbf{int} \mid \varphi\}] \vdash \mathbf{ifz}\, x\, \mathbf{then}\, e_1\, \mathbf{else}\, e_2 : \tau \Rightarrow \Gamma'} \quad \text{(T-If)}$$

$$\frac{\Theta \mid \mathcal{L} \mid \Gamma[y \hookleftarrow \tau_1], x : (\tau_2 \wedge_x x =_{\tau_2} y)\, \mathbf{ref}^1 \vdash e : \tau \Rightarrow \Gamma' \quad x \notin dom(\Gamma')}{\Theta \mid \mathcal{L} \mid \Gamma[y : \tau_1 + \tau_2] \vdash \mathbf{let}\, x = \mathbf{mkref}\, y\, \mathbf{in}\, e : \tau \Rightarrow \Gamma'} \quad \text{(T-MkRef)}$$

$$\frac{\Theta \mid \mathcal{L} \mid \Gamma \vdash e_1 : \tau' \Rightarrow \Gamma' \quad \Theta \mid \mathcal{L} \mid \Gamma' \vdash e_2 : \tau'' \Rightarrow \Gamma''}{\Theta \mid \mathcal{L} \mid \Gamma \vdash e_1\, ;\, e_2 : \tau'' \Rightarrow \Gamma''} \quad \text{(T-Seq)}$$

$$\frac{\tau' = \begin{cases}\tau_1 \wedge_y y =_{\tau_1} x & r > 0 \\ \tau_1 & r = 0\end{cases} \quad \Theta \mid \mathcal{L} \mid \Gamma[y \hookleftarrow \tau'\, \mathbf{ref}^r], x : \tau_2 \vdash e : \tau \Rightarrow \Gamma' \quad x \notin dom(\Gamma')}{\Theta \mid \mathcal{L} \mid \Gamma[y : (\tau_1 + \tau_2)\, \mathbf{ref}^r] \vdash \mathbf{let}\, x = *y\, \mathbf{in}\, e : \tau \Rightarrow \Gamma'} \quad \text{(T-Deref)}$$

$$\frac{\Gamma \models \varphi \qquad \epsilon \mid \Gamma \vdash_{WF} \varphi \quad \Theta \mid \mathcal{L} \mid \Gamma \vdash e : \tau \Rightarrow \Gamma'}{\Theta \mid \mathcal{L} \mid \Gamma \vdash \mathbf{assert}(\varphi)\, ;\, e : \tau \Rightarrow \Gamma'} \quad \text{(T-Assert)}$$

Fig. 7. Expression typing rules.

refinement types. The typing rules are given in Figures 7 and 8. A typing judgment takes the form $\Theta \mid \mathcal{L} \mid \Gamma \vdash e : \tau \Rightarrow \Gamma'$, which indicates that e is well-typed under a function type environment Θ, typing context $\mathcal{L}$, and type environment Γ, and evaluates to a value of type τ and modifies the input environment according to Γ'. Any valid typing derivation must have $\mathcal{L} \vdash_{WF} \Gamma$, $\mathcal{L} \vdash_{WF} \Gamma'$, and $\mathcal{L} \mid \Gamma' \vdash_{WF} \tau$, i.e., the input and output type environments and result type must be well-formed.

The typing rules in Figure 7 handle the relatively standard features in our language. The rule T-Seq for sequential composition is fairly straightforward except that the output type environment for e_1 is the input type environment for e_2. T-LetInt is also straightforward; since x is bound to a constant, it is given type $\{\nu : \mathbf{int} \mid \nu = n\}$ to indicate x is exactly n. The output type environment Γ' cannot mention x (expressed with $x \notin dom(\Gamma)$) to prevent x from escaping its scope. This requirement can be met by applying the subtyping rule (see below) to weaken refinements to no longer mention x. As in other refinement type systems [40], this requirement is critical for ensuring soundness.

Rule T-Let is crucial to understanding our ownership type system. The body of the let expression e is typechecked under a type environment where the type of y in Γ is linearly split into two types: τ_1 for y and τ_2 for the newly created binding x. This splitting is expressed using the $+$ operator. If y is a reference type, the split operation distributes some portion of y's ownership information to its new alias x. The split operation also distributes refinement information between the two types. For example, type $\{\nu : \mathbf{int} \mid \nu > 0\}\, \mathbf{ref}^1$ can be split into (1) $\{\nu : \mathbf{int} \mid \nu > 0\}\, \mathbf{ref}^r$ and $\{\nu : \mathbf{int} \mid \nu > 0\}\, \mathbf{ref}^{(1-r)}$ (for $r \in (0, 1)$),

i.e., two *immutable* references with non-trivial refinement information, or (2) $\{\nu : \mathbf{int} \mid \nu > 0\}\,\mathbf{ref}^1$ and $\{\nu : \mathbf{int} \mid \top\}\,\mathbf{ref}^0$, where one of the aliases is mutable and the other provides no refinement information. How a type is split depends on the usage of x and y in e. Formally, we define the type addition operator as the least commutative partial operation that satisfies the following rules:

$$\{\nu : \mathbf{int} \mid \varphi_1\} + \{\nu : \mathbf{int} \mid \varphi_2\} = \{\nu : \mathbf{int} \mid \varphi_1 \wedge \varphi_2\} \qquad \text{(TADD-INT)}$$

$$\tau_1\,\mathbf{ref}^{r_1} + \tau_2\,\mathbf{ref}^{r_2} = (\tau_1 + \tau_2)\,\mathbf{ref}^{r_1 + r_2} \qquad \text{(TADD-REF)}$$

Viewed another way, type addition describes how to combine two types for the same value such that the combination soundly incorporates all information from the two original types. Critically, the type addition operation cannot create or destroy ownership and refinement information, only combine or divide it between types. Although not explicit in the rules, by ownership well-formedness, if the entirety of a reference's ownership is transferred to another type during a split, all refinements in the remaining type must be $\top$.

The additional bits $\wedge_y y =_{\tau_1} x$ and $\wedge_x x =_{\tau_2} y$ express equality between x and y as refinements. We use the strengthening operation $\tau \wedge_x \varphi$ and typed equality proposition $x =_\tau y$, defined respectively as:

$$\{\nu : \mathbf{int} \mid \varphi\} \wedge_y \varphi' = \{\nu : \mathbf{int} \mid \varphi \wedge [\,\nu\,/\,y\,]\,\varphi'\} \qquad (x =_{\{\nu : \mathbf{int} \mid \varphi\}} y) = (x = y)$$

$$\tau\,\mathbf{ref}^r \wedge_y \varphi' = \tau\,\mathbf{ref}^r \qquad (x =_{\tau\,\mathbf{ref}^r} y) = \top$$

We do not track equality between references or between the contents of aliased reference cells as doing so would violate our refinement well-formedness condition. These operations are also used in other rules that can introduce equality.

Rule T-MKREF is very similar to T-LET, except that x is given a reference type of ownership 1 pointing to τ_2, which is obtained by splitting the type of y. In T-DEREF, the content type of y is split and distributed to x. The strengthening is *conditionally* applied depending on the ownership of the dereferenced pointer, that is, if $r = 0$, τ' has to be a maximal type $\top_i$.

Our type system also tracks path information; in the T-IF rule, we update the refinement on the condition variable within the respective branches to indicate whether the variable must be zero. By requiring both branches to produce the same output type environment, we guarantee that these conflicting refinements are rectified within the type derivations of the two branches.

The type rule for assert statements has the precondition $\Gamma \models \varphi$ which is defined to be $\models [\![\Gamma]\!] \implies \varphi$, i.e., the logical formula $[\![\Gamma]\!] \implies \varphi$ is valid in the chosen theory. $[\![\Gamma]\!]$ lifts the refinements on the integer valued variables into a proposition in the logic used for verification. This denotation operation is defined as:

$$[\![\bullet]\!] = \top \qquad [\![\{\nu : \mathbf{int} \mid \varphi\}]\!]_y = [\,y\,/\,\nu\,]\,\varphi$$

$$[\![\Gamma, x : \tau]\!] = [\![\Gamma]\!] \wedge [\![\tau]\!]_x \qquad [\![\tau'\,\mathbf{ref}^r]\!]_y = \top$$

If the formula $[\![\Gamma]\!] \implies \varphi$ is valid, then in any context and under any valuation of program variables that satisfy the refinements in $[\![\Gamma]\!]$, the predicate φ must be true and the assertion must not fail. This intuition forms the foundation of our soundness claim (Section 3.4).

$$\frac{\begin{array}{c}\text{(The shapes of } \tau' \text{ and } \tau_2 \text{ are similar)}\\ \Theta \mid \mathcal{L} \mid \Gamma[x \hookleftarrow \tau_1][y \hookleftarrow (\tau_2 \wedge_y y =_{\tau_2} x)\,\mathbf{ref}^1] \vdash e : \tau \Rightarrow \Gamma'\end{array}}{\Theta \mid \mathcal{L} \mid \Gamma[x : \tau_1 + \tau_2][y : \tau'\,\mathbf{ref}^1] \vdash y := x\,;e : \tau \Rightarrow \Gamma'} \quad \text{(T-Assign)}$$

$$\frac{\begin{array}{c}(\tau_1\,\mathbf{ref}^{r_1} + \tau_2\,\mathbf{ref}^{r_2}) \approx (\tau_1'\,\mathbf{ref}^{r_1'} + \tau_2'\,\mathbf{ref}^{r_2'})\\ \Theta \mid \mathcal{L} \mid \Gamma[x \hookleftarrow \tau_1'\,\mathbf{ref}^{r_1'}][y \hookleftarrow \tau_2'\,\mathbf{ref}^{r_2'}] \vdash e : \tau \Rightarrow \Gamma'\end{array}}{\Theta \mid \mathcal{L} \mid \Gamma[x : \tau_1\,\mathbf{ref}^{r_1}][y : \tau_2\,\mathbf{ref}^{r_2}] \vdash \mathbf{alias}(x = y)\,;e : \tau \Rightarrow \Gamma'} \quad \text{(T-Alias)}$$

$$\frac{\begin{array}{c}(\tau_1\,\mathbf{ref}^{r_1} + \tau_2\,\mathbf{ref}^{r_2}) \approx (\tau_1'\,\mathbf{ref}^{r_1'} + \tau_2'\,\mathbf{ref}^{r_2'})\\ \Theta \mid \mathcal{L} \mid \Gamma[x \hookleftarrow \tau_1'\,\mathbf{ref}^{r_1}][y \hookleftarrow (\tau_2'\,\mathbf{ref}^{r_2'})\,\mathbf{ref}^{r}] \vdash e : \tau \Rightarrow \Gamma'\end{array}}{\Theta \mid \mathcal{L} \mid \Gamma[x : \tau_1\,\mathbf{ref}^{r_1}][y : (\tau_2\,\mathbf{ref}^{r_2})\,\mathbf{ref}^{r}] \vdash \mathbf{alias}(x = *y)\,;e : \tau \Rightarrow \Gamma'} \quad \text{(T-AliasPtr)}$$

$$\frac{\Gamma \leq \Gamma' \qquad \Theta \mid \mathcal{L} \mid \Gamma' \vdash e : \tau \Rightarrow \Gamma'' \qquad \Gamma'', \tau \leq \Gamma''', \tau'}{\Theta \mid \mathcal{L} \mid \Gamma \vdash e : \tau' \Rightarrow \Gamma'''} \quad \text{(T-Sub)}$$

$$\tau_1 \approx \tau_2 \text{ iff } \bullet \vdash \tau_1 \leq \tau_2 \text{ and } \bullet \vdash \tau_2 \leq \tau_1.$$

Fig. 8. Pointer manipulation and subtyping

$$\frac{\Gamma \models \varphi_1 \implies \varphi_2}{\Gamma \vdash \{\nu : \mathbf{int} \mid \varphi_1\} \leq \{\nu : \mathbf{int} \mid \varphi_2\}} \quad \text{(S-Int)} \qquad \frac{\forall\, x \in dom(\Gamma').\Gamma \vdash \Gamma(x) \leq \Gamma'(x)}{\Gamma \leq \Gamma'} \quad \text{(S-TyEnv)}$$

$$\frac{r_1 \geq r_2 \qquad \Gamma \vdash \tau_1 \leq \tau_2}{\Gamma \vdash \tau_1\,\mathbf{ref}^{r_1} \leq \tau_2\,\mathbf{ref}^{r_2}} \quad \text{(S-Ref)} \qquad \frac{\Gamma, x : \tau \leq \Gamma', x : \tau' \qquad x \notin dom(\Gamma)}{\Gamma, \tau \leq \Gamma, \tau'} \quad \text{(S-Res)}$$

Fig. 9. Subtyping rules.

Destructive Updates, Aliasing, and Subtyping We now discuss the handling of assignment, aliasing annotations, and subtyping as described in Figure 8. Although apparently unrelated, all three concern updating the refinements of (potentially) aliased reference cells.

Like the binding forms discussed above, T-Assign splits the assigned value's type into two types via the type addition operator, and distributes these types between the right hand side of the assignment and the mutated reference contents. Refinement information in the fresh contents *may* be inconsistent with any previous refinement information; only the shapes must be the same. In a system with unrestricted aliasing, this typing rule would be unsound as it would admit writes that are inconsistent with refinements on aliases of the left hand side. However, the assignment rule requires that the updated reference has an ownership of 1. By the ownership type invariant, all aliases with the updated reference have 0 ownership, and by ownership well-formedness may only contain the $\top$ refinement.

Example 3. We can type the program as follows:

```
let x = mkref 5 in        // x : {ν : int | ν = 5} ref^1
let y = x in              // x : ⊤_1, y : {ν : int | ν = 5} ref^1
y := 4; assert(*y = 4)    // x : ⊤_1, y : {ν : int | ν = 4} ref^1
```

In this and later examples, we include type annotations within comments. We stress that these annotations are for expository purposes only; our tool can infer these types automatically with no manual annotations.

As described thus far, the type system is quite strict: if ownership has been completely transferred from one reference to another, the refinement information found in the original reference is effectively useless. Additionally, once a mutable pointer has been split through an assignment or let expression, there is no way to recover mutability. The typing rule for must alias assertions, T-Alias and T-AliasPtr, overcomes this restriction by exploiting the must-aliasing information to "shuffle" or redistribute ownerships *and refinements* between two aliased pointers. The typing rule assigns two fresh types $\tau_1' \, \mathbf{ref}^{r_1'}$ and $\tau_2' \, \mathbf{ref}^{r_2'}$ to the two operand pointers. The choice of τ_1', r_1', τ_2', and r_2' is left open provided that the sum of the new types, $(\tau_1' \, \mathbf{ref}^{r_1'}) + (\tau_2' \, \mathbf{ref}^{r_2'})$ is equivalent (denoted $\approx$) to the sum of the original types. Formally, $\approx$ is defined as in Figure 8; it implies that any refinements in the two types must be logically equivalent and that ownerships must also be equal. This redistribution is sound precisely because the two references are assumed to alias; the total ownership for the single memory cell pointed to by both references cannot be increased by this shuffling. Further, any refinements that hold for the contents of one reference must necessarily hold for contents of the other and vice versa.

Example 4 (Shuffling ownerships and refinements). Let $\varphi_{=n}$ be $\nu = n$.

```
let x = mkref 5 in     // x : {ν : int | φ=5} ref^1
let y = x in           // x : ⊤1, y : {ν : int | φ=5} ref^1
y := 4; alias(x = y)   // x : {ν : int | φ=4} ref^0.5, y : {ν : int | φ=4} ref^0.5
```

The final type assignment for x and y is justified by

$$\top_1 + \{\nu : \mathbf{int} \mid \varphi_{=4}\} \, \mathbf{ref}^1 = \{\nu : \mathbf{int} \mid \top \wedge \varphi_{=4}\} \, \mathbf{ref}^1 \approx$$
$$\{\nu : \mathbf{int} \mid \varphi_{=4} \wedge \varphi_{=4}\} \, \mathbf{ref}^1 = \{\nu : \mathbf{int} \mid \varphi_{=4}\} \, \mathbf{ref}^{0.5} + \{\nu : \mathbf{int} \mid \varphi_{=4}\} \, \mathbf{ref}^{0.5} .$$

The aliasing rules give fine-grained control over ownership information. This flexibility allows mutation through two or more aliased references within the same scope. Provided sufficient aliasing annotations, the type system may shuffle ownerships between one or more live references, enabling and disabling mutability as required. Although the reliance on these annotations appears to decrease the practicality of our type system, we expect these aliasing annotations can be inserted by a conservative must-aliasing analysis. Further, empirical experience from our prior work [49] indicates that only a small number of annotations are required for larger programs.

Example 5 (Shuffling Mutability). Let $\varphi_{=n}$ again be $\nu = n$. The following program uses two live, aliased references to mutate the same memory location:

```
let x = mkref 0 in
let y = x in                 // x : {ν : int | φ=0} ref^1, y : ⊤1
  x := 1; alias(x = y);      // x : ⊤1, y : {ν : int | φ=1} ref^1
  y := 2; alias(x = y);      // x : {ν : int | φ=2} ref^0.5, y : {ν : int | φ=2} ref^0.5
  assert(*x = 2)
```

$$\frac{\Theta(f) = \forall\lambda.\,\langle x_1:\tau_1,\ldots,x_n:\tau_n\rangle \to \langle x_1:\tau'_1,\ldots,x_n:\tau'_n \mid \tau\rangle \quad \sigma_\alpha = [\ell:\mathcal{L}/\lambda] \quad \sigma_x = [y_1/x_1]\cdots[y_n/x_n] \quad \Theta \mid \mathcal{L} \mid \Gamma[y_i \hookleftarrow \sigma_\alpha\,\sigma_x\,\tau'_i], x:\sigma_\alpha\,\sigma_x\,\tau \vdash e:\tau' \Rightarrow \Gamma' \quad x \notin dom(\Gamma')}{\Theta \mid \mathcal{L} \mid \Gamma[y_i:\sigma_\alpha\,\sigma_x\,\tau_i] \vdash \mathbf{let}\ x = f^\ell(y_1,\ldots,y_n)\ \mathbf{in}\ e:\tau' \Rightarrow \Gamma'} \quad \text{(T-Call)}$$

$$\frac{\Theta(f) = \forall\lambda.\,\langle x_1:\tau_1,\ldots,x_n:\tau_n\rangle \to \langle x_1:\tau'_1,\ldots,x_n:\tau'_n \mid \tau\rangle \quad \Theta \mid \lambda \mid x_1:\tau_1,\ldots,x_n:\tau_n \vdash e:\tau \Rightarrow x_1:\tau'_1,\ldots,x_n:\tau'_n}{\Theta \vdash f \mapsto (x_1,\,..,x_n)e} \quad \text{(T-FunDef)}$$

$$\frac{\forall f \mapsto (x_1,\,..,x_n)e \in D.\Theta \vdash f \mapsto (x_1,\,..,x_n)e \quad dom(D) = dom(\Theta)}{\Theta \vdash D} \quad \text{(T-Funs)}$$

$$\frac{\Theta \vdash D \quad \vdash_{WF} \Theta \quad \Theta \mid \epsilon \mid \bullet \vdash e:\tau \Rightarrow \Gamma}{\vdash \langle D, e\rangle} \quad \text{(T-Prog)}$$

Fig. 10. Program typing rules

After the first aliasing statement the type system shuffles the (exclusive) mutability between x and y to enable the write to y. After the second aliasing statement the ownership in y is split with x; note that transferring all ownership from y to x would also yield a valid typing.

Finally, we describe the subtyping rule. The rules for subtyping types and environments are shown in Figure 9. For integer types, the rules require the refinement of a supertype is a logical consequence of the subtype's refinement conjoined with the lifting of Γ. The subtype rule for references is *covariant* in the type of reference contents. It is widely known that in a language with unrestricted aliasing and mutable references such a rule is unsound: after a write into the coerced pointer, reads from an alias may yield a value disallowed by the alias' type [37]. However, as in the assign case, ownership types prevent unsoundness; a write to the coerced pointer requires the pointer to have ownership 1, which guarantees any aliased pointers have the maximal type and provide no information about their contents beyond simple types.

3.3 Interprocedural Fragment and Context-Sensitivity

We now turn to a discussion of the interprocedural fragment of our language, and how our type system propagates context information. The remaining typing rules for our language are shown in Figure 10. These rules concern the typing of function calls, function bodies, and entire programs.

We first explain the T-Call rule. The rule uses two substitution maps. σ_x translates between the parameter names used in the function type and actual argument names at the call-site. σ_α instantiates all occurrences of λ in the callee type with $\ell : \mathcal{L}$, where ℓ is the label of the call-site and $\mathcal{L}$ the typing context of the call. The types of the arguments y_i's are required to match the parameter

types (post substitution). The body of the let binding is then checked with the argument types updated to reflect the changes in the function call (again, post substitution). This update is well-defined because we require all function arguments be distinct as described in Section 2.1. Intuitively, the substitution σ_α represents incrementally refining the behavior of the callee function with partial context information. If $\mathcal{L}$ is itself a context variable λ', this substitution effectively transforms any context prefix queries over λ in the argument/return/output types into a queries over $\ell : \lambda'$. In other words, while the exact concrete execution context of the callee is unknown, the context must at least begin with ℓ which can potentially rule out certain behaviors.

Rule T-FunDef type checks a function definition $f \mapsto (x_1, .., x_n)e$ against the function type given in Θ. As a convenience we assume that the parameter names in the function type match the formal parameters in the function definition. The rule checks that under an initial environment given by the argument types the function body produces a value of the return type and transforms the arguments according to the output types. As mentioned above, functions may be executed under many different contexts, so type checking the function body is performed under the context variable λ that occurs in the function type.

Finally, the rule for typing programs (T-Prog) checks that all function definitions are well typed under a well-formed function type environment, and that the entry point e is well typed in an empty type environment and the typing context ϵ, i.e., the initial context.

Example 6 (1-CFA). Recall the program in Figure 3 in Section 1; assume the function calls are labeled as follows:

```
p := get^ℓ1(p) + 1;
// ...
q := get^ℓ2(q) + 1;
```

Taking τ_p to be the type shown in Example 2:

$$\{\nu : \mathbf{int} \mid (\ell_1 \subseteq \lambda \implies \nu = 3) \wedge (\ell_2 \subseteq \lambda \implies \nu = 5)\}$$

we can give `get` the type $\forall\lambda.\, \langle z : \tau_p\, \mathbf{ref}^1 \rangle \to \langle z : \tau_p\, \mathbf{ref}^1 \mid \tau_p \rangle$.

Example 7 (2-CFA). To see how context information propagates across multiple calls, consider the following change to the code considered in Example 6:

```
get_real(z) { *z }
get(z) { get_real^ℓ3(z) }
```

The type of `get` remains as in Example 6, and taking τ to be

$$\{\nu : \mathbf{int} \mid (\ell_3\, \ell_1 \subseteq \lambda' \implies \nu = 3) \wedge (\ell_3\, \ell_2 \subseteq \lambda' \implies \nu = 5)\}$$

the type of `get_real` is: $\forall\lambda'.\, \langle z : \tau\, \mathbf{ref}^1 \rangle \to \langle z : \tau\, \mathbf{ref}^1 \mid \tau \rangle$.

We focus on the typing of the call to `get_real` in `get`; it is typed in context λ and a type environment where `p` is given type τ_p from Example 6.

Applying the substitution $[\ell_3 : \lambda/\lambda']$ to the argument type of `get_real` yields:

$$\{\nu : \mathbf{int} \mid (\ell_3\,\ell_1 \subseteq \ell_3 : \lambda \implies \nu = 3) \wedge (\ell_3\,\ell_2 \subseteq \ell_3 : \lambda \implies \nu = 5)\}\,\mathbf{ref}^1 \approx \{\nu : \mathbf{int} \mid (\ell_1 \subseteq \lambda \implies \nu = 3) \wedge (\ell_2 \subseteq \lambda \implies \nu = 5)\}\,\mathbf{ref}^1$$

which is exactly the type of `p`. A similar derivation applies to the return type of `get_real` and thus `get`.

3.4 Soundness

We have proven that any program that type checks according to the rules above will never experience an assertion failure. We formalize this claim with the following soundness theorem.

Theorem 1 (Soundness). *If* $\vdash \langle D, e\rangle$*, then* $\langle \emptyset, \emptyset, \cdot, e\rangle \not\longrightarrow_D^* \mathbf{AssertFail}$.

Further, any well-typed program either diverges, halts in the configuration **AliasFail***, or halts in a configuration* $\langle H, R, \cdot, x\rangle$ *for some* H, R *and* x*, i.e., evaluation does not get stuck.*

Proof (Sketch). By standard progress and preservation lemmas; the full proof has been omitted for space reasons and can be found in the accompanying appendix.

4 Inference and Extensions

We now briefly describe the inference algorithm implemented in our tool ConSORT. We sketch some implemented extensions needed to type more interesting programs and close with a discussion of current limitations of our prototype.

4.1 Inference

Our tool first runs a standard, simple type inference algorithm to generate type templates for every function parameter type, return type, and for every live variable at each program point. For a variable x of simple type $\tau_S ::= \mathbf{int} \mid \tau_S\,\mathbf{ref}$ at program point p ConSORT generates a type template $[\![\tau_S]\!]_{x,0,p}$ as follows:

$$[\![\mathbf{int}]\!]_{x,n,p} = \{\nu : \mathbf{int} \mid \varphi_{x,n,p}(\nu; \mathbf{FV}_p)\} \quad [\![\tau_S\,\mathbf{ref}]\!]_{x,n,p} = [\![\tau_S]\!]_{x,n+1,p}\,\mathbf{ref}^{r_{x,n,p}}$$

$\varphi_{x,n,p}(\nu; \mathbf{FV}_p)$ denotes a fresh relation symbol applied to ν and the free variables of simple type **int** at program point p (denoted $\mathbf{FV}_p$). $r_{x,n,p}$ is a fresh ownership variable. For each function f, there are two synthetic program points, f^b and f^e for the beginning and end of the function respectively. At both points, ConSORT generates type template for each argument, where $\mathbf{FV}_{f^b}$ and $\mathbf{FV}_{f^e}$ are the names of integer typed parameters. At f^e, ConSORT also generates a type template for the return value using a synthetic variable name. We write Γ^p to indicate the type environment at point p, where every variable is mapped to the appropriate type template. Then $[\![\Gamma^p]\!]$ is equivalent to $\bigwedge_{x \in \mathbf{FV}_p} \varphi_{x,0,p}(x; \mathbf{FV}_p)$.

When generating these type templates, our implementation also generates ownership well-formedness constraints. Specifically, for a type template of the form $\{\nu : \mathbf{int} \mid \varphi_{x,n+1,p}(\nu; \mathbf{FV}_p)\}\,\mathbf{ref}^{r_{x,n,p}}$ ConSORT emits the constraint: $r_{x,n,p} = 0 \implies \varphi_{x,n+1,p}(\nu; \mathbf{FV}_p)$ and for a type template $(\tau\,\mathbf{ref}^{r_{x,n+1,p}})\,\mathbf{ref}^{r_{x,n,p}}$ ConSORT emits the constraint $r_{x,n,p} = 0 \implies r_{x,n+1,p} = 0$.

ConSORT then walks the program, generating constraints between relation symbols and ownership variables according to the typing rules. These constraints take three forms, ownership constraints, subtyping constraints, and assertion constraints. Ownership constraints are simple linear (in)equalities over ownership variables and constants, according to conditions imposed by the typing rules. For example, if variable x has the type template $\tau\,\mathbf{ref}^{r_{x,0,p}}$ for the expression $x := y\,;\, e$ at point p, ConSORT generates the constraint $r_{x,0,p} = 1$.

ConSORT emits subtyping constraints between the relation symbols at related program points according to the rules of the type system. For example, consider the let binding $\mathbf{let}\, x = y\, \mathbf{in}\, e$ at program point p, where e is at program point p', and x has simple type $\mathbf{int}\,\mathbf{ref}$. ConSORT generates the following subtyping constraint:

$$[\![\Gamma^p]\!] \wedge \varphi_{y,1,p}(\nu; \mathbf{FV}_p) \implies \varphi_{y,1,p'}(\nu; \mathbf{FV}_{p'}) \wedge \varphi_{x,1,p'}(\nu; \mathbf{FV}_{p'})$$

in addition to the ownership constraint $r_{y,0,p} = r_{y,0,p'} + r_{x,0,p'}$.

Finally, for each $\mathbf{assert}(\varphi)$ in the program, ConSORT emits an assertion constraint of the form: $[\![\Gamma^p]\!] \implies \varphi$ which requires the refinements on integer typed variables in scope are sufficient to prove φ.

Encoding Context Sensitivity To make inference tractable, we require the user to fix *a priori* the maximum length of prefix queries to a constant k (this choice is easily controlled with a command line parameter to our tool). We supplement the arguments in *every* predicate application with a set of integer context variables $c_1, \ldots, c_k$; these variables do not overlap with any program variables.

ConSORT uses these variables to infer context sensitive refinements as follows. Consider a function call $\mathbf{let}\, x = f^{\ell}(y_1, \ldots, y_n)\, \mathbf{in}\, e$ at point p where e is at point p'. ConSORT generates the following constraint for a refinement $\varphi_{y_i,n,p}(\nu, c_1, \ldots, c_k; \mathbf{FV}_p)$ which occurs in the type template of y_i:

$$\begin{aligned}
&\varphi_{y_i,n,p}(\nu, c_0, \ldots, c_k; \mathbf{FV}_p) \implies \sigma_x\, \varphi_{x_i,n,f^b}(\nu, \ell, c_0, \ldots, c_{k-1}; \mathbf{FV}_{f^b}) \\
&\sigma_x\, \varphi_{x_i,n,f^e}(\nu, \ell, c_0, \ldots, c_{k-1}; \mathbf{FV}_{f^e}) \implies \varphi_{y_i,n,p'}(\nu, c_0, \ldots, c_k; \mathbf{FV}_{p'}) \\
&\sigma_x = [y_1/x_1] \cdots [y_n/x_n]
\end{aligned}$$

Effectively, we have encoded $\ell_1 \ldots \ell_k \subseteq \lambda$ as $\wedge_{0<i\leq k} c_i = \ell_i$. In the above, the shift from $c_0, \ldots, c_k$ to $\ell, c_0, \ldots, c_{k-1}$ plays the role of σ_α in the T-Call rule. The above constraint effectively determines the value of c_0 within the body of the function f. If f calls another function g, the above rule propagates this value of c_0 to c_1 within g and so on. The solver may then instantiate relation symbols with predicates that are conditional over the values of c_i.

Solving Constraints The results of the above process are two systems of constraints; real arithmetic constraints over ownership variables and constrained Horn clauses (CHC) over the refinement relations. The inference process generates subtyping constraints for every variable at each program point so the number of constraints, relations, and ownership variables is quadratic in the size of the program in the worst case.[4] These systems are not independent; the relation constraints may mention the value of ownership variables due to the well-formedness constraints described above. The ownership constraints are first solved with Z3 [15]. These constraints are non-linear but Z3 appears particularly well-engineered to quickly find solutions for the instances generated by ConSORT. We also constrain Z3 to maximize the number of non-zero ownership variables; this ensures as few refinements as possible are constrained to be $\top$ due to ownership well-formedness.

The values of ownership variables inferred by Z3 are then substituted into the constrained Horn clauses, and the resulting system is checked for satisfiability with an off-the-shelf CHC solver. Our implementation generates constraints in the industry standard SMT-Lib2 format [7]; any solver that accepts this format can be used as a backend for ConSORT. Our implementation currently supports Spacer [33] (part of the Z3 solver [15]), HoICE [12], and Eldarica [41] (adding a new backend requires only a handful of lines of glue code). We found that different solvers are better tuned to different problems; we also implemented *parallel mode* which runs all supported solvers in parallel and uses the result of the first solver to complete.

4.2 Extensions

We now sketch some language extensions supported by our implementation.

Primitive Operations As defined in Section 2, our language can compare integers to zero and load and store them from memory, but can perform no meaningful computation over these numbers. To promote the flexibility of our type system and simplify our soundness statement, we do not fix a set of primitive operations and their static semantics. Instead, we assume any set of primitive operations used in a program are given sound function types in Θ. For example, under the assumption that $+$ has its usual semantics and the underlying logic supports $+$, we can give $+$ the type $\forall\lambda.\langle x:\top, y:\top\rangle \to \langle x:\top, y:\top \mid \{\nu : \mathbf{int} \mid \nu = x + y\}\rangle$. Interactions with a nondeterministic environment or unknown program inputs can then be modeled with a primitive that returns integers refined with $\top$.

Dependent Tuples Our implementation supports types of the form: $(x_1 : \tau_1, \ldots, x_n : \tau_n)$, where x_i can appear within τ_j $(j \neq i)$ if τ_i is an integer type. For example, $(x : \{\nu : \mathbf{int} \mid \top\}, y : \{\nu : \mathbf{int} \mid \nu > x\})$ is the type of tuples whose second element is strictly greater than the first. We also extend the language with tuple constructors as a new value form, and let bindings with tuple patterns as the LHS.

[4] In practice we can simplify away many "intermediate" relations to improve performance.

The extension to type checking is relatively straightforward; the only significant extensions are to the subtyping rules. Specifically, the subtyping check for a tuple element $x_i : \tau_i$ is performed in a type environment elaborated with the types and names of other tuple elements. The extension to type inference is also straightforward; the arguments for a predicate symbol include any enclosing dependent tuple names and the environment in subtyping constraints is likewise extended.

Recursive Types Our language also supports some unbounded heap structures via recursive reference types. To keep inference tractable, we forbid nested recursive types, multiple occurrences of a type variable bound by the μ binder and additionally fix the shape of refinements that occur within a recursive type. For recursive refinements that fit the above restriction, our approach for refinements is broadly similar to that in [32], and we use the ownership scheme of [49] for handling ownership. We first use simple type inference to infer the shape of the recursive types, and automatically insert fold/unfold annotations into the source program. As in [32], the refinements within an unfolding of a recursive type may refer to dependent tuple names bound by the enclosing type. These recursive types can express the invariant of a mutable, sorted list. Adapting the approach in [49], all recursive types are unfolded once before assigning ownership variables. Further unfoldings copy existing ownership variables.

As in Java or C++, our language does not support sum types, and thus any instantiation of a recursive type must use a null pointer. Our implementation therefore supports an **ifnull** construct and a distinguished null constant. Our implementation allows any refinement to hold for the null constant, including $\bot$. Currently, our implementation currently does *not* detect null pointer dereferences, and all soundness guarantees are made modulo freedom of null dereferences. As $[\![\Gamma]\!]$ omits the refinements on reference types, null pointer refinements do not affect the verification of programs free of null pointer dereferences.

Arrays Our implementation supports arrays of integers. Each array is given an ownership describing the ownership of memory allocated for the entire array. The array type contains two refinements: the first refines the length of the array itself, and the second refines the entire array contents. The content refinement may refer to a symbolic index variable for precise, per-index refinements. At reads and writes to the array, ConSORT instantiates the refinement's symbolic index variable with the concrete variable used at the read/write.

As in [49], our restriction to arrays of integers stems from the difficulty of ownership inference. Sound handling of pointer arrays requires index-wise tracking of ownerships which significantly complicates automated inference. We leave supporting arrays of pointers to future work.

4.3 Limitations

Our current approach is not complete; there are safe programs that will be rejected by our type system. As mentioned in Section 3.1, our well-formedness condition forbids refinements that refer to memory locations. As a result, ConSORT

Table 1. Description of benchmark suite adapted from JayHorn. **Java** are programs that test Java-specific features. **Inc** are tests that cannot be handled by ConSORT, e.g., null checking, etc. **Bug** includes a "safe" program we discovered was actually incorrect.

Set	Orig.	Adapted	Java	Inc	Bug
Safe	41	32	6	2	1
Unsafe	41	26	13	2	0

cannot in general express, e.g., that the contents of two references are equal. Further, due to our reliance on automated theorem provers we are restricted to logics with sound but potentially incomplete decision procedures. ConSORT also does not support conditional or context-sensitive ownerships, and therefore cannot precisely handle conditional mutation or aliasing.

5 Experiments

We now present the results of some preliminary experiments performed with the implementation described in Section 4.[5] The goal of these experiments was to answer the following questions:

1. Is the type system (and extensions of Section 4) expressive enough to type and verify non-trivial programs?
2. Is type inference feasible?

To answer these questions, we evaluated our prototype implementation on two sets of benchmarks. The first set is adapted from JayHorn [29, 30], a verification tool for Java. This test suite contains a combination of 82 safe and unsafe programs written in Java. We chose this benchmark suite as, like ConSORT, JayHorn is concerned with the automated verification of programs in a language with mutable, aliased memory cells. Further, although some of their benchmark programs tested Java specific features, most could be adapted into our low-level language. The tests we could adapt provide a comparison with existing state-of-the-art verification techniques. A detailed breakdown of the adapted benchmark suite can be found in Table 1.

Remark 2. The original JayHorn paper includes two additional benchmark sets, Mine Pump and CBMC. Both our tool and recent JayHorn versions time out on the Mine Pump benchmark. Further, the CBMC tests were either subsumed by our own test programs, tested Java specific features, or tested program synthesis functionality. We therefore omitted both of these benchmarks from our evaluation.

The second benchmark set consists of data structure implementations and microbenchmarks written directly in our low-level imperative language. We

[5] Our experiments and the ConSORT source code are available at `https://www.fos.kuis.kyoto-u.ac.jp/projects/consort/`.

developed this suite to test the expressive power of our type system and inference. The programs included in this suite are:

- **Array-List** Implementation of an unbounded list backed by an array.
- **Sorted-List** Implementation of a mutable, sorted list maintained with an in-place insertion sort algorithm.
- **Shuffle** Multiple live references are used to mutate the same location in program memory as in Example 5.
- **Mut-List** Implementation of general linked lists with a clear operation.
- **Array-Inv** A program which allocates a length n array and writes the value i at every index i.
- **Intro2** The motivating program shown in Figure 2 in Section 1.

We introduced unsafe mutations to these programs to check our tool for unsoundness and translated these programs into Java for further comparison with JayHorn.

Our benchmarks and JayHorn's require a small number of trivially identified alias annotations. The adapted JayHorn benchmarks contain a total of 6 annotations; the most for any individual test was 3. The number of annotations required for our benchmark suite are shown in column **Ann.** of Table 2.

Experimental Setup We first ran ConSORT on each program in our benchmark suite and ran the most recent[6] development build of JayHorn[7] on the corresponding Java version. We recorded the final verification result for both our tool and JayHorn. We also collected the end-to-end runtime of ConSORT for each test; we do not give a performance comparison with JayHorn given the many differences in target languages. For the JayHorn suite, we first ran our tool on the adapted version of each test program and ran JayHorn on the original Java version. We also did not collect runtime information for this set of experiments because our goal is a comparison of tool precision, not performance. All tests were run on a machine with 16 GB RAM and 4 CPUs at 2GHz and with a timeout of 60 seconds (the same timeout was used in [29]). We used ConSORT's parallel backend (Section 4) with Z3 version 4.8.4, HoICE version 1.8.1, and Eldarica version 2.0.1 and JayHorn's Eldarica backend.[8]

5.1 Results

The results of our experiments are shown in Table 2. On the JayHorn benchmark suite ConSORT performs competitively with JayHorn, correctly identifying 29 of the 32 safe programs as such. For all 3 tests on which ConSORT timed out after 60 seconds, JayHorn also timed out or encountered an exception (columns *T/O* and *Err.*). For the unsafe programs, ConSORT correctly identified all

[6] As of October 15, 2019

[7] We did not use the most recent official release of JayHorn as it does not include critical fixes for soundness bugs `https://github.com/jayhorn/jayhorn/issues/154`.

[8] We also tried the Spacer backend included with JayHorn, but it required an extremely old version of Z3 (4.3.2) and crashed the JVM on several of our tests.

Table 2. Comparison of ConSORT to JayHorn on the benchmark set of [29] (top) and our custom benchmark suite (bottom). *T/O* indicates a time out.

		ConSORT		**JayHorn**			
Set	**N. Tests**	*Correct*	*T/O*	*Correct*	*T/O*	*Err.*	*Imp.*
Safe	32	29	3	27	1	3	0
Unsafe	26	26	0	20	1	0	6

Name	**Safe?**	**Time(s)**	**Ann**	**JH**	**Name**	**Safe?**	**Time(s)**	**Ann**	**JH**
Array-Inv	✓	8.89	0	T/O	**Array-Inv-BUG**	X	6.62	0	T/O
Array-List	✓	16.84	0	T/O	**Array-List-BUG**	X	1.02	0	T/O
Intro2	✓	0.08	0	T/O	**Intro2-BUG**	X	0.02	0	T/O
Mut-List	✓	1.08	3	T/O	**Mut-List-BUG**	X	0.40	3	T/O
Shuffle	✓	0.07	3	✓	**Shuffle-BUG**	X	0.08	3	X
Sorted-List	✓	1.88	3	T/O	**Sorted-List-BUG**	X	1.36	3	T/O

programs as unsafe within 60 seconds; JayHorn timed out on one test, and answered Unknown for 6 others (column *Imp.*).

On our own benchmark set, ConSORT correctly verifies all safe versions of the programs within 60 seconds. For the unsafe variants, ConSORT was able to quickly and definitively determine these programs unsafe. JayHorn times out on all tests except for **Shuffle** and **ShuffleBUG** (column **JH**). We investigated the cause of time outs and discovered that after verification failed with an unbounded heap model, JayHorn attempts verification on increasingly larger bounded heaps. In every case, JayHorn exceeded the 60 second timeout before reaching a pre-configured limit on the heap bound. This result suggests JayHorn struggles in the presence of per-object invariants and unbounded allocations; the only two tests JayHorn successfully analyzed contain just a single object allocation.

We do not believe this struggle is indicative of a shortcoming in JayHorn's implementation, but stems from the fundamental limitations of JayHorn's memory representation. Like many verification tools (see Section 6), JayHorn uses a single, unchanging invariant to for every object allocated at the same syntactic location; effectively, all objects allocated at the same location are assumed to alias with one another. This representation cannot, in general, handle programs with different invariants for distinct objects that evolve over time. We hypothesize other tools that adopt a similar approach will exhibit the same difficulty.

6 Related Work

The difficulty in handling programs with mutable references and aliasing has been well-studied. Like JayHorn, many approaches model the heap explicitly at verification time, approximating concrete heap locations with allocation site labels [13, 19, 29, 30, 39]; each *abstract location* is also associated with a refinement. As abstract locations summarize many concrete locations, this approach does not in

general admit strong updates and flow-sensitivity; in particular, the refinement associated with an abstract location is fixed for the lifetime of the program. The techniques cited above include various workarounds for this limitation. For example, [13, 39] temporarily allows breaking these invariants through a distinguished program name as long as the abstract location is not accessed through another name. The programmer must therefore eventually bring the invariant back in sync with the summary location. As a result, these systems ultimately cannot precisely handle programs that require evolving invariants on mutable memory.

A similar approach was taken in CQual [22] by Aiken et al. [2]. They used an explicit *restrict* binding for pointers. Strong updates are permitted through pointers bound with *restrict*, but the program is forbidden from using any pointers which share an allocation site while the restrict binding is live.

A related technique used in the field of object-oriented verification is to declare object invariants at the class level and allow these invariants on object fields to be broken during a limited period of time [6, 21]. In particular, the work on Spec# [6] uses an ownership system which tracks whether object a owns object b; like ConSORT's ownership system, these ownerships are used contain the effects of mutation. However, Spec#'s ownership is quite strict and does not admit references to b outside of the owning object a.

F*, a dependently typed dialect of ML, includes an update/select theory of heaps and requires explicit annotations summarizing the heap effects of a method [38, 50, 51]. This approach enables modular reasoning and precise specification of pre- and post-conditions with respect to the heap, but precludes full automation.

The work on rely–guarantee reference types by Gordon et al. [25, 26] uses refinement types in a language mutable references and aliasing. Their approach extends reference types with rely/guarantee predicates; the rely predicate describes possible mutations via aliases, and the guarantee predicate describes the admissible mutations through the current reference. If two references may alias, then the guarantee predicate of one reference implies the rely predicate of the other and vice versa. This invariant is maintained with a splitting operation that is similar to our + operator. Further, their type system allows strong updates to reference refinements provided the new refinements are preserved by the rely predicate. Thus, rely-guarantee refinement support multiple mutable, aliased references with non-trivial refinement information. Unfortunately this expressiveness comes at the cost of automated inference and verification; an embedding of this system into Liquid Haskell [55] described in [26] was forced to sacrifice strong updates.

Work by Degen et al. [16] introduced linear *state annotations* to Java. To effect strong updates in the presence of aliasing, like ConSORT, their system requires annotated memory locations are mutated only through a distinguished reference. Further, all aliases of this mutable reference give no information about the state of the object much like our 0 ownership pointers. However, their system cannot handle multiple, immutable aliases with non-trivial annotation information; *only* the mutable reference may have non-trivial annotation information.

The fractional ownerships in ConSORT and their counterparts in [48, 49] have a clear relation to linear type systems. Many authors have explored the

use of linear type systems to reason in contexts with aliased mutable references [17, 18, 45], and in particular with the goal of supporting strong updates [1]. A closely related approach is RustHorn by Matsushita et al. [35]. Much like ConSORT, RustHorn uses CHC and linear aliasing information for the sound and—unlike ConSORT—complete verification of programs with aliasing and mutability. However, their approach depends on Rust's strict *borrowing discipline*. In particular, it cannot handle programs where multiple aliased references are used to mutate memory in the same lexical region. In contrast, ConSORT supports fine-grained, per-statement changes in mutability and even further control with **alias** annotations, which allows it to verify larger classes of programs.

The idea of using a rational number to express permissions to access a reference dates back to the type system of *fractional permissions* by Boyland [11]. His work used fractional permissions to verify race freedom of a concurrent program without a may-alias analysis. Later, Terauchi [52] proposed a type-inference algorithm that reduces typing constraints to a set of linear inequalities over rational numbers. Boyland's idea also inspired a variant of separation logic for a concurrent programming language [10] to express sharing of read permissions among several threads. Our previous work [48, 49], inspired by Boyland's work and Terauchi's work, proposed methods for type-based verification of resource-leak freedom, in which a rational number expresses an *obligation* to deallocate certain resource, not just a permission.

The issue of context-sensitivity (sometimes called *polyvariance*) is well-studied in the field of abstract interpretation (e.g., [27, 31, 36, 43, 44], see [24] for a recent survey). Polyvariance has also been used in type systems to assign different behaviors to the same function depending on its call site [3, 5, 56]. In the area of refinement type systems, Zhu and Jagannathan developed a context-sensitive dependent type system for a functional language [59] that indexed function types by unique labels attached to call-sites. Our context-sensitivity approach was inspired by this work. In fact, we could have formalized context-polymorphism within the framework of full dependent types, but chose the current presentation for simplicity.

7 Conclusion

We presented ConSORT, a novel type system for safety verification of imperative programs with mutability and aliasing. Our type system is built upon the novel combination of fractional ownership types and refinement types. Ownership types flow-sensitively and precisely track the existence of mutable aliases. ConSORT admits sound strong updates by discarding refinement information on mutably-aliased references as indicated by ownership types. Our type system is amenable to automatic type inference; we have implemented a prototype of this inference tool and found it can verify several non-trivial programs and outperforms a state-of-the-art program verifier. We plan to investigate using fractional ownership types to soundly allow refinements that mention memory locations. In addition, we plan to use the type system and language presented here as a backend language for the verification of high-level languages like Java, C++, etc.

Bibliography

[1] Amal Ahmed, Matthew Fluet, and Greg Morrisett. L^3: a linear language with locations. *Fundamenta Informaticae*, 77(4):397–449, 2007.

[2] Alex Aiken, Jeffrey S. Foster, John Kodumal, and Tachio Terauchi. Checking and inferring local non-aliasing. In *Proceedings of the Conference on Programming Language Design and Implementation*, pages 129–140, 2003.

[3] Torben Amtoft and Franklyn Turbak. Faithful translations between polyvariant flows and polymorphic types. In *European Symposium on Programming*, pages 26–40. Springer, 2000.

[4] Thomas Ball, Vladimir Levin, and Sriram K Rajamani. A decade of software model checking with SLAM. *Communications of the ACM*, 54(7):68–76, 2011.

[5] Anindya Banerjee. A modular, polyvariant and type-based closure analysis. In *Proceedings of the International Conference on Functional Programming*, pages 1–10, 1997.

[6] Mike Barnett, Manuel Fähndrich, K Rustan M Leino, Peter Müller, Wolfram Schulte, and Herman Venter. Specification and verification: the Spec# experience. *Commun. ACM*, 54(6):81–91, 2011.

[7] Clark Barrett, Pascal Fontaine, and Cesare Tinelli. The Satisfiability Modulo Theories Library (SMT-LIB). `www.SMT-LIB.org`, 2016.

[8] Jesper Bengtson, Karthikeyan Bhargavan, Cédric Fournet, Andrew D. Gordon, and Sergio Maffeis. Refinement types for secure implementations. *ACM Transactions on Programming Languages and Systems (TOPLAS)*, 33(2):8, 2011.

[9] Karthikeyan Bhargavan, Barry Bond, Antoine Delignat-Lavaud, Cédric Fournet, Chris Hawblitzel, Cătălin Hriţcu, Samin Ishtiaq, Markulf Kohlweiss, Rustan Leino, Jay Lorch, Kenji Maillard, Jianyang Pan, Bryan Parno, Jonathan Protzenko, Tahina Ramananandro, Ashay Rane, Aseem Rastogi, Nikhil Swamy, Laure Thompson, Peng Wang, Santiago Zanella-Béguelin, and Jean-Karim Zinzindohoué. Everest: Towards a verified, drop-in replacement of HTTPS. In *2nd Summit on Advances in Programming Languages (SNAPL 2017)*. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik, 2017.

[10] Richard Bornat, Cristiano Calcagno, Peter W. O'Hearn, and Matthew J. Parkinson. Permission accounting in separation logic. In *POPL*, pages 259–270, 2005.

[11] John Boyland. Checking interference with fractional permissions. In *SAS 2003*, pages 55–72. Springer, 2003.

[12] Adrien Champion, Naoki Kobayashi, and Ryosuke Sato. HoIce: An ICE-based non-linear Horn clause solver. In *Asian Symposium on Programming Languages and Systems*, pages 146–156. Springer, 2018.

[13] Ravi Chugh, David Herman, and Ranjit Jhala. Dependent types for JavaScript. In *Proceedings of the ACM International Conference on Object Oriented Programming Systems Languages and Applications*, pages 587–606, 2012.

[14] Patrick Cousot, Radhia Cousot, Jerôme Feret, Laurent Mauborgne, Antoine Miné, David Monniaux, and Xavier Rival. The ASTRÉE analyzer. In *European Symposium on Programming*, pages 21–30. Springer, 2005.

[15] Leonardo De Moura and Nikolaj Bjørner. Z3: An efficient SMT solver. In *International conference on Tools and Algorithms for the Construction and Analysis of Systems*, pages 337–340. Springer, 2008.

[16] Markus Degen, Peter Thiemann, and Stefan Wehr. Tracking linear and affine resources with JAVA(X). In *European Conference on Object-Oriented Programming*, pages 550–574. Springer, 2007.

[17] Robert DeLine and Manuel Fähndrich. Enforcing high-level protocols in low-level software. In *Proceedings of the Conference on Programming Language Design and Implementation*, pages 59–69, 2001.

[18] Manuel Fähndrich and Robert DeLine. Adoption and focus: Practical linear types for imperative programming. In *Proceedings of the Conference on Programming Language Design and Implementation*, pages 13–24, 2002.

[19] Stephen J. Fink, Eran Yahav, Nurit Dor, G. Ramalingam, and Emmanuel Geay. Effective typestate verification in the presence of aliasing. *ACM Transactions on Software Engineering and Methodology (TOSEM)*, 17(2):9, 2008.

[20] Cormac Flanagan. Hybrid type checking. In *Principles of Programming Languages*, pages 245–256, 2006.

[21] Cormac Flanagan, K. Rustan M. Leino, Mark Lillibridge, Greg Nelson, James B. Saxe, and Raymie Stata. Extended static checking for Java. In *Proceedings of the Conference on Programming Language Design and Implementation*, pages 234–245, 2002.

[22] Jeffrey S. Foster, Tachio Terauchi, and Alex Aiken. Flow-sensitive type qualifiers. In *Proceedings of the Conference on Programming Language Design and Implementation*, pages 1–12, 2002.

[23] Tim Freeman and Frank Pfenning. Refinement types for ML. In *Programming Language Design and Implementation*, pages 268–277, 1991.

[24] Thomas Gilray and Matthew Might. A survey of polyvariance in abstract interpretations. In *International Symposium on Trends in Functional Programming*, pages 134–148. Springer, 2013.

[25] Colin S. Gordon, Michael D. Ernst, and Dan Grossman. Rely-guarantee references for refinement types over aliased mutable data. In *Proceedings of the Conference on Programming Language Design and Implementation*, pages 73–84, 2013.

[26] Colin S Gordon, Michael D Ernst, Dan Grossman, and Matthew J. Parkinson. Verifying invariants of lock-free data structures with rely-guarantee and refinement types. *ACM Transactions on Programming Languages and Systems (TOPLAS)*, 39(3):11, 2017.

[27] Ben Hardekopf, Ben Wiedermann, Berkeley Churchill, and Vineeth Kashyap. Widening for control-flow. In *International Conference on Verification, Model Checking, and Abstract Interpretation*, 2014.

[28] Chris Hawblitzel, Jon Howell, Manos Kapritsos, Jacob R Lorch, Bryan Parno, Michael L Roberts, Srinath Setty, and Brian Zill. IronFleet: proving practical distributed systems correct. In *Proceedings of the 25th Symposium on Operating Systems Principles*, pages 1–17. ACM, 2015.

[29] Temesghen Kahsai, Rody Kersten, Philipp Rümmer, and Martin Schäf. Quantified heap invariants for object-oriented programs. In *LPAR-21, May 7–12, 2017, Maun, Botswana*, pages 368–384, 2017.

[30] Temesghen Kahsai, Philipp Rümmer, Huascar Sanchez, and Martin Schäf. JayHorn: A framework for verifying Java programs. In *International Conference on Computer Aided Verification*, pages 352–358. Springer, 2016.

[31] Vineeth Kashyap, Kyle Dewey, Ethan A. Kuefner, John Wagner, Kevin Gibbons, John Sarracino, Ben Wiedermann, and Ben Hardekopf. JSAI: a static analysis platform for JavaScript. In *Proceedings of the Conference on Foundations of Software Engineering*, 2014.

[32] Ming Kawaguchi, Patrick Rondon, and Ranjit Jhala. Type-based data structure verification. In *PLDI*, 2009.

[33] Anvesh Komuravelli, Arie Gurfinkel, Sagar Chaki, and Edmund M. Clarke. Automatic abstraction in SMT-based unbounded software model checking. In *International Conference on Computer Aided Verification*, pages 846–862. Springer, 2013.

[34] K. Rustan M. Leino. Dafny: An automatic program verifier for functional correctness. In *International Conference on Logic for Programming Artificial Intelligence and Reasoning*, pages 348–370. Springer, 2010.

[35] Yusuke Matsushita, Takeshi Tsukada, and Naoki Kobayashi. RustHorn: CHC-based verification for Rust programs. Submitted for publication, 2019.

[36] Ana Milanova, Atanas Rountev, and Barbara G. Ryder. Parameterized object sensitivity for points-to analysis for Java. *ACM Transactions on Software Engineering and Methodology (TOSEM)*, 14(1):1–41, 2005.

[37] Benjamin C Pierce. *Types and programming languages*. MIT press, 2002.

[38] Jonathan Protzenko, Jean-Karim Zinzindohoué, Aseem Rastogi, Tahina Ramananandro, Peng Wang, Santiago Zanella-Béguelin, Antoine Delignat-Lavaud, Cătălin Hriţcu, Karthikeyan Bhargavan, Cédric Fournet, and Nikhil Swamy. Verified low-level programming embedded in F*. *Proceedings of the ACM on Programming Languages*, 1(ICFP):17, 2017.

[39] Patrick Rondon, Ming Kawaguchi, and Ranjit Jhala. Low-level liquid types. In *Proceedings of the Symposium on Principles of Programming Languages*, pages 131–144, 2010.

[40] Patrick M. Rondon, Ming Kawaguci, and Ranjit Jhala. Liquid types. In *Proceedings of the Conference on Programming Language Design and Implementation*, pages 159–169, 2008.

[41] Philipp Rümmer, Hossein Hojjat, and Viktor Kuncak. Disjunctive interpolants for Horn-clause verification. In *International Conference on Computer Aided Verification*, pages 347–363. Springer, 2013.

[42] Micha Sharir and Amir Pnueli. Two approaches to interprocedural data flow analysis. In Stephen S. Muchnick and Neil D. Jones, editors, *Program Flow Analysis: Theory and Applications*, chapter 7, pages 189–223. Prentice Hall, 1981.

[43] Olin Shivers. *Control-flow analysis of higher-order languages*. PhD thesis, Carnegie Mellon University, 1991.

[44] Yannis Smaragdakis, Martin Bravenboer, and Ondřej Lhoták. Pick your contexts well: Understanding object-sensitivity. In *Proceedings of the Symposium on Principles of Programming Languages*, pages 17–30, 2011.

[45] Frederick Smith, David Walker, and Greg Morrisett. Alias types. In *European Symposium on Programming*, pages 366–381. Springer, 2000.

[46] Johannes Späth, Karim Ali, and Eric Bodden. Context-, flow-, and field-sensitive data-flow analysis using synchronized pushdown systems. *Proceedings of the ACM on Programming Languages*, 3(POPL):48, 2019.

[47] Johannes Späth, Lisa Nguyen Quang Do, Karim Ali, and Eric Bodden. Boomerang: Demand-driven flow-and context-sensitive pointer analysis for Java. In *30th European Conference on Object-Oriented Programming (ECOOP 2016)*. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik, 2016.

[48] Kohei Suenaga, Ryota Fukuda, and Atsushi Igarashi. Type-based safe resource deallocation for shared-memory concurrency. In *Proceedings of the ACM International Conference on Object Oriented Programming Systems Languages and Applications*, pages 1–20, 2012.

[49] Kohei Suenaga and Naoki Kobayashi. Fractional ownerships for safe memory deallocation. In *Asian Symposium on Programming Languages and Systems*, pages 128–143. Springer, 2009.

[50] Nikhil Swamy, Cătălin Hriţcu, Chantal Keller, Aseem Rastogi, Antoine Delignat-Lavaud, Simon Forest, Karthikeyan Bhargavan, Cédric Fournet, Pierre-Yves Strub, Markulf Kohlweiss, Jean-Karim Zinzindohoué, and Santiago Zanella-Béguelin. Dependent types and multi-monadic effects in F*. In *Proceedings of the Symposium on Principles of Programming Languages*, pages 256–270, 2016.

[51] Nikhil Swamy, Joel Weinberger, Cole Schlesinger, Juan Chen, and Benjamin Livshits. Verifying higher-order programs with the Dijkstra monad. In *Proceedings of the Conference on Programming Language Design and Implementation*, pages 387–398, 2013.

[52] Tachio Terauchi. Checking race freedom via linear programming. In *Proceedings of the Conference on Programming Language Design and Implementation*, pages 1–10, 2008.

[53] Hiroshi Unno and Naoki Kobayashi. Dependent type inference with interpolants. In *Proceedings of the Conference on Principles and Practice of Declarative Programming*, pages 277–288. ACM, 2009.

[54] Niki Vazou, Patrick M. Rondon, and Ranjit Jhala. Abstract refinement types. In *European Symposium on Programming*, pages 209–228. Springer, 2013.

[55] Niki Vazou, Eric L. Seidel, Ranjit Jhala, Dimitrios Vytiniotis, and Simon Peyton-Jones. Refinement types for Haskell. In *Proceedings of the International Conference on Functional Programming*, pages 269–282, 2014.

[56] J. B. Wells, Allyn Dimock, Robert Muller, and Franklyn Turbak. A calculus with polymorphic and polyvariant flow types. *Journal of Functional Programming*, 12(3):183–227, 2002.

[57] Hongwei Xi and Frank Pfenning. Dependent types in practical programming. In *Proceedings of the Symposium on Principles of Programming Languages*, pages 214–227. ACM, 1999.

[58] Pamela Zave. Using lightweight modeling to understand Chord. *ACM SIGCOMM Computer Communication Review*, 42(2):49–57, 2012.

[59] He Zhu and Suresh Jagannathan. Compositional and lightweight dependent type inference for ML. In *International Conference on Verification, Model Checking, and Abstract Interpretation*, pages 295–314. Springer, 2013.

$$\frac{\Theta \mid [] : \tau \Rightarrow \Gamma \mid \mathcal{L} \vdash_{ectx} E : \tau' \Rightarrow \Gamma' \qquad \Theta \mid \mathcal{L} \mid \Gamma' \vdash e : \tau'' \Rightarrow \Gamma''}{\Theta \mid [] : \tau \Rightarrow \Gamma \mid \mathcal{L} \vdash_{ectx} E\,;e : \tau'' \Rightarrow \Gamma''} \quad \text{(TE-Seq)}$$

$$\frac{}{\Theta \mid [] : \tau \Rightarrow \Gamma \mid \mathcal{L} \vdash_{ectx} [] : \tau \Rightarrow \Gamma} \quad \text{(TE-Hole)}$$

$$\frac{\begin{array}{c}\Theta \mid [] : \tau' \Rightarrow \Gamma' \mid \mathcal{L} \vdash_{ectx} E : \tau'' \Rightarrow \Gamma'' \\ \Theta \mid \mathcal{L} \mid \Gamma, x : \tau \vdash e : \tau' \Rightarrow \Gamma' \\ x \notin dom(\Gamma')\end{array}}{\Theta \mid [] : \tau \Rightarrow \Gamma \mid \mathcal{L} \vdash_{ectx} E[\mathbf{let}\ x = []^{\ell}\ \mathbf{in}\ e] : \tau'' \Rightarrow \Gamma''} \quad \text{(TE-Stack)}$$

$$\begin{array}{c}(E\,;e)[e'] = E[e']\,;e \\ [][e'] = e' \\ E[\mathbf{let}\ y = []^{\ell}\ \mathbf{in}\ e][x] = E[\mathbf{let}\ y = x\ \mathbf{in}\ e]\end{array}$$

Fig. 11. Context typing and substitution

A Proof of Type Soundness (Theorem 1)

We first define a typing relation for machine configurations $\left\langle H, R, \vec{F}, e \right\rangle$ as shown in Figure 12. The critical component of this typing relation is the consistency relation **Cons**. Intuitively, **Cons** expresses that the current heap and registers are consistent with the ownership and refinement information implied by Γ. We say triple (H, R, Γ) is *consistent*, and write $\mathbf{Cons}\,(H, R, \Gamma)$. In the definitions for **own** we write $\{a \mapsto r\}$ to denote a function $\mathbf{Addr} \to [0, 1]$ which returns r for a, and 0 otherwise. We write $\emptyset$ to denote a constant function $\mathbf{Addr} \to [0, 1]$ which always returns 0. We define the addition between two functions $O_1, O_2 : \mathbf{Addr} \to [0, 1]$ as: $(O_1 + O_2)(a) = O_1(a) + O_2(a)$. Finally, if a summation Σ has no summands, we take its result to be $\emptyset$.

The proof of Theorem 1 requires the following four key lemmas. These lemmas are stated with respect to some well-typed program $\langle D, e\rangle$, i.e. $\vdash \langle D, e\rangle$.

Lemma 1. $\vdash^{D}_{conf} \langle \emptyset, \emptyset, \cdot, e\rangle$

Proof. Trivial, taking $\Gamma = \bullet$ and by inversion on $\vdash \langle D, e\rangle$.

Lemma 2. $\vdash^{D}_{conf} \mathbf{C}$ *implies* $\mathbf{C} \neq \mathbf{AssertFail}$

Proof. Simple proof by contradiction, as the **AssertFail** is not well-typed.

Lemma 3. *If* $\vdash^{D}_{conf} \left\langle H, R, \vec{F}, e \right\rangle$ *and* $\left\langle H, R, \vec{F}, e \right\rangle \longrightarrow_D \mathbf{C}$*, then* $\vdash^{D}_{conf} \mathbf{C}$

$$\frac{\vec{\ell} = \mathbf{Trace}\,(\vec{F}) \qquad n = |\vec{\ell}| = |\vec{F}| \qquad \Theta \vdash D \qquad \forall j \in \{1..n\}\,.\vec{\ell}_j = tail^{n-j+1}(\vec{\ell}) \\ \mathbf{Cons}\,(H, R, \Gamma) \qquad \forall i \in \{1..n\}\,.\Theta \mid [] : \tau_i \Rightarrow \Gamma_i \mid \vec{\ell}_i \vdash_{ectx} F_i : \tau_{i-1} \Rightarrow \Gamma_{i-1} \\ \vec{F} = F_n : \cdots : F_1 : \cdot \qquad \Theta \mid \vec{\ell} \mid \Gamma \vdash e : \tau_n \Rightarrow \Gamma_n}{\vdash^D_{conf} \left\langle H, R, \vec{F}, e \right\rangle}$$

$$\vdash^D_{conf} \mathbf{AliasFail}$$

$$\mathbf{Trace}\,(\cdot) = \epsilon$$
$$\mathbf{Trace}\,(E[\mathbf{let}\ x = []^\ell\ \mathbf{in}\ e] : \vec{F}) = \ell : \mathbf{Trace}\,(\vec{F})$$

$$\mathbf{Cons}\,(H, R, \Gamma) \overset{\text{def}}{\iff} \mathbf{SAT}\,(H, R, \Gamma) \wedge \forall\, a \in dom(H).\mathbf{Own}\,(H, R, \Gamma)(a) \leq 1$$
$$\mathbf{SAT}\,(H, R, \Gamma) \overset{\text{def}}{\iff} \forall\, x \in dom(\Gamma).x \in dom(R) \wedge \mathbf{SATv}(H, R, R(x), \Gamma(x))$$
$$\mathbf{SATv}(H, R, v, \tau) \overset{\text{def}}{\iff} \begin{cases} v \in \mathbb{Z} \wedge [R]\,[v/\nu]\varphi & \tau = \{\nu : \mathbf{int} \mid \varphi\} \\ a \in dom(H) \wedge \mathbf{SATv}(H, R, H(a), \tau') & \tau = \tau'\ \mathbf{ref}^r \wedge v\ =\ a \end{cases}$$
$$[\emptyset]\,\varphi = \varphi$$
$$[R\{y \mapsto n\}]\,\varphi = [R]\,[n/y]\varphi$$
$$[R\{y \mapsto a\}]\,\varphi = [R]\,\varphi$$
$$\mathbf{Own}\,(H, R, \Gamma) = \Sigma_{x \in dom(\Gamma)}\,\mathbf{own}\,(H, R(x), \Gamma(x))$$
$$\mathbf{own}\,(H, v, \tau) = \begin{cases} \{a \mapsto r\} + \mathbf{own}\,(H, H(a), \tau') & v\ =\ a \wedge a \in dom(H) \wedge \tau = \tau'\ \mathbf{ref}^r \\ \emptyset & o.w. \end{cases}$$

Fig. 12. Machine state typing

Lemma 4. *If* $\vdash^{D}_{conf} \mathbf{C}$*, then, one of the following conditions hold:*

1. $\exists \mathbf{C}', \mathbf{C} \longrightarrow_D \mathbf{C}'$*, or*
2. $\mathbf{C} = \mathbf{AssertFail}$*, or*
3. $\mathbf{C} = \langle H, R, \cdot, x \rangle$

Lemmas 3 and 4 are the heart of proof effort, we give their proofs in Appendices C and D respectively.

We can now prove Theorem 1:

Proof (Theorem 1: Soundness). From Lemmas 1 and 3 and an inductive argument, any configuration reachable from the initial state must be well-typed. Then, by Lemma 2 every configuration reachable from the initial state cannot be **AssertFail**, i.e., a well-typed program never experiences an assertion failure. This completes the first part of the proof.

To prove the second portion of the theorem, it suffices to show that any configuration reachable from the initial state can step or is a final configuration. Again from Lemmas 1 and 3 and a simple inductive argument, we must have that for any state $\mathbf{C}$ such that $\langle \emptyset, \emptyset, \cdot, e \rangle \longrightarrow^{*}_{D} \mathbf{C} \vdash^{D}_{conf} \mathbf{C}$. Then by Lemma 4 we have the configuration may step or is one of the final configurations.

The remainder of this appendix proves Lemmas 3 and 4. We introduce some auxiliary definitions and lemmas in Appendix B, give the proof of Lemma 3 in Appendix C, and prove Lemma 4 in Appendix D.

B Auxiliary Lemmas and Definitions

The well-formedness rules omitted from the main paper are found in Figure 13. We write $\mathcal{L} \vdash_{WF} \tau \Rightarrow \Gamma$ as shorthand for $\mathcal{L} \vdash_{WF} \Gamma$ and $\mathcal{L} \mid \Gamma \vdash_{WF} \tau$.

We first prove that the subtyping relations are transitive.

Lemma 5.

1. *If* $\Gamma \leq \Gamma'$ *then* $\models [\![\Gamma]\!] \implies [\![\Gamma']\!]$.
2. *If* $\Gamma \vdash \tau_1 \leq \tau_2$ *and* $\Gamma \vdash \tau_2 \leq \tau_3$*, then* $\Gamma \vdash \tau_1 \leq \tau_3$
3. *If* $\Gamma \leq \Gamma'$ *and* $\Gamma' \vdash \tau_1 \leq \tau_2$*, then* $\Gamma \vdash \tau_1 \leq \tau_2$
4. *If* $\Gamma \leq \Gamma'$*,* $\Gamma \vdash \tau_1 \leq \tau_2$*, and* $\Gamma' \vdash \tau_2 \leq \tau_3$*, then* $\Gamma \vdash \tau_1 \leq \tau_3$.
5. *If* $\Gamma \leq \Gamma'$ *and* $\Gamma' \leq \Gamma''$*, then* $\Gamma \leq \Gamma''$.

Proof.

1. It suffices to show that $\models [\![\Gamma]\!] \implies [x/\nu]\,\varphi'$ for any $x \in dom(\Gamma')$ where $\Gamma'(x) = \{\nu : \mathbf{int} \mid \varphi'\}$. From $\Gamma \leq \Gamma'$ we have $\models [\![\Gamma]\!] \wedge \varphi \implies \varphi'$ where $\Gamma(x) = \{\nu : \mathbf{int} \mid \varphi\}$. We must then have $\models [\![\Gamma]\!] \wedge [x/\nu]\,\varphi \implies [x/\nu]\,\varphi'$. From the definition of $[\![\Gamma]\!]$ we have $[\![\Gamma]\!] \wedge [x/\nu]\,\varphi \iff [\![\Gamma]\!]$, giving the desired result.

$$\frac{\forall\, x \in dom(\Gamma).\mathcal{L} \mid \Gamma \vdash_{WF} \Gamma(x)}{\mathcal{L} \vdash_{WF} \Gamma}\ \text{(WF-Env)} \qquad \frac{\mathcal{L} \mid \Gamma \vdash_{WF} \varphi}{\mathcal{L} \mid \Gamma \vdash_{WF} \{\nu : \mathbf{int} \mid \varphi\}}\ \text{(WF-Int)} \qquad \frac{\mathcal{L} \mid \Gamma \vdash_{WF} \tau}{\mathcal{L} \mid \Gamma \vdash_{WF} \tau\, \mathbf{ref}^r}\ \text{(WF-Ref)}$$

$$\frac{\forall\, x \in \mathbf{FPV}(\varphi) \setminus \{\nu\}\,.\Gamma(x) = \{\nu : \mathbf{int} \mid \_\} \quad \mathbf{FCV}(\varphi) \subseteq \mathbf{CV}(\mathcal{L})}{\mathcal{L} \mid \Gamma \vdash_{WF} \varphi}\ \text{(WF-Phi)} \qquad \frac{\mathcal{L} \mid \Gamma \vdash_{WF} \tau \qquad \mathcal{L} \vdash_{WF} \Gamma}{\mathcal{L} \vdash_{WF} \tau \Rightarrow \Gamma}\ \text{(WF-Result)}$$

$$\frac{\lambda \vdash_{WF} x_1 : \tau_1, \ldots, x_n : \tau_n \quad \lambda \vdash_{WF} \tau \Rightarrow x_1 : \tau_1', \ldots, x_n : \tau_n'}{\vdash_{WF} \forall \lambda.\, \langle x_1 : \tau_1, \ldots, x_n : \tau_n \rangle \rightarrow \langle x_1 : \tau_1', \ldots, x_n : \tau_n' \mid \tau \rangle}\ \text{(WF-FunType)} \qquad \frac{\forall\, f \in dom(\Theta).\, \vdash_{WF} \Theta(f)}{\vdash_{WF} \Theta}\ \text{(WF-FunEnv)}$$

$$\begin{aligned}
\text{Free Ctxt Vars}\quad \mathbf{FCV}(\varphi_1 \vee \varphi_2) &= \mathbf{FCV}(\varphi_1) \cup \mathbf{FCV}(\varphi_2)\\
\mathbf{FCV}(\neg\varphi) &= \mathbf{FCV}(\varphi)\\
\mathbf{FCV}(\widehat{v}_1 = \widehat{v}_2) &= \mathbf{FCV}(\phi(\widehat{v}_1, .., \widehat{v}_n)) = \emptyset\\
\mathbf{FCV}(\vec{\ell} \subseteq \mathcal{C}) &= \mathbf{FCV}(\mathcal{C})\\
\mathbf{FCV}(\ell : \mathcal{C}) &= \mathbf{FCV}(\mathcal{C})\\
\mathbf{FCV}(\mathcal{L}) &= \mathbf{CV}(\mathcal{L})\\
\text{Ctxt Vars}\quad \mathbf{CV}(\vec{\ell}) &= \emptyset\\
\mathbf{CV}(\lambda) &= \{\lambda\}
\end{aligned}$$

Fig. 13. Well-formedness of types and environments.

2. By induction on $\Gamma \vdash \tau_1 \leq \tau_2$. We only consider the base case where $\tau_1 = \{\nu : \mathbf{int} \mid \varphi_1\}$ and $\tau_2 = \{\nu : \mathbf{int} \mid \varphi_2\}$, the case for reference types follows from the induction hypothesis. By further inversion on $\Gamma \vdash \tau_2 \leq \tau_3$ we therefore have:
   $\tau_3 = \{\nu : \mathbf{int} \mid \varphi_3\}$
   $\models [\![\Gamma]\!] \wedge \varphi_1 \implies \varphi_2 \qquad \models [\![\Gamma]\!] \wedge \varphi_2 \implies \varphi_3$
   From which it is immediate that we must have $\models [\![\Gamma]\!] \wedge \varphi_1 \implies \varphi_3$, whereby S-Int gives $\Gamma \vdash \tau_1 \leq \tau_3$.
3. By induction on $\Gamma' \vdash \tau_1 \leq \tau_2$. The case for reference types is immediate from the inductive hypothesis, we focus on the base case where $\tau_1 = \{\nu : \mathbf{int} \mid \varphi_1\}$ and $\tau_2 = \{\nu : \mathbf{int} \mid \varphi_2\}$, and where $\models [\![\Gamma']\!] \wedge \varphi_1 \implies \varphi_2$. From $\Gamma \leq \Gamma'$ and Item 1 above, we have $\models [\![\Gamma]\!] \implies [\![\Gamma']\!]$ from which we can derive $[\![\Gamma]\!] \wedge \varphi_1 \implies \varphi_2$, i.e., $\Gamma \vdash \tau_1 \leq \tau_2$.
4. Immediate from Items 2 and 3.
5. Immediate corollary of Item 4.

Definition 1. *A value v reaches an integer with n dereferences in heap H when it is in the relation $H \vdash v \Downarrow n$ defined as the smallest relation closed under the following rules:*

1. *If $v \in \mathbb{Z}$ then $H \vdash v \Downarrow 0$*
2. *If $H \vdash v \Downarrow n$ and $H(a) = v$ then $H \vdash a \Downarrow n + 1$*

We will write $H \vdash v \Downarrow |\tau|$ to indicate a value v is shape consistent with τ in heap H, where $|\tau|$ is the number of reference constructors in the type τ.

We also prove a standard inversion lemma to handle the fact our typing rules are not syntax directed.

Lemma 6 (Inversion). *If* $\Theta \mid \mathcal{L} \mid \Gamma \vdash e_0 : \tau \Rightarrow \Gamma'$*, then there exists some* Γ_p*,* τ_p*, and* Γ'_p *such that* $\Gamma \leq \Gamma_p$*,* $\vec{\ell} \vdash_{WF} \Gamma_p$*,* $\Gamma'_p, \tau_p \leq \Gamma', \tau$*, and:*

1. *If* $e_0 = x$ *then* $\Gamma_p(x) = \tau_p + \tau'$*,* $\Gamma'_p = \Gamma_p[x \hookleftarrow \tau']$*.*
2. *If* $e_0 = \mathbf{let}\, x = y \,\mathbf{in}\, e$*, then* $\Theta \mid \mathcal{L} \mid \Gamma_p[y \hookleftarrow \tau_1 \wedge_y y =_{\tau_1} x], x : (\tau_2 \wedge_x x =_{\tau_2} y) \vdash e : \tau_p \Rightarrow \Gamma'_p$ *and* $x \notin dom(\Gamma'_p)$ *where* $\Gamma_p(y) = \tau_1 + \tau_2$*.*
3. *If* $e_0 = \mathbf{let}\, x = n \,\mathbf{in}\, e$ *then* $\Theta \mid \mathcal{L} \mid \Gamma_p, x : \{\nu : \mathbf{int} \mid \nu = n\} \vdash e : \tau_p \Rightarrow \Gamma'_p$ *and* $x \notin dom(\Gamma'_p)$*.*
4. *If* $e_0 = \mathbf{ifz}\, x \,\mathbf{then}\, e_1 \,\mathbf{else}\, e_2$ *then:*
   - $\Gamma_p(x) = \{\nu : \mathbf{int} \mid \varphi\}$
   - $\Theta \mid \mathcal{L} \mid \Gamma_p[x \hookleftarrow \{\nu : \mathbf{int} \mid \varphi \wedge \nu = 0\}] \vdash e_1 : \tau_p \Rightarrow \Gamma'_p$
   - $\Theta \mid \mathcal{L} \mid \Gamma_p[x \hookleftarrow \{\nu : \mathbf{int} \mid \varphi \wedge \nu \neq 0\}] \vdash e_2 : \tau_p \Rightarrow \Gamma'_p$
5. *If* $e_0 = \mathbf{let}\, x = \mathbf{mkref}\, y \,\mathbf{in}\, e$*, then* $\Gamma_p(y) = \tau_1 + \tau_2$*,* $\Theta \mid \mathcal{L} \mid \Gamma[y \hookleftarrow \tau_1], x : (\tau_2 \wedge_x x =_{\tau_2} y)\, \mathbf{ref}^1 \vdash e : \tau \Rightarrow \Gamma'_p$*, and* $x \notin dom(\Gamma'_p)$
6. *If* $e_0 = \mathbf{let}\, x = *y \,\mathbf{in}\, e$*, then:*
   - $\Gamma_p(y) = \tau_1 + \tau_2\, \mathbf{ref}^r$
   - $\Theta \mid \mathcal{L} \mid \Gamma_p[y \hookleftarrow \tau''\, \mathbf{ref}^r], x : \tau_2 \vdash e : \tau_p \Rightarrow \Gamma'_p$
   - $x \notin dom(\Gamma'_p)$
   - $$\tau'' = \begin{cases} (\tau_1 \wedge_y y =_{\tau_1} x) & r > 0 \\ \tau_1 & r = 0 \end{cases}$$
7. *If* $e_0 = \mathbf{let}\, x = f^{\ell}(y_1, \ldots, y_n) \,\mathbf{in}\, e$ *then:*
   - $\Gamma_p(y_i) = \sigma_\alpha\, \sigma_x\, \tau_i$ *for each* $i \in \{1, \ldots, n\}$
   - $\Theta \mid \mathcal{L} \mid \Gamma_p[y_i \hookleftarrow \sigma_\alpha\, \sigma_x\, \tau'_i], x : \sigma_\alpha\, \sigma_x\, \tau \vdash e : \tau_p \Rightarrow \Gamma'_p$
   - $\Theta(f) = \forall \lambda.\, \langle x_1 : \tau_1, \ldots, x_n : \tau_n \rangle \to \langle x_1 : \tau'_1, \ldots, x_n : \tau'_n \mid \tau \rangle$
   - $\sigma_\alpha = [\ell : \mathcal{L}/\lambda]$
   - $\sigma_x = [y_1/x_1] \cdots [y_n/x_n]$
   - $x \notin dom(\Gamma'_p)$
8. *If* $e_0 = y := x\,; e$ *then:*
   - $\Gamma_p(x) = \tau_1 + \tau_2$
   - $\Gamma_p(y) = \tau'\, \mathbf{ref}^1$
   - $\Theta \mid \mathcal{L} \mid \Gamma_p[x \hookleftarrow \tau_1][y \hookleftarrow (\tau_2 \wedge_y y =_{\tau_2} x)\, \mathbf{ref}^1] \vdash e : \tau_p \Rightarrow \Gamma'_p$
   - *The shapes of* τ' *and* τ_2 *are similar, i.e,* $|\tau'| = |\tau_2|$*.*
9. *If* $e_0 = \mathbf{alias}(x = y)\,; e$ *then there exist some* $\tau_1, \tau_2, \tau'_1, \tau'_2, r_1, r_2, r'_1, r'_2$ *such that:*
   - $\tau_1\, \mathbf{ref}^{r_1} + \tau_2\, \mathbf{ref}^{r_2} \approx \tau'_1\, \mathbf{ref}^{r'_1} + \tau'_2\, \mathbf{ref}^{r'_2}$
   - $\Gamma_p(x) = \tau_1\, \mathbf{ref}^{r_1}$ *and* $\Gamma_p(y) = \tau_2\, \mathbf{ref}^{r_2}$
   - $\Theta \mid \mathcal{L} \mid \Gamma[x \hookleftarrow \tau'_1\, \mathbf{ref}^{r'_1}][y \hookleftarrow \tau'_2\, \mathbf{ref}^{r'_2}] \vdash e : \tau_p \Rightarrow \Gamma'_p$
10. *If* $e_0 = \mathbf{alias}(x = *y)\,; e$ *then there exist some* $\tau_1, \tau_2, \tau'_1, \tau'_2, r_1, r_2, r'_1, r'_2, r$*, such that:*
    - $\tau_1\, \mathbf{ref}^{r_1} + \tau_2\, \mathbf{ref}^{r_2} \approx \tau'_1\, \mathbf{ref}^{r'_1} + \tau'_2\, \mathbf{ref}^{r'_2}$
    - $\Gamma_p(x) = \tau_1\, \mathbf{ref}^{r_1}$ *and* $\Gamma_p(y) = (\tau_2\, \mathbf{ref}^{r_2})\, \mathbf{ref}^{r}$

- $\Theta \mid \mathcal{L} \mid \Gamma[x \hookleftarrow \tau_1' \, \mathbf{ref}^{r_1'}][y \hookleftarrow (\tau_2' \, \mathbf{ref}^{r_2'}) \, \mathbf{ref}^{r}] \vdash e : \tau_p \Rightarrow \Gamma_p'$

11. *If* $e_0 = e_1 \,;\, e_2$ *then* $\Theta \mid \mathcal{L} \mid \Gamma_p \vdash e_1 : \tau_1 \Rightarrow \Gamma_1$ *and* $\Theta \mid \mathcal{L} \mid \Gamma_1 \vdash e_2 : \tau_p \Rightarrow \Gamma_p'$
12. *If* $e_0 = x \,;\, e'$ *then* $\Theta \mid \mathcal{L} \mid \Gamma_p[x : \tau' + \tau_0] \vdash x : \tau_1 \Rightarrow \Gamma_p[x \hookleftarrow \tau_0]$ *and* $\Theta \mid \mathcal{L} \mid \Gamma_p[x \hookleftarrow \tau_0] \vdash e' : \tau_p \Rightarrow \Gamma_p'$
13. *If* $e_0 = \mathbf{assert}(\varphi) \,;\, e$ *then* $\Gamma_p \models \varphi$ *and* $\Theta \mid \mathcal{L} \mid \Gamma_p \vdash e : \tau_p \Rightarrow \Gamma_p'$

Proof. By straightforward induction on the typing relation and the transitivity of the subtyping relation Lemma 5.

The only case of note is the case for $e_0 = x \,;\, e_2$. If the subderivation for x has applications of T-Sub then the subtypings on the output environment can be pushed into application subtyping on input environments when typing e'. Similarly, any input subtypings on the input environment of the derivation of x can be pushed into T-Sub rules such that $\Gamma \leq \Gamma_p[x : \tau' + \tau_0]$.

Lemmas 7 and 8 prove some standard properties of execution contexts: any decomposition of a well-typed expression into a execution context and redex can be well-typed, and substituting a well-typed expression matching a context's hole type yields a well-typed expression

Lemma 7. *For any* E *and* e' *such that* $E[e'] = e$ *where* $\Theta \mid \mathcal{L} \mid \Gamma \vdash e : \tau \Rightarrow \Gamma'$ *there exists some* τ_0, Γ_0 *such that* $\Theta \mid [] : \tau_0 \Rightarrow \Gamma_0 \mid \mathcal{L} \vdash_{ectx} E : \tau \Rightarrow \Gamma'$ *and* $\Theta \mid \mathcal{L} \mid \Gamma \vdash e' : \tau_0 \Rightarrow \Gamma_0$.

Proof. By induction on the structure of E.

Case $E = []$**:**

Trivial, by taking $\tau_0 = \tau$ and $\Gamma_0 = \Gamma'$.

Case $E = E' \,;\, e''$**:**

Then $E[e'] = E'[e'] \,;\, e'' = e$. By Lemma 6 we have

$$\begin{array}{ll} \Theta \mid \mathcal{L} \mid \Gamma_p \vdash E'[e'] : \tau_1 \Rightarrow \Gamma_1 & \Theta \mid \mathcal{L} \mid \Gamma_1 \vdash e'' : \tau_p \Rightarrow \Gamma_p' \\ \Gamma \leq \Gamma_p & \Gamma_p', \tau_p \leq \Gamma', \tau \end{array}$$

for some Γ_p, Γ_p', and τ_p.

By the induction hypothesis we then have $\Theta \mid \mathcal{L} \mid \Gamma_p \vdash e' : \tau_0 \Rightarrow \Gamma_0$ and $\Theta \mid [] : \tau_0 \Rightarrow \Gamma_0 \mid \mathcal{L} \vdash_{ectx} E' : \tau_1 \Rightarrow \Gamma_1$. for some τ_0 and Γ_0.

Next, as $\Gamma_p', \tau_p \leq \Gamma', \tau$ by an application of T-Sub, we have $\Theta \mid \mathcal{L} \mid \Gamma_1 \vdash e'' : \tau \Rightarrow \Gamma'$. By TE-Seq, we therefore have: $\Theta \mid [] : \tau_0 \Rightarrow \Gamma_0 \mid \mathcal{L} \vdash_{ectx} E' \,;\, e'' : \tau \Rightarrow \Gamma'$.

Finally, from $\Gamma \leq \Gamma_p$ and $\Theta \mid \mathcal{L} \mid \Gamma_p \vdash e' : \tau_0 \Rightarrow \Gamma_0$, and application of T-Sub, we have $\Theta \mid \mathcal{L} \mid \Gamma \vdash e' : \tau_0 \Rightarrow \Gamma_0$.

Lemma 8. *If* $\Theta \mid [] : \tau \Rightarrow \Gamma' \mid \mathcal{L} \vdash_{ectx} E : \tau'' \Rightarrow \Gamma''$ *and* $\Theta \mid \mathcal{L} \mid \Gamma \vdash e : \tau \Rightarrow \Gamma'$ *for some* Γ*, then* $\Theta \mid \mathcal{L} \mid \Gamma \vdash E[e] : \tau'' \Rightarrow \Gamma''$.

Proof. By induction on the typing derivation of E.

Case TE-Seq**:** $E = E' ; e'$
$E[e] = E'[e] ; e'$
$\Theta \mid [] : \tau \Rightarrow \Gamma' \mid \mathcal{L} \vdash_{ectx} E' : \tau_0 \Rightarrow \Gamma_0$
$\Theta \mid \mathcal{L} \mid \Gamma_0 \vdash e' : \tau'' \Rightarrow \Gamma''$

By the induction hypothesis we have $\Theta \mid \mathcal{L} \mid \Gamma \vdash E'[e] : \tau_0 \Rightarrow \Gamma_0$. We then have our result via an application of T-Seq.

Case TE-Hole**:**

Trivial, as $\tau = \tau''$ and $\Gamma' = \Gamma''$ and $E[e] = e$.

Lemma 9 (Context Variable Substitution).

1. *If* $\tau_3 = \tau_1 + \tau_2$ *then* $[\mathcal{L}/\lambda]\,\tau_3 = [\mathcal{L}/\lambda]\,\tau_1 + [\mathcal{L}/\lambda]\,\tau_2$
2. *For any* $\vec{\ell}$*:*
   (a) *If* $\lambda \vdash_{WF} \Gamma$ *then* $\vec{\ell} \vdash_{WF} [\vec{\ell}/\lambda]\Gamma$
   (b) *If* $\lambda \mid \Gamma \vdash_{WF} \tau$ *then* $\vec{\ell} \mid [\vec{\ell}/\lambda]\Gamma \vdash_{WF} [\vec{\ell}/\lambda]\,\tau$
   (c) *If* $\lambda \vdash_{WF} \tau \Rightarrow \Gamma$ *then* $\vec{\ell} \vdash_{WF} [\vec{\ell}/\lambda]\,\tau \Rightarrow [\vec{\ell}/\lambda]\Gamma$
3. *For any* Γ, τ_1, τ_2, λ *and* $\vec{\ell}$, *If* $\Gamma \vdash \tau_1 \leq \tau_2$, *then* $[\vec{\ell}/\lambda]\Gamma \vdash [\vec{\ell}/\lambda]\,\tau_1 \leq [\vec{\ell}/\lambda]\,\tau_2$
4. *If* $\Gamma \models \varphi$ *where* $\lambda \notin \mathbf{FCV}\,(\varphi)$ *then* $[\vec{\ell}/\lambda]\Gamma \models \varphi$
5. *If* $\Theta \mid \lambda \mid \Gamma \vdash e : \tau \Rightarrow \Gamma'$ *then* $\Theta \mid \vec{\ell} \mid [\vec{\ell}/\lambda]\Gamma \vdash e : [\vec{\ell}/\lambda]\,\tau \Rightarrow [\vec{\ell}/\lambda]\Gamma'$

Proof.

1. By straightforward induction on the definition of $\tau_1 + \tau_2 = \tau_3$.
2. Observe that any substitution of context variables cannot change simple types within Γ and thus all types and refinements remain well-formed with respect to integer variables in Γ. It thus suffices to show that $\mathbf{FCV}\,([\vec{\ell}/\lambda]\,\varphi) \subseteq \mathbf{CV}\,(\vec{\ell}) = \emptyset$ for any refinement φ appearing in τ or a type in Γ. By the assumed well-formedness of τ with respect to context variable λ (resp. Γ), after substitution all free context variables in τ (resp. the types in Γ) will be replaced with $\vec{\ell}$. Thus, post-substitution no free context variables appear in the refinement of $[\vec{\ell}/\lambda]\,\tau$ (resp. refinements of types in $[\vec{\ell}/\lambda]\Gamma$), trivially satisfying our requirements.
3. If λ does not appear free in τ_1, τ_2 or Γ, then the result trivially holds. Let us then assume λ appears free. We prove the result by induction on the subtyping derivation.

   Case S-Ref**:** $\tau_1 = \tau_1'\,\mathbf{ref}^{r_1}$ $\quad$ $\tau_2 = \tau_2'\,\mathbf{ref}^{r_2}$
   $[\vec{\ell}/\lambda]\,\tau_1 = ([\vec{\ell}/\lambda]\,\tau_1')\,\mathbf{ref}^{r_1}$ $\quad$ $[\vec{\ell}/\lambda]\,\tau_2 = ([\vec{\ell}/\lambda]\,\tau_2')\,\mathbf{ref}^{r_2}$
   $\Gamma \vdash \tau_1' \leq \tau_2'$ $\quad$ $r_2 \leq r_1$

   We must show that $[\vec{\ell}/\lambda]\Gamma \vdash [\vec{\ell}/\lambda]\,\tau_1' \leq [\vec{\ell}/\lambda]\,\tau_2'$ which holds immediately from the induction hypothesis.

Case S-Int**:** $\tau_1 = \{\nu : \mathbf{int} \mid \varphi_1\}$ $\quad \tau_2 = \{\nu : \mathbf{int} \mid \varphi_2\}$
$[\vec{\ell}/\lambda]\,\tau_1 = \left\{\nu : \mathbf{int} \mid [\vec{\ell}/\lambda]\,\varphi_1\right\}$ $\quad [\vec{\ell}/\lambda]\,\tau_2 = \left\{\nu : \mathbf{int} \mid [\vec{\ell}/\lambda]\,\varphi_2\right\}$
$\Gamma \models \varphi_1 \implies \varphi_2$

We must show that $[\vec{\ell}/\lambda]\Gamma \models [\vec{\ell}/\lambda]\,\varphi_1 \implies [\vec{\ell}/\lambda]\,\varphi_2$, i.e. $\models \left[\!\left[[\vec{\ell}/\lambda]\Gamma\right]\!\right] \wedge [\vec{\ell}/\lambda]\,\varphi_1 \implies [\vec{\ell}/\lambda]\,\varphi_2$. From our assumption that $\Gamma \models \varphi_1 \implies \varphi_2$ we have that $\models [\![\Gamma]\!] \wedge \varphi_1 \implies \varphi_2$ is valid, whereby the formula $[\![\Gamma]\!] \wedge \varphi_1 \implies \varphi_2$ is true for any possible concrete valuation of the free context variable λ. As $[\vec{\ell}/\lambda]\ [\![\Gamma]\!]$ is equivalent to $\left[\!\left[[\vec{\ell}/\lambda]\Gamma\right]\!\right]$ we have the formula $\left[\!\left[[\vec{\ell}/\lambda]\Gamma\right]\!\right] \wedge [\vec{\ell}/\lambda]\,\varphi_1 \implies [\vec{\ell}/\lambda]\,\varphi_2$ must also be valid.

4. If λ does not appear free in $[\![\Gamma]\!]$, then the result trivially holds. Otherwise $\models [\![\Gamma]\!] \implies \varphi$ holds for any concrete valuation of the free context variable λ. Then the formula $\models \left[\!\left[[\vec{\ell}/\lambda]\Gamma\right]\!\right] \implies \varphi$ must be valid from the equivalence of $[\vec{\ell}/\lambda]\ [\![\Gamma]\!]$ and $\left[\!\left[[\vec{\ell}/\lambda]\Gamma\right]\!\right]$.
5. By induction on the typing derivation $\Theta \mid \lambda \mid \Gamma \vdash e : \tau \Rightarrow \Gamma'$. In every case, that $\vec{\ell} \vdash_{WF} [\vec{\ell}/\lambda]\,\tau \Rightarrow [\vec{\ell}/\lambda]\Gamma'$ and $\vec{\ell} \vdash_{WF} [\vec{\ell}/\lambda]\Gamma$ holds from Item 2.

   Case T-Var**:** $e = x$ $\quad \tau = \tau_2$
   $\Gamma = \Gamma_0[x : \tau_1 + \tau_2]$ $\quad \Gamma' = \Gamma_0[x \hookleftarrow \tau_2]$

   By application of Item 1.

   Case T-LetInt**:** $e = \mathbf{let}\, x = n\, \mathbf{in}\, e'$ $\quad \Theta \mid \lambda \mid \Gamma, x : \{\nu : \mathbf{int} \mid \nu = n\} \vdash e' : \tau \Rightarrow \Gamma'$
   $x \notin dom(\Gamma')$

   The induction hypothesis gives

   $$\Theta \mid \vec{\ell} \mid [\vec{\ell}/\lambda]\Gamma, x : \{\nu : \mathbf{int} \mid \nu = n\} \vdash e : [\vec{\ell}/\lambda]\,\tau \Rightarrow [\vec{\ell}/\lambda]\Gamma'$$

   We conclude $\Theta \mid \vec{\ell} \mid [\vec{\ell}/\lambda]\Gamma \vdash \mathbf{let}\, x = n\, \mathbf{in}\, e : [\vec{\ell}/\lambda]\,\tau \Rightarrow [\vec{\ell}/\lambda]\Gamma'$ as required.

   Case T-Let**:** $e = \mathbf{let}\, x = y\, \mathbf{in}\, e'$ $\quad x \notin dom(\Gamma')$
   $\Theta \mid \lambda \mid \Gamma_1 \vdash e' : \tau \Rightarrow \Gamma'$ $\quad \Gamma_1 = \Gamma[y \hookleftarrow (\tau_1 \wedge_y y =_{\tau_1} x)], x : (\tau_2 \wedge_x x =_{\tau_2} y)$
   $\Gamma \quad (y) = y : \tau_1 + \tau_2$

   By Item 1, $([\vec{\ell}/\lambda]\Gamma)(y) = [\vec{\ell}/\lambda]\,(\tau_1 + \tau_2) = ([\vec{\ell}/\lambda]\,\tau_1 + [\vec{\ell}/\lambda]\,\tau_2)$. We must then show that $\Theta \mid \vec{\ell} \mid \Gamma_1' \vdash e' : [\vec{\ell}/\lambda]\,\tau \Rightarrow [\vec{\ell}/\lambda]\Gamma'$ where

   $$\Gamma_1' = ([\vec{\ell}/\lambda]\Gamma)[y \hookleftarrow [\vec{\ell}/\lambda]\,\tau_1 \wedge_y y = x], x : ([\vec{\ell}/\lambda]\,\tau_2 \wedge_x x = y)$$

   As $\Gamma_1' = [\vec{\ell}/\lambda]\Gamma_1$ the induction hypothesis gives the required typing judgment.

   Case T-If,T-Seq**:**

   By trivial application of the inductive hypothesis.

   Case T-MkRef,T-Deref**:**

   By reasoning similar to T-Let.

Case T-Call**:** $e = \mathbf{let}\, x = f^{\ell}(y_1, \ldots, y_n)\, \mathbf{in}\, e'$
$\sigma_x = [y_1/x_1] \cdots [y_n/x_n]$
$\sigma_\alpha = [\ell : \lambda/\lambda']$
$\Theta \mid \lambda \mid \Gamma_1 \vdash e' : \tau \Rightarrow \Gamma'$
$y \notin dom(\Gamma')$
$\Theta(f) = \forall \lambda'. \langle x_1 : \tau_1, \ldots, x_n : \tau_n \rangle \rightarrow \langle x_1 : \tau_1', \ldots, x_n : \tau_n' \mid \tau' \rangle$
$\Gamma_1 = \Gamma[y_i \hookleftarrow \sigma_\alpha\, \sigma_x\, \tau_i'], x : \sigma_\alpha\, \sigma_x\, \tau'$

We must first show that for $\sigma_\alpha' = [\ell : \vec{\ell}/\lambda']$:

$$\Theta \mid \vec{\ell} \mid \Gamma_3 \vdash e' : [\vec{\ell}/\lambda]\, \tau \Rightarrow [\vec{\ell}/\lambda]\Gamma'$$

where $\Gamma_3 = ([\vec{\ell}/\lambda]\Gamma)[y_i \hookleftarrow \sigma_\alpha'\, \sigma_x\, \tau_i'], x : \sigma_\alpha'\, \sigma_x\, \tau'$.
We first observe that $\Gamma_3 = [\vec{\ell}/\lambda]\Gamma_1$ (this follows from the equivalence of $[\vec{\ell}/\lambda]\,[\ell : \lambda/\lambda']$ and $[\ell : \vec{\ell}/\lambda']$) whereby the induction hypothesis gives the required typing derivation.
We must also show that $\forall i \in \{1..n\}\,.([\ell : \vec{\ell}/\lambda]\Gamma)(y_i) = \sigma_\alpha'\, \sigma_x\, \tau_i$. From the assumed well-typing of the term under λ we have that $\forall i \in \{1..n\}\,.\Gamma(y_i) = \sigma_\alpha\, \sigma_x\, \tau_i$. Recall that σ_α' is equivalent to $[\vec{\ell}/\lambda]\, \sigma_\alpha$, whereby we have $[\vec{\ell}/\lambda]\Gamma(y_i) = [\vec{\ell}/\lambda]\, \sigma_\alpha\, \sigma_x\, \tau_i = \sigma_\alpha'\, \sigma_x\, \tau_i$ for any i as equality is preserved by consistent substitution.

Case T-Assign**:**

By the inductive hypothesis and application of Item 1.

Case T-Alias**:** $\Theta \mid \lambda \mid \Gamma[x : \tau_1\, \mathbf{ref}^{r_1}][y : \tau_2\, \mathbf{ref}^{r_2}] \vdash \mathbf{alias}(x = y)\,; e : \tau \Rightarrow \Gamma$
$\tau_1\, \mathbf{ref}^{r_1} + \tau_2\, \mathbf{ref}^{r_2} \approx \tau_1'\, \mathbf{ref}^{r_1'} + \tau_2'\, \mathbf{ref}^{r_2'}$
$\Theta \mid \lambda \mid \Gamma[x \hookleftarrow \tau_1'\, \mathbf{ref}^{r_1'}][y \hookleftarrow \tau_2'\, \mathbf{ref}^{r_2'}] \vdash e : \tau \Rightarrow \Gamma$

From Item 1 we have that $[\vec{\ell}/\lambda]\,(\tau_1\, \mathbf{ref}^{r_1}) + [\vec{\ell}/\lambda]\,(\tau_2\, \mathbf{ref}^{r_2}) = [\vec{\ell}/\lambda]\,(\tau_1\, \mathbf{ref}^{r_1} + \tau_2\, \mathbf{ref}^{r_2})$ and similarly for $\tau_1'\, \mathbf{ref}^{r_1'} + \tau_2'\, \mathbf{ref}^{r_2'}$. It therefore remains to show that:

$$[\vec{\ell}/\lambda]\,(\tau_1\, \mathbf{ref}^{r_1} + \tau_2\, \mathbf{ref}^{r_2}) \approx [\vec{\ell}/\lambda]\,(\tau_1'\, \mathbf{ref}^{r_1'} + \tau_2'\, \mathbf{ref}^{r_2'})$$

For which it suffices to show that $\bullet \vdash [\vec{\ell}/\lambda]\,(\tau_1\, \mathbf{ref}^{r_1} + \tau_2\, \mathbf{ref}^{r_2}) \leq [\vec{\ell}/\lambda]\,(\tau_1'\, \mathbf{ref}^{r_1'} + \tau_2'\, \mathbf{ref}^{r_2'})$ and $\bullet \vdash [\vec{\ell}/\lambda]\,(\tau_1'\, \mathbf{ref}^{r_1'} + \tau_2'\, \mathbf{ref}^{r_2'}) \leq [\vec{\ell}/\lambda]\,(\tau_1\, \mathbf{ref}^{r_1} + \tau_2\, \mathbf{ref}^{r_2})$. From $\tau_1\, \mathbf{ref}^{r_1} + \tau_2\, \mathbf{ref}^{r_2} \approx \tau_1'\, \mathbf{ref}^{r_1'} + \tau_2'\, \mathbf{ref}^{r_2'}$ these both follow from Item 3, whereby the result follows from the inductive hypothesis.

Case T-AliasPtr**:**

By similar reasoning to the T-Alias case.

Case T-Sub**:** $\Theta \mid \lambda \mid \Gamma_1 \vdash e : \tau_1 \Rightarrow \Gamma_2 \qquad \Gamma \leq \Gamma_1$
$\Gamma_2, \tau_1 \leq \Gamma', \tau$

By the induction hypothesis we have that: $\Theta \mid \vec{\ell} \mid [\vec{\ell}/\lambda]\Gamma_1 \vdash e : [\vec{\ell}/\lambda]\, \tau_1 \Rightarrow [\vec{\ell}/\lambda]\Gamma_2$.
If we show that $[\vec{\ell}/\lambda]\Gamma \leq [\vec{\ell}/\lambda]\Gamma_1$ and $[\vec{\ell}/\lambda]\Gamma_2, [\vec{\ell}/\lambda]\, \tau_1 \leq [\vec{\ell}/\lambda]\Gamma', [\vec{\ell}/\lambda]\, \tau$ we will have the required result. To show the first requirement, for any $x \in dom(\Gamma)$ we have that $[\vec{\ell}/\lambda]\Gamma \vdash [\vec{\ell}/\lambda]\Gamma(x) \leq [\vec{\ell}/\lambda]\Gamma_1(x)$ from Item 3 so

we have $[\vec{\ell}/\lambda]\Gamma \leq [\vec{\ell}/\lambda]\Gamma_1$. To show the latter requirement, we observe that $[\vec{\ell}/\lambda]\Gamma_2, [\vec{\ell}/\lambda]\,\tau_1 \leq [\vec{\ell}/\lambda]\Gamma', [\vec{\ell}/\lambda]\,\tau$ is equivalent to showing $[\vec{\ell}/\lambda](\Gamma_2, x : \tau_1) \leq [\vec{\ell}/\lambda](\Gamma', x : \tau)$ for some $x \notin dom(\Gamma_2)$, whereby we have the required subtyping relationship from the application of Item 3.

Case T-ASSERT**:** $\Theta \mid \lambda \mid \Gamma \vdash \mathbf{assert}(\varphi)\,; e : \tau \Rightarrow \Gamma'$ $\quad \Gamma \models \varphi$
$\Theta \mid \lambda \mid \Gamma \vdash e : \tau \Rightarrow \Gamma'$ $\quad \epsilon \mid \Gamma \vdash_{WF} \varphi$

By induction hypothesis, the result holds if we can show $[\vec{\ell}/\lambda]\Gamma \models \varphi$ which follows from Item 4 (that $\lambda \notin \mathbf{FCV}\,(\varphi)$ follows from the well-formedness of φ with respect to ϵ).

Lemma 10 (Substitution). *If* $\Theta \mid \mathcal{L} \mid \Gamma \vdash e : \tau \Rightarrow \Gamma'$ *and* $x' \notin dom(\Gamma)$*, then* $\Theta \mid \mathcal{L} \mid [x'/x]\,\Gamma \vdash e : [x'/x]\,\tau \Rightarrow [x'/x]\,\Gamma'$.

Proof. By straightforward induction of typing rules.

We now prove that if every variable satisfies its refinement in a type environment Γ, we must have $\models [R]\,[\![\Gamma]\!]$.

Lemma 11. *If* $\mathbf{SAT}\,(H, R, \Gamma)$ *then* $\models [R]\,[\![\Gamma]\!]$.

Proof. To show $\models [R]\,[\![\Gamma]\!]$, it suffices to show that for any $x \in dom(\Gamma)$ where $\Gamma(x) = \{\nu : \mathbf{int} \mid \varphi\} \models [R]\,[x/\,\nu\,]\,\varphi$ holds. From $\mathbf{SAT}\,(H, R, \Gamma)$, we must have $\mathbf{SATv}(H, R, R(x), \Gamma(x))$, whereby we have $R(x) \in \mathbb{Z}$ and $[R]\,[R(x)/\nu]\varphi$. As $[R]\,[x/\,\nu\,]\,\varphi$ is equivalent to $[R]\,[R(x)/\nu]\varphi$, and we have the desired result.

We prove that subtyping preserves the consistency relation in the following sense.

Lemma 12. *If* $\Gamma \leq \Gamma'$ *and* $\mathbf{Cons}\,(H, R, \Gamma)$ *then:*

1. *For any* $x \in dom(\Gamma')$, $\forall\, a \in dom(H).\mathbf{own}\,(H, R(x), \Gamma(x))(a) \leq \mathbf{own}\,(H, R(x), \Gamma'(x))(a)$
2. $\forall\, a \in dom(H).\mathbf{Own}\,(H, R, \Gamma')(a) \leq 1$
3. *If* $\Gamma \vdash \tau \leq \tau'$ *and* $\mathbf{SATv}(H, R, v, \tau)$ *then* $\mathbf{SATv}(H, R, v, \tau')$
4. $\mathbf{SAT}\,(H, R, \Gamma')$
5. $\mathbf{Cons}\,(H, R, \Gamma')$

Proof.

1. By induction on $\Gamma \vdash \Gamma(x) \leq \Gamma'(x)$.
2. Direct consequence of 1 and that $\forall\, a \in dom(H).\mathbf{Own}\,(H, R, \Gamma)(a) \leq 1$ from $\mathbf{Cons}\,(H, R, \Gamma)$.
3. From $\mathbf{Cons}\,(H, R, \Gamma)$ we have $\mathbf{SAT}\,(H, R, \Gamma)$ which by Lemma 11 we have $\models [R]\,[\![\Gamma]\!]$. We now proceed by induction on $\Gamma \vdash \tau \leq \tau'$.

   Case: $\tau = \{\nu : \mathbf{int} \mid \varphi\}$ $\quad \tau' = \{\nu : \mathbf{int} \mid \varphi'\}$
   $\models [\![\Gamma]\!] \wedge \varphi \implies \varphi'$

   From $\mathbf{SATv}(H, R, v, \tau)$ we have $\models [R]\,[v/\nu]\varphi$. We must show that $\models [R]\,[v/\nu]\varphi'$. From $\models [\![\Gamma]\!] \wedge \varphi \implies \varphi'$ we must have $\models [R]\,[v/\nu]\,[\![\Gamma]\!] \wedge [R]\,[v/\nu]\varphi \implies [R]\,[v/\nu]\varphi'$ is valid. As ν does not appear free in $[\![\Gamma]\!]$, we have $\models [R]\,[\![\Gamma]\!] \wedge [R]\,[v/\nu]\varphi \implies [R]\,[v/\nu]\varphi'$ is valid whereby the result is immediate.

Case: $\tau = \tau_p\,\mathbf{ref}^{r_1} \qquad \tau' = \tau'_p\,\mathbf{ref}^{r_2}$
$r_2 \leq r_1$

Immediate from the induction hypothesis.

4. Immediate consequence of 3 and that $\Gamma \leq \Gamma'$ implies that $\Gamma \vdash \Gamma(x) \leq \Gamma'(x)$ for any $x \in dom(\Gamma')$.
5. Immediate from 2 and 4.

To show consistency is preserved during evaluation, Lemmas 13 and 14 show types equivalent according to $\approx$ are equivalent for the purposes of **own** and **SATv**. Then Lemmas 15 and 16 show that the type addition operator + "distributes" over **SATv** and **own**.

Lemma 13 (Type Equivalence Preserves Satisfiability). *If* $\tau_1 \approx \tau_2$*, then* $\mathbf{SATv}(H, R, v, \tau_1) \iff \mathbf{SATv}(H, R, v, \tau_2)$.

Proof. We prove the forward case by induction on $\bullet \vdash \tau_1 \leq \tau_2$ as implied by $\tau_1 \approx \tau_2$. The inductive case follows from the IH. In the the base case where $\tau_1 = \{\nu : \mathbf{int} \mid \varphi_1\}$ and $\tau_2 = \{\nu : \mathbf{int} \mid \varphi_2\}$, from $\bullet \vdash \tau_1 \leq \tau_2$ we have that $\models \varphi_1 \implies \varphi_2$ is valid, from which we must have $\models [R]\,[v/\nu]\varphi_1 \implies [R]\,[v/\nu]\varphi_2$, where from the definition of $\mathbf{SATv}(H, R, v, \tau_1)$ we must then have $\mathbf{SATv}(H, R, v, \tau_2)$.

The backwards case follows similar reasoning by induction on $\bullet \vdash \tau_2 \leq \tau_1$.

Lemma 14. *If* $\tau_1 \approx \tau_2$*, then* $\mathbf{own}\,(H, v, \tau_1) = \mathbf{own}\,(H, v, \tau_2)$.

Proof. By reasoning similar to that in Lemma 13.

Lemma 15. *If* $\tau_p = \tau_1 + \tau_2$*, then* $\mathbf{own}\,(H, v, \tau_p) = \mathbf{own}\,(H, v, \tau_1) + \mathbf{own}\,(H, v, \tau_2)$.

Proof. By induction on the rules used to derive $\tau_1 + \tau_2 = \tau_p$.

Case Tadd-Int**:**

We have $\mathbf{own}\,(H, v, \tau_p) = \mathbf{own}\,(H, v, \tau_1 + \tau_2)$, where $\tau_1 + \tau_2 = \{\nu : \mathbf{int} \mid \varphi_1 \wedge \varphi_2\}$, $\mathbf{own}\,(H, v, \tau_1)$ and $\mathbf{own}\,(H, v, \tau_2)$, where $\tau_1 = \{\nu : \mathbf{int} \mid \varphi_1\}, \tau_2 = \{\nu : \mathbf{int} \mid \varphi_2\}$.

From the definition of ownership, we have $\mathbf{own}\,(H, v, \tau_p) = \mathbf{own}\,(H, v, \tau_1) = \mathbf{own}\,(H, v, \tau_2) = \emptyset$. It is thus trivial that $\mathbf{own}\,(H, v, \tau_p) = \mathbf{own}\,(H, v, \tau_1) + \mathbf{own}\,(H, v, \tau_2)$.

Case Tadd-Ref**:**

We assume $v = a$ and $a \in dom(H)$, otherwise the result trivially holds.

We have $\mathbf{own}\,(H, v, \tau_p) = \mathbf{own}\,(H, v, \tau_1 + \tau_1)$, where $\tau_1 + \tau_2 = (\tau'_1 + \tau'_2)\,\mathbf{ref}^{r_1 + r_2}$, and $\tau_1 = \tau'_1\,\mathbf{ref}^{r_1}$, $\tau_2 = \tau'_2\,\mathbf{ref}^{r_2}$.

From the definition of ownership, we have $\mathbf{own}\,(H, v, \tau_p) = \{a \mapsto r_1 + r_2\} + \mathbf{own}\,(H, H(v), \tau'_1 + \tau'_2)$ and:

$$\begin{aligned}\mathbf{own}\,(H, v, \tau_1) + \mathbf{own}\,(H, v, \tau_2) &= \{a \mapsto r_1\} + \mathbf{own}\,(H, H(v), \tau'_1) + \{a \mapsto r_2\} + \mathbf{own}\,(H, H(v), \tau'_2) \\ &= \{a \mapsto r_1 + r_2\} + \mathbf{own}\,(H, H(v), \tau'_1) + \mathbf{own}\,(H, H(v), \tau'_2)\end{aligned}$$

By the induction hypothesis, have that $\mathbf{own}\,(H, H(v), \tau'_1 + \tau'_2) = \mathbf{own}\,(H, H(v), \tau'_1) + \mathbf{own}\,(H, H(v), \tau'_2)$ and can conclude that $\mathbf{own}\,(H, v, \tau_p) = \mathbf{own}\,(H, v, \tau_1) + \mathbf{own}\,(H, v, \tau_2)$.

Lemma 16. *If* $\tau_p = \tau_1 + \tau_2$*, we have* $\mathbf{SATv}(H, R, v, \tau_p)$ *iff.* $\mathbf{SATv}(H, R, v, \tau_1)$ *and* $\mathbf{SATv}(H, R, v, \tau_2)$

Proof. By induction on the rules used to derive $\tau_1 + \tau_2$. In the following we only prove the forward direction of the implication; the backwards direction is symmetric.

Case TADD-INT**:**

We have $\mathbf{SATv}(H, R, v, \tau_1 + \tau_2)$, where $\tau_1 + \tau_2 = \{\nu : \mathbf{int} \mid \varphi_1 \wedge \varphi_2\}$, we must show $\mathbf{SATv}(H, R, v, \tau_1)$ and $\mathbf{SATv}(H, R, v, \tau_2)$, where $\tau_1 = \{\nu : \mathbf{int} \mid \varphi_1\}, \tau_2 = \{\nu : \mathbf{int} \mid \varphi_2\}$.

From the definition of **SATv**, we must show $[R]\,[v/\nu]\varphi_1$ and $[R]\,[v/\nu]\varphi_2$. From $\mathbf{SATv}(H, R, v, \tau_1 + \tau_2)$ we have $[R]\,[v/x](\varphi_1 \wedge \varphi_2)$. It is immediate that for any value v such that $[R]\,[v/\nu](\varphi_1 \wedge \varphi_2)$, we must have $[R]\,[v/\nu]\varphi_1$ and $[R]\,[v/\nu]\varphi_2$. We then conclude $\mathbf{SATv}(H, R, v, \tau_1 + \tau_2)$ implies $\mathbf{SATv}(H, R, v, \tau_1)$ and $\mathbf{SATv}(H, R, v, \tau_2)$.

Case TADD-REF**:**

Immediate from the definition of **SATv** and the inductive hypothesis.

Definition 2. *The valid substitution relation, written* $R \vdash_{vs} \tau$ *is the smallest relation closed under the following rules:*

$$\frac{\forall\, x \in \mathbf{FPV}\,(\varphi) \setminus \{\nu\}\,.\exists\, n.R(x) = n}{R \vdash_{vs} \{\nu : \mathbf{int} \mid \varphi\}}$$

$$\frac{R \vdash_{vs} \tau}{R \vdash_{vs} \tau\,\mathbf{ref}^r}$$

Lemma 17. *If* $\vec{\ell} \vdash_{WF} \Gamma$ *and* $\mathbf{Cons}\,(H, R, \Gamma)$*, then* $\forall\, x \in dom(\Gamma).R \vdash_{vs} \Gamma(x)$*.*

Proof. By $\mathbf{Cons}\,(H, R, \Gamma)$, all integer variables in Γ must be in the domain of R and must be an integer. From $\vec{\ell} \vdash_{WF} \Gamma$, any free variables in any refinement of any type in Γ must be an integer valued variable in Γ, which gives the required result.

Definition 3. *We will write* $R \sqsubseteq R'$ *to denote two register files such that:*

1. $dom(R) \subseteq dom(R')$*, and*
2. $\forall\, x \in dom(R).R(x) = R'(x)$

Definition 4. *Two heaps* H *and* H' *are* equivalent modulo a*, written* $H \approx_a H'$ *if:*

1. $dom(H) = dom(H')$
2. $\forall\, a' \in dom(H).a' \neq a \implies H(a) = H(a')$
3. *For any* n*,* $H \vdash a \Downarrow n$ *iff* $H' \vdash a \Downarrow n$

Lemma 18. *For any type* $\tau = \top_n$*,* $H, v,$ $\mathbf{own}\,(H, v, \top_n) = \emptyset$*.*

Proof. By induction on τ. In the base case, the result is trivial. Then consider the case where $\tau = \top_{n-1}\,\mathbf{ref}^0$. If $v \notin \mathbf{Addr}$, or if $v = a$ and $a \notin dom(H)$, then the result trivially holds. Otherwise the result holds from the inductive hypothesis, the definition of $+$ and $\{a \mapsto 0\}$.

Lemma 19 (Heap Update Ownership Preservation). *If* $H \approx_a H'$ *and* $\mathbf{own}\,(H, v, \tau)(a) = 0$*, then* $\mathbf{own}\,(H, v, \tau) = \mathbf{own}\,(H', v, \tau)$.

Proof. By induction on the shape of τ. If $\tau = \{\nu : \mathbf{int} \mid \varphi\}$ then the result trivially holds. Otherwise, $\tau = \tau'\,\mathbf{ref}^r$. We assume that $v = a''$ and $a'' \in dom(H)$ (otherwise the result trivially holds, as $dom(H) = dom(H')$ by $H \approx_a H'$). Consider the case where $a'' = a$. By definition $\mathbf{own}\,(H, a, \tau) = \{a \mapsto r\} + \mathbf{own}\,(H, H(a), \tau')$, and by the assumption that $\mathbf{own}\,(H, a, \tau)(a) = 0$ we must have that $r = 0$. Further, by the ownership well-formedness of types, we must have $\tau' = \top_n$ for some n, thus by Lemma 18 we have $\mathbf{own}\,(H, v, \tau) = \emptyset = \{a \mapsto 0\} + \mathbf{own}\,(H', H'(a), \top_n) = \mathbf{own}\,(H', v, \tau)$.

Finally, consider the case where $a'' \neq a$. Then from the definition of $\mathbf{own}\,(H, a'', \tau)$ and our assumption that $\mathbf{own}\,(H, a'', \tau)(a) = 0$, we have $\mathbf{own}\,(H, H(a''), \tau')(a) = 0$, and the result holds from the inductive hypothesis and that $H(a'') = H'(a'')$.

Lemma 20. *For any* H*,* R*,* v*, and* τ*, if* $\mathbf{SATv}(H, R, v, \tau)$ *then* $H \vdash v \Downarrow |\tau|$

Proof. By induction on τ and the definition of $\mathbf{SATv}$.

Lemma 21. *For any* n*, if* $H \vdash v \Downarrow n$ *then for any* R*,* $\mathbf{SATv}(H, R, v, \top_n)$.

Proof. By induction on n. In the base case, by inversion on $H \vdash v \Downarrow 0$ we have $v \in \mathbb{Z}$ and as $[R]\,[v/\nu]\top \implies \top$, we conclude $\mathbf{SATv}(H, R, v, \top_0)$.

For $n > 0$, by inversion on $H \vdash v \Downarrow n$ we have that $v = a$, $a \in dom(H)$, and $H \vdash H(a) \Downarrow n - 1$, whereby the result holds from the inductive hypothesis.

Lemma 22 (Heap Update Consistency Preservation). *If* $H \approx_a H'$ *and* $\mathbf{own}\,(H, v, \tau)(a) = 0$ *and* $\mathbf{SATv}(H, R, v, \tau)$*, then* $\mathbf{SATv}(H', R, v, \tau)$.

Proof. By induction on the shape of τ. The base case where $\tau = \{\nu : \mathbf{int} \mid \varphi\}$ is trivial. We therefore consider the case where $v = a'$ and $\tau = \tau'\,\mathbf{ref}^r$.

We first consider the case where $a' = a$, then by our assumption that $\mathbf{own}\,(H, a, \tau)(a) = 0$, we must have that $\tau = \tau'\,\mathbf{ref}^0$, whereby $\tau = \top_n$ for some n. From $\mathbf{SATv}(H, R, a, \tau)$ and Lemma 20, we must have that $H \vdash a \Downarrow |\tau|$, and from $H \approx_a H'$, we therefore have that $H' \vdash a \Downarrow |\tau|$ whereby the result holds from Lemma 21.

Otherwise, we have that $a' \neq a$, and by definition we must have that $\mathbf{own}\,(H, H(a), \tau')(a) = 0$ and $H'(a) = H(a)$ hence the result follows from the inductive hypothesis.

Lemma 23 (Register Weakening). *If* $\mathbf{SATv}(H, R, v, \tau)$ *and* $R \vdash_{vs} \tau$*, then for any* R' *such that* $R \sqsubseteq R'$*,* $\mathbf{SATv}(H, R', v, \tau)$.

Proof. By induction on the shape of τ. If $\tau = \tau' \, \mathbf{ref}^r$, then the result follows from the inductive hypothesis. We therefore consider the case where $\tau = \{\nu : \mathbf{int} \mid \varphi\}$. Without loss of generality, we consider the case where $dom(R') \setminus dom(R) = \{x\}$, and $R'(x) = n$. (If $R'(x) = a$, the extra binding also has no effect, and the case where more than one binding is added follows from n applications of the following argument.)

From $\mathbf{SATv}(H, R, v, \tau)$, we conclude that $v \in \mathbb{Z}$ and that $[R]\,[v/\nu]\varphi$. If $x \notin \mathbf{FPV}\,(\varphi)$ then $[R]\,[v/\nu]\varphi \iff [R']\,[v/\nu]\varphi$ and the result holds trivially. Otherwise, if $x \in \mathbf{FPV}\,(\varphi)$ and $x \notin dom(R)$ then R is not a valid substitution, violating our assumption.

Lemma 24 (Heap Extension Consistency Preservation). *If we have heap H, such that $\mathbf{SATv}(H, R, v, \tau)$, for any heap $H' = H\{a \mapsto v'\}, a \notin dom(H)$, then we have $\mathbf{SATv}(H', R, v, \tau)$.*

Proof. By induction on the shape of τ. The base case where $\tau = \{\nu : \mathbf{int} \mid \varphi\}$ is trivial. Next, we consider the case where $\tau = \tau' \, \mathbf{ref}^r$. We must show that $v \in dom(H')$ and $\mathbf{SATv}(H', R, H'(v), \tau')$. The first condition is immediately satisfied by inversion on $\mathbf{SATv}(H, R, v, \tau')$, and from $a \notin dom(H)$, we have $v \neq a$, which gives that $H'(v) = H(v)$. That is we must show $\mathbf{SATv}(H', R, H(v), \tau')$, which is follows from the induction hypothesis.

Lemma 25 (Heap Extension Ownership Preservation). *If $\mathbf{SATv}(H, R, v, \tau)$, then for any $a \notin dom(H)$ $\mathbf{own}\,(H, v, \tau) = \mathbf{own}\,(H\{a \mapsto v'\}, v, \tau)$ for any value v'.*

Proof. By induction on τ. The base case is trivial as $\mathbf{own}\,(H, v, \{\nu : \mathbf{int} \mid \varphi\}) = \emptyset = \mathbf{own}\,(H\{a \mapsto v\}, v, \{\nu : \mathbf{int} \mid \varphi\})$. We therefore consider the case where $\tau = \tau' \, \mathbf{ref}^r$.

From $\mathbf{SATv}(H, R, v, \tau)$ we must have that $v = a'$ and $a' \in dom(H)$ (and by extension $a' \in dom(H\{a \mapsto v'\})$). From the definition of the ownership function, we have that $\mathbf{own}\,(H, v, \tau) = \mathbf{own}\,(H, H(a), \tau') + \{a' \mapsto r\}$. and $\mathbf{own}\,(H\{a \mapsto v'\}, v, \tau) = \mathbf{own}\,(H\{a \mapsto v'\}, H\{a \mapsto v'\}(a'), \tau') + \{a' \mapsto r\}$ Then from our requirement that $a \notin dom(H)$, we have $a \neq a'$ and therefore $H(a') = H\{a \mapsto v'\}(a')$, whereby the result holds from the inductive hypothesis.

Lemma 26 (Environment Weakening). *Define the partial operation $\Gamma_1 \uplus \Gamma_2$ for two environments such $dom(\Gamma_1) \cap dom(\Gamma_2) = \emptyset$:*

$$(\Gamma_1 \uplus \Gamma_2) \quad (x) = \begin{cases} \Gamma_1 \quad (x) & x \in dom(\Gamma_1) \\ \Gamma_2 \quad (x) & x \in dom(\Gamma_2) \\ undef & o.w. \end{cases}$$

Then, for any Γ and Γ'' where $dom(\Gamma) \cap dom(\Gamma'') = \emptyset$:

1. *$\Gamma \vdash \tau_1 \leq \tau_2$ implies $\Gamma \uplus \Gamma'' \vdash \tau_1 \leq \tau_2$*
2. *$\Gamma \leq \Gamma'$ implies $\Gamma \uplus \Gamma'' \leq \Gamma' \uplus \Gamma''$*

3. *If* $\Theta \mid \mathcal{L} \mid \Gamma \vdash e : \tau \Rightarrow \Gamma'$, $\mathcal{L} \vdash_{WF} \Gamma \uplus \Gamma''$, *and* $\mathcal{L} \vdash_{WF} \Gamma' \uplus \Gamma''$, *then* $\Theta \mid \mathcal{L} \mid \Gamma \uplus \Gamma'' \vdash e : \tau \Rightarrow \Gamma'' \uplus \Gamma$

Proof.

1. As in the proof of Lemma 9 (part 3), at the root of the subtyping derivation is a logical judgment of the form $\models [\![\Gamma]\!] \wedge \varphi_1 \implies \varphi_2$ which can be shown to be valid. We must then show that $\models [\![\Gamma \uplus \Gamma'']\!] \wedge \varphi_1 \implies \varphi_2$ is valid. As $[\![\Gamma'' \uplus \Gamma]\!] \wedge \varphi_1 = [\![\Gamma'']\!] \wedge [\![\Gamma]\!] \wedge \varphi$ only strengthens the pre-condition $[\![\Gamma]\!] \wedge \varphi_1$, $\models [\![\Gamma'' \uplus \Gamma]\!] \wedge \varphi_1 \implies \varphi_2$ must also be valid.
2. It suffices to show that $\Gamma \uplus \Gamma'' \vdash (\Gamma \uplus \Gamma'')(x) \leq (\Gamma' \uplus \Gamma'')(x)$ for any arbitrary $x \in dom(\Gamma' \uplus \Gamma'')$. If $x \in dom(\Gamma')$ then we must have $\Gamma \vdash \Gamma(x) \leq \Gamma'(x)$ by inversion on $\Gamma \leq \Gamma'$, whereby $\Gamma \uplus \Gamma'' \vdash (\Gamma \uplus \Gamma'')(x) = \Gamma(x) \leq \Gamma' \uplus \Gamma''(x) = \Gamma'(x)$ from part 1.
   If $x \notin dom(\Gamma')$, then we must show $(\Gamma \uplus \Gamma'') \vdash \Gamma''(x) \leq \Gamma''(x)$, which trivially holds.
3. By induction on the typing derivation of $\Theta \mid \mathcal{L} \mid \Gamma \vdash e : \tau \Rightarrow \Gamma'$. We assume that the variables bound in any let expressions that appear within e are not in the domain of Γ''; this requirement can be easily enforced with consistent renaming and is preserved during evaluation. The only interesting cases are T-Sub and T-Assert and the let bindings; the other cases follow from the induction hypothesis.
   We now prove the relevant cases.

   Case T-Let**:** $\Theta \mid \mathcal{L} \mid \Gamma \vdash \mathbf{let}\ x = y\ \mathbf{in}\ e : \tau \Rightarrow \Gamma'$
   $\Theta \mid \mathcal{L} \mid \Gamma[y \hookleftarrow \tau_1 \wedge_y y =_{\tau_1} x], x : \tau_2 \wedge_x x =_{\tau_2} y \vdash e : \tau \Rightarrow \Gamma'$
   $\Gamma(y) = \tau_1 + \tau_2 \qquad x \notin dom(\Gamma')$

   Let $\Gamma''' = \Gamma'' \uplus \Gamma[y \hookleftarrow \tau_1 \wedge_y y =_{\tau_1} x], x : \tau_2 \wedge_x x =_{\tau_2} y$. To use the inductive hypothesis, we must show that $\mathcal{L} \vdash_{WF} \Gamma'''$ and $\mathcal{L} \vdash_{WF} \Gamma' \uplus \Gamma''$. The latter holds by assumption. To show the former, it suffices to show $\mathcal{L} \mid \Gamma''' \vdash_{WF} \tau_1 \wedge_y y =_{\tau_1} x$ and $\mathcal{L} \mid \Gamma''' \vdash_{WF} \tau_2 \wedge_x x =_{\tau_2} y$. From the assumed well-formedness $\mathcal{L} \vdash_{WF} \Gamma \uplus \Gamma''$, we must have $\mathcal{L} \mid \Gamma \uplus \Gamma'' \vdash_{WF} \tau_1 + \tau_2$, and in particular $\mathcal{L} \mid \Gamma \uplus \Gamma'' \vdash_{WF} \tau_1$ and $\mathcal{L} \mid \Gamma \uplus \Gamma'' \vdash_{WF} \tau_2$. From this we conclude both conditions hold. To show the well-typing of the overall let expression, we must show $x \notin dom(\Gamma' \uplus \Gamma'')$, which follows from our assumption and $x \notin dom(\Gamma')$. Finally, we must also show that $\mathcal{L} \mid \Gamma' \uplus \Gamma'' \vdash_{WF} \tau$. From $\mathcal{L} \mid \Gamma' \vdash_{WF} \tau$ and the fact that$\forall x \in dom(\Gamma').\Gamma'(x) = \{\nu : \mathbf{int} \mid _\}$ iff $(\Gamma' \uplus \Gamma'')(x) = \{\nu : \mathbf{int} \mid \_\}$, we must have $\mathcal{L} \mid \Gamma' \uplus \Gamma'' \vdash_{WF} \tau$.

   Cases T-LetInt**,** T-Mkref**,** T-Mkref**,** T-Deref**,** T-Call**:**

   Similar to the reasoning in T-Let.

   Case T-Sub**:** $\Gamma \leq \Gamma_1 \qquad \Theta \mid \mathcal{L} \mid \Gamma_1 \vdash e : \tau_2 \Rightarrow \Gamma_2$
   $\Gamma_2, \tau_2 \leq \Gamma', \tau$

   From the rules for subtyping, we must have $dom(\Gamma_1) \subseteq dom(\Gamma)$ and $dom(\Gamma') \subseteq dom(\Gamma_2)$. A simple inductive argument gives that $dom(\Gamma_2) \subseteq dom(\Gamma_1)$, therefore we have $dom(\Gamma') \subseteq dom(\Gamma_1)$. Let $\mathcal{LV}$ be the set of free variables in the

refinements of Γ'' that are not in the domain of Γ''. From the assumed well-formedness of $\mathcal{L} \vdash_{WF} \Gamma' \uplus \Gamma''$, we must have that $\forall x \in \mathcal{LV}.x \in dom(\Gamma') \wedge \Gamma'(x) = \{\nu : \mathbf{int} \mid \_\}$. Thus, $\mathcal{LV} \subseteq \Gamma_1$ and $\mathcal{LV} \subseteq \Gamma_2$. Further, by definition, for any $\Gamma_p \leq \Gamma_q$, if $\Gamma_q(x) = \{\nu : \mathbf{int} \mid \_\}$ then $\Gamma_p(x) = \{\nu : \mathbf{int} \mid \_\}$, i.e. subtyping preserves simple types. We conclude that $\mathcal{L} \vdash_{WF} \Gamma_1 \uplus \Gamma''$ and $\mathcal{L} \vdash_{WF} \Gamma_2 \uplus \Gamma''$, whereby the inductive hypothesis gives $\Theta \mid \mathcal{L} \mid \Gamma_1 \uplus \Gamma'' \vdash e : \tau_2 \Rightarrow \Gamma_2 \uplus \Gamma''$. To prove the overall result, we must show that $\Gamma \uplus \Gamma'' \leq \Gamma_1 \uplus \Gamma''$ and $\Gamma_2 \uplus \Gamma'', \tau_2 \leq \Gamma' \uplus \Gamma'', \tau$ which follow from parts 1 and 2 above. That $\mathcal{L} \mid \Gamma_2 \uplus \Gamma'' \vdash_{WF} \tau_2$ follows by reasoning to the case for T-Let above.

Case T-Assert**:**

We must show that $\models [\![\Gamma'' \uplus \Gamma]\!] \implies \varphi$ which is equivalent to $\models [\![\Gamma'']\!] \wedge [\![\Gamma]\!] \implies \varphi$. As the source term was well typed, $\models [\![\Gamma]\!] \implies \varphi$ is valid, we must then have $\models [\![\Gamma'']\!] \wedge [\![\Gamma]\!] \implies \varphi$ whereby the inductive hypothesis gives the required result.

C Proof of Lemma 3

We first prove two additional lemmas. Lemmas 27 and 28 give key facts used in the return and call cases respectively; we have separated them into separate lemmas for clarity.

Lemma 27. *For any Γ_0 such that $\Theta \mid \ell : \vec{\ell} \mid \Gamma_0 \vdash x : \tau_1 \Rightarrow \Gamma_1$ and $\Theta \mid [] : \tau_1 \Rightarrow \Gamma_1 \mid \vec{\ell} \vdash_{ectx} E[\mathbf{let}\, y = []^{\ell}\, \mathbf{in}\, e] : \tau_2 \Rightarrow \Gamma_2$ then $\Theta \mid \vec{\ell} \mid \Gamma_0 \vdash E[\mathbf{let}\, y = []^{\ell}\, \mathbf{in}\, e][x] : \tau_2 \Rightarrow \Gamma_2$.*

Proof. It suffices to show that $\Theta \mid \vec{\ell} \mid \Gamma_0 \vdash \mathbf{let}\, y = x\, \mathbf{in}\, e : \tau_1' \Rightarrow \Gamma_1'$ and $\Theta \mid [] : \tau_1' \Rightarrow \Gamma_1' \mid \vec{\ell} \vdash_{ectx} E : \tau_2 \Rightarrow \Gamma_2$ for some τ_1' and Γ_1' whereby the result will hold from Lemma 8.

By inversion on $\Theta \mid [] : \tau_1 \Rightarrow \Gamma_1 \mid \vec{\ell} \vdash_{ectx} E[\mathbf{let}\, y = []^{\ell}\, \mathbf{in}\, e] : \tau_2 \Rightarrow \Gamma_2$ we have

$$\Theta \mid \vec{\ell} \mid \Gamma_1, y : \tau_1 \vdash e : \tau_1'' \Rightarrow \Gamma_1'' \tag{1}$$
$$\Theta \mid [] : \tau_1'' \Rightarrow \Gamma_1'' \mid \mathcal{L} \vdash_{ectx} E : \tau_2 \Rightarrow \Gamma_2 \tag{2}$$
$$y \notin dom(\Gamma_1'') \tag{3}$$

We take $\Gamma_1' = \Gamma_1''$, $\tau_1' = \tau_1''$, and then Equation (2) gives the necessary typing for E.

It remains to to show that

$$\Theta \mid \vec{\ell} \mid \Gamma_0 \vdash \mathbf{let}\, y = x\, \mathbf{in}\, e : \tau_1' \Rightarrow \Gamma_1'$$

(That $\vec{\ell} \vdash_{WF} \tau_1' \Rightarrow \Gamma_1'$ follows from Equation (1))

By Lemma 6, from $\Theta \mid \ell : \vec{\ell} \mid \Gamma_0 \vdash x : \tau_1 \Rightarrow \Gamma_1$ we conclude there exists some Γ_p, τ_p, and Γ'_p such:

$$\Gamma_0 \leq \Gamma_p \tag{4}$$
$$\Gamma'_p, \tau_p \leq \Gamma_1, \tau_1 \tag{5}$$
$$\Gamma'_p = \Gamma_p[x \hookleftarrow \tau'_p] \tag{6}$$
$$\Gamma_p(x) = \tau_p + \tau'_p \tag{7}$$
$$\vec{\ell} \vdash_{WF} \Gamma_p \tag{8}$$

We first apply T-Sub with Equations (4) and (8), so it remains to show

$$\Theta \mid \vec{\ell} \mid \Gamma_p[x : \tau_p + \tau'_p] \vdash \mathbf{let}\, y = x \,\mathbf{in}\, e : \tau'_1 \Rightarrow \Gamma'_1$$

which, by T-Let holds if we show that:

$$\Theta \mid \vec{\ell} \mid \Gamma_p[x \hookleftarrow \tau'_p \wedge_x x =_{\tau'_p} y], y : \tau_p \wedge_y y =_{\tau_p} x \vdash e : \tau'_1 \Rightarrow \Gamma'_1$$

($y \notin dom(\Gamma'_1)$ follows from Equation (3), and the well-formedness of $\Gamma_p[x \hookleftarrow \tau'_p \wedge_x x =_{\tau'_p} y], y : \tau_p \wedge_y y =_{\tau_p} x$ follows from the well-formedness of Γ_p, τ_p and τ'_p and that x and y appear in the refinements iff they are mapped to integer types in the new type environment.)

We can use T-Sub to weaken the type environment to:

$$\Theta \mid \vec{\ell} \mid \Gamma_p[x \hookleftarrow \tau'_p], y : \tau_p \vdash e : \tau'_1 \Rightarrow \Gamma'_1$$

From Equation (5) above, we have that $\Gamma_p[x \hookleftarrow \tau'_p], y : \tau_p \leq \Gamma_1, y : \tau_1$, whereby one final application of T-Sub allows us to use to Equation (1) above.

Lemma 28. *Let* $E[\mathbf{let}\, x = f^{\ell}(y_1, \ldots, y_n)\, \mathbf{in}\, e']$ *be a term such that:*

$\Theta \mid \vec{\ell} \mid \Gamma_0 \vdash \mathbf{let}\, x = f^{\ell}(y_1, \ldots, y_n)\, \mathbf{in}\, e' : \tau_1 \Rightarrow \Gamma_1$ $\quad$ $\sigma_\alpha = [\ell : \vec{\ell}/\lambda]$

$\Theta \mid [] : \tau_1 \Rightarrow \Gamma_1 \mid \vec{\ell} \vdash_{ectx} E : \tau_2 \Rightarrow \Gamma_2$ $\quad$ $\sigma_x = [y_1/x_1] \cdots [y_n/x_n]$

$f \mapsto (x_1, .., x_n) e \in D$ $\quad$ $\Theta \vdash f \mapsto (x_1, .., x_n) e$

$\vdash_{WF} \Theta$

where $\Theta(f) = \forall \lambda.\, \langle x_1 : \tau_1, \ldots, x_n : \tau_n \rangle \rightarrow \langle x_1 : \tau'_1, \ldots, x_n : \tau'_n \mid \tau_q \rangle$.

Then there exist some τ_3 *and* Γ_3*:*

$$\Theta \mid \ell : \vec{\ell} \mid \Gamma_0 \vdash \sigma_x\, e : \tau_3 \Rightarrow \Gamma_3$$
$$\Theta \mid [] : \tau_3 \Rightarrow \Gamma_3 \mid \vec{\ell} \vdash_{ectx} E[\mathbf{let}\, x = []^{\ell}\, \mathbf{in}\, e] : \tau_2 \Rightarrow \Gamma_2$$

Proof. From Lemma 6 on $\Theta \mid \vec{\ell} \mid \Gamma_0 \vdash \mathbf{let}\, x = f^{\ell}(y_1, \ldots, y_n)\, \mathbf{in}\, e' : \tau_1 \Rightarrow \Gamma_1$ we have, for some $\Gamma_p, \tau_p, \Gamma'_p$, that:

$$\Gamma_0 \leq \Gamma_p \tag{9}$$
$$\Gamma'_p, \tau_p \leq \Gamma_1, \tau_1 \tag{10}$$
$$\Gamma_p(y_i) = \sigma_\alpha\, \sigma_x\, \tau_i \tag{11}$$
$$\vec{\ell} \vdash_{WF} \Gamma_p \tag{12}$$
$$\Theta \mid \vec{\ell} \mid \Gamma_p[y_i \hookleftarrow \sigma_\alpha\, \sigma_x\, \tau'_i], x : \sigma_\alpha\, \sigma_x\, \tau_q \vdash e' : \tau_p \Rightarrow \Gamma'_p \tag{13}$$
$$x \notin dom(\Gamma'_p) \tag{14}$$

To prove the first part, from the well-typing of the function body, we have $\Theta \mid \lambda \mid x_1 : \tau_1, \ldots, x_n : \tau_n \vdash e : \tau_q \Rightarrow x_1 : \tau'_1, \ldots, x_n : \tau'_n$. From our assumption that all variable names are distinct, by n applications of the substitution lemma (Lemma 10) we have: $\Theta \mid \lambda \mid y_1 : \sigma_x\, \tau_1, \ldots, y_n : \sigma_x\, \tau_n \vdash \sigma_x e : \sigma_x\, \tau_q \Rightarrow y_1 : \sigma_x\, \tau'_1, \ldots, y_n : \sigma_x\, \tau'_n$. By Lemma 9 (part 5) we then have $\Theta \mid \ell : \vec{\ell} \mid y_1 : \sigma_\alpha\, \sigma_x\, \tau_1, \ldots, y_n : \sigma_\alpha\, \sigma_x\, \tau_n \vdash \sigma_x e : \sigma_\alpha\, \sigma_x\, \tau_q \Rightarrow y_1 : \sigma_\alpha\, \sigma_x\, \tau'_1, \ldots, y_n : \sigma_\alpha\, \sigma_x\, \tau'_n$. We take $\tau_3 = \sigma_\alpha\, \sigma_x\, \tau_q$ and $\Gamma_3 = \Gamma_p[y_i \hookleftarrow \sigma_\alpha\, \sigma_x\, \tau'_i]$.

By the well-formedness of function types and well-formedness of Γ_p, we must have that $\ell : \vec{\ell} \vdash_{WF} \Gamma_3$. Then by Equations (11) and (12) and lemma 26 we have $\Theta \mid \ell : \vec{\ell} \mid \Gamma_p \vdash \sigma_x e : \tau_3 \Rightarrow \Gamma_3$, whereby Equation (9) and an application of T-Sub gives $\Theta \mid \ell : \vec{\ell} \mid \Gamma_0 \vdash \sigma_x e : \tau_3 \Rightarrow \Gamma_3$, i.e., the first result.

To prove the second part, from the typing rule for TE-Stack we must show:

$$\Theta \mid [] : \tau_1 \Rightarrow \Gamma_1 \mid \vec{\ell} \vdash_{ectx} E : \tau_2 \Rightarrow \Gamma_2 \tag{15}$$

$$x \notin dom(\Gamma_1) \tag{16}$$

$$\Theta \mid \vec{\ell} \mid \Gamma_3, x : \tau_3 \vdash e' : \tau_1 \Rightarrow \Gamma_1 \tag{17}$$

Equation (15) holds by assumption, and Equation (16) follows from Equation (14) and that $\Gamma'_p \leq \Gamma_1$ implies $dom(\Gamma_1) \subseteq dom(\Gamma'_p)$. $\vec{\ell} \vdash_{WF} \tau_1 \Rightarrow \Gamma_1$ follows from the well-typing of the function call term, and $\vec{\ell} \vdash_{WF} \Gamma_3, x : \tau_3$ follows from Equation (13).

From Equations (10) and (13) we then have Equation (17) via an application of T-Sub.

Proof (Preservation; Lemma 3). The proof is organized by cases analysis on the transition rule used of e, and showing that the output configuration is well typed by $\vdash^D_{conf}$. We must therefore find a Γ'' that is consistent with H' and R' and also satisfies the other conditions imposed by the definition of $\vdash^D_{conf}$. Here Γ'', H', R' represent the type environment, heap and register after the transition respectively. We identify the heap and register file before transition with H and R respectively. In order to show that the ownership invariant is preserved, we need to prove that $\forall\, a \in dom(H).\mathbf{Own}\,(H, R, \Gamma'')(a) \leq 1$. In many cases, we will show that $\mathbf{Own}\,(H, R, \Gamma) = \mathbf{Own}\,(H', R', \Gamma'')$, whereby from the assumption that $\mathbf{Cons}\,(H, R, \Gamma)$ as implied by $\vdash^D_{conf}$ we have $\forall\, a \in dom(H).\mathbf{Own}\,(H, R, \Gamma)(a) \leq 1$, giving the desired result.

Case R-Var**:** $\quad \vdash^D_{conf} \left\langle H, R, F_{n-1} : \vec{F}, x \right\rangle, \left\langle H, R, F_{n-1} : \vec{F}, x \right\rangle \longrightarrow_D \left\langle H, R, \vec{F}, F_{n-1}[x] \right\rangle$

By inversion on configuration typing $\vdash^D_{conf} \left\langle H, R, F_{n-1} : \vec{F}, x \right\rangle$, we have:

$$\Theta \mid \vec{\ell} \mid \Gamma \vdash x : \tau_n \Rightarrow \Gamma_n$$

$$\forall i \in \{1..n\}\, .\Theta \mid [] : \tau_i \Rightarrow \Gamma_i \mid \vec{\ell}_{i-1} \vdash_{ectx} F_{i-1} : \tau_{i-1} \Rightarrow \Gamma_{i-1}$$

Using Lemma 27, we can conclude that $\Theta \mid \vec{\ell}_{n-1} \mid \Gamma \vdash F_{n-1}[x] : \tau_{n-1} \Rightarrow \Gamma_{n-1}$. We therefore take $\Gamma'' = \Gamma$.

It remains to show that $\mathbf{Cons}\,(H, R, \Gamma'')$ which follows immediately from $\mathbf{Cons}\,(H, R, \Gamma)$.

Case R-Deref: $\vdash^{D}_{conf} \left\langle H, R, \vec{F}, E[\mathbf{let}\, x = *y\, \mathbf{in}\, e]\right\rangle$
$\left\langle H, R, \vec{F}, E[\mathbf{let}\, x = *y\, \mathbf{in}\, e]\right\rangle \longrightarrow_D \left\langle H, R\{x' \mapsto v\}, \vec{F}, E[[x'/x]e]\right\rangle$
$H(a) = v \qquad R(y) = a \qquad R' = R\{x' \mapsto v\}$

By inversion on the configuration typing relationship, we have that:

$$\Theta \mid \vec{\ell} \mid \Gamma_0 \vdash E[\mathbf{let}\, x = *y\, \mathbf{in}\, e] : \tau_n \Rightarrow \Gamma_n \qquad\qquad \mathbf{Cons}\,(H, R, \Gamma_0)$$

By Lemma 7, we have some τ, Γ'_0 such that:

$$\Theta \mid [] : \tau \Rightarrow \Gamma'_0 \mid \vec{\ell} \vdash_{ectx} E : \tau_n \Rightarrow \Gamma_n \qquad \Theta \mid \vec{\ell} \mid \Gamma_0 \vdash \mathbf{let}\, x = *y\, \mathbf{in}\, e : \tau \Rightarrow \Gamma'_0$$

Using Lemma 6, we have some Γ_p, Γ'_p and τ_p such that:

$$\Gamma_0 \leq \Gamma_p \qquad \vec{\ell} \vdash_{WF} \Gamma_p \qquad \Gamma'_p, \tau_p \leq \Gamma'_0, \tau$$
$$\Gamma_p(y) = (\tau_1 + \tau_2)\,\mathbf{ref}^r \qquad x \notin dom(\Gamma'_p)$$
$$\tau'' = \begin{cases} (\tau_1 \wedge_y y =_{\tau_1} x) & r > 0 \\ \tau_1 & r = 0 \end{cases}$$
$$\Theta \mid \vec{\ell} \mid \Gamma_p[y \hookleftarrow \tau''\,\mathbf{ref}^r], x : \tau_2 \vdash e : \tau_p \Rightarrow \Gamma'_p$$

From Lemma 12, we then have $\mathbf{Cons}\,(H, R, \Gamma_p)$. We will now show that:

$$\mathbf{Cons}\,(H, R\{x' \mapsto v\}, \Gamma'') \tag{18}$$
$$\Theta \mid \vec{\ell} \mid \Gamma'' \vdash [x'/x]e : [x'/x]\,\tau_p \Rightarrow [x'/x]\,\Gamma'_p \tag{19}$$

where $\Gamma'' = \Gamma_p[y \hookleftarrow ([x'/x]\,\tau'')\,\mathbf{ref}^r], x' : \tau_2$.

Together these give our desired result. To see how, from $\Theta \mid \vec{\ell} \mid \Gamma_p[y \hookleftarrow \tau''\,\mathbf{ref}^r], x : \tau_2 \vdash e : \tau_p \Rightarrow \Gamma'_p$ above, we must have that $\vec{\ell} \vdash_{WF} \tau_p \Rightarrow \Gamma'_p$. From $x \notin dom(\Gamma'_p)$ we must therefore have that $[x'/x]\,\tau_p = \tau_p$ and $[x'/x]\,\Gamma'_p = \Gamma'_p$. As $\Gamma'_p, \tau_p \leq \Gamma'_0, \tau$ an application of T-Sub gives $\Theta \mid \vec{\ell} \mid \Gamma'' \vdash [x'/x]e : \tau \Rightarrow \Gamma'_0$. Then Lemma 8 will give that $\Theta \mid \vec{\ell} \mid \Gamma'' \vdash E[[x'/x]e] : \tau_n \Rightarrow \Gamma_n$.

As E and the stack $\vec{F}$ remained unchanged, combined with Equation (18) this gives $\vdash^{D}_{conf} \left\langle H, R\{x' \mapsto v\}, \vec{F}, E[[x'/x]e]\right\rangle$ as required. As the above argument is used almost completely unchanged in all of the following cases, we will invert the redex without regard for the T-Sub rule, with the understanding that the subtyping rule is handled with an argument identical to the above.

We now show that $\Theta \mid \vec{\ell} \mid \Gamma'' \vdash [x'/x]e : [x/x']\,\tau_p \Rightarrow [x'/x]\,\Gamma'$ and $\mathbf{Cons}\,(H, R\{x' \mapsto v\}, \Gamma'')$. The first is easy to obtain using Lemma 10; from $\mathbf{Cons}\,(H, R, \Gamma)$, we have that $\forall x \in dom(\Gamma).x \in dom(R)$, whereby from $x \notin dom(R)$ we have $x' \notin dom(\Gamma)$. It therefore remains to show $\mathbf{Cons}\,(H, R\{x' \mapsto v\}, \Gamma'')$.

To show $\mathbf{SAT}\,(H, R, \Gamma'')$, it suffices to show that $\mathbf{SATv}(H, R', R'(x'), \tau_2)$ and $\mathbf{SATv}(H, R', H(R'(y)), \tau'')$ (that $\mathbf{SATv}$ holds for all other variables z follows from $\Gamma(z) = \Gamma''(z)$ and Lemmas 17 and 23). If τ_1 is an integer type and $r > 0$, then by the definition of the strengthening operator, the latter is equivalent to show that $\mathbf{SATv}(H, R', H(R'(y)), \tau_1)$ and that $R'(x') = H(R'(y)) = H(R(y))$, which is immediate from the definition of R-DEREF. If τ_1 is not an integer or if $r = 0$, then we must only show that $\mathbf{SATv}(H, R', H(R'(y)), \tau_1)$.

From $\mathbf{Cons}\,(H, R, \Gamma_p)$, we know that $\mathbf{SAT}\,(H, R, \Gamma_p)$, in particular, $\mathbf{SATv}(H, R, R(y), \Gamma_p(y))$. Then by Lemmas 16, 17 and 23, from $R \sqsubseteq R'$, $\vec{\ell} \vdash_{WF} \Gamma_p$, and $\mathbf{SATv}(H, R, v, \tau_1 + \tau_2)$ we obtain that $\mathbf{SATv}(H, R', v, \tau_1)$ and $\mathbf{SATv}(H, R', v, \tau_2)$, where $v = H(R(y))$. We thus have that $\mathbf{SATv}(H, R', R'(x'), \tau_2)$ and $\mathbf{SATv}(H, R', H(R'(y)), \tau_1)$ are satisfied.

We must also show that the ownership invariant is preserved. Then, it's to show $\forall\, a \in dom(H).\mathbf{Own}\,(H, R', \Gamma'')(a) \leq 1$. Define O'_0 and O'_1 as follows:

$$\begin{aligned}
\mathbf{Own}\,(H, R, \Gamma_p) &= O'_0 + \mathbf{own}\,(H, R(y), \Gamma_p(y)) \\
\mathbf{Own}\,(H, R', \Gamma'') &= O'_1 + \mathbf{own}\,(H, R'(y), \Gamma''(y)) + \mathbf{own}\,(H, R'(x'), \Gamma''(x')) \\
O'_0 &= \Sigma_{z \in dom(\Gamma) \setminus \{y\}}\, \mathbf{own}\,(H, R(z), \Gamma_p(z)) \\
O'_1 &= \Sigma_{z \in dom(\Gamma'') \setminus \{y, x'\}}\, \mathbf{own}\,(H, R'(z'), \Gamma''(z'))
\end{aligned}$$

By Lemma 19, $O'_0 = O'_1$ holds. Then, it suffices to show that $\mathbf{own}\,(H, R(y), \Gamma_p(y)) = \mathbf{own}\,(H, R'(y), \Gamma''(y)) + \mathbf{own}\,(H, R'(x'), \Gamma''(x'))$.

As $R'(x') = H(R'(y)) = H(R(y))$ and from the definition of Γ'', we have:

$$\begin{aligned}
\mathbf{own}\,(H, R'(x'), \Gamma''(x')) &= \mathbf{own}\,(H, H(R(y)), \tau_2) \\
\mathbf{own}\,(H, R'(y), \Gamma''(y)) &= \{a \mapsto r\} + \mathbf{own}\,(H, H(R(y)), \tau_1)
\end{aligned}$$

From the definition of the ownership function, we have that

$$\mathbf{own}\,(H, R(y), \Gamma_p(y)) = \{a \mapsto r\} + \mathbf{own}\,(H, H(R(y)), \tau_1 + \tau_2)$$

which, by Lemma 15, is equivalent to:

$$\{a \mapsto r\} + \mathbf{own}\,(H, H(R(y)), \tau_1) + \mathbf{own}\,(H, H(R(y)), \tau_2)$$

We therefore have $\mathbf{own}\,(H, R(y), \Gamma(y)) = \mathbf{own}\,(H, R'(y), \Gamma''(y)) + \mathbf{own}\,(H, R'(x'), \Gamma''(x'))$, and conclude that $\mathbf{Own}\,(H, R, \Gamma) = \mathbf{Own}\,(H, R', \Gamma'')$.

Case R-SEQ**:** $\quad \vdash^{D}_{conf} \left\langle H, R, \vec{F}, E[x\,;e] \right\rangle$
$\left\langle H, R, \vec{F}, E[x\,;e] \right\rangle \longrightarrow_D \left\langle H, R, \vec{F}, E[e] \right\rangle$

By inversion (see R-DEREF) we have for some Γ that:

$$\begin{aligned}
&\Theta \mid \vec{\ell} \mid \Gamma[x : \tau_0 + \tau_1] \vdash x : \tau_0 \Rightarrow \Gamma[x \hookleftarrow \tau_1] \\
&\Theta \mid \vec{\ell} \mid \Gamma[x \hookleftarrow \tau_1] \vdash e : \tau' \Rightarrow \Gamma' \\
&\mathbf{Cons}\,(H, R, \Gamma)
\end{aligned}$$

We take $\Gamma'' = \Gamma[x \hookleftarrow \tau_1]$.

It suffices to show (see R-Deref) that $\Theta \mid \vec{\ell} \mid \Gamma'' \vdash e : \tau' \Rightarrow \Gamma'$, and $\mathbf{Cons}\,(H, R, \Gamma'')$. The first is immediate from the inversion above, and $\mathbf{Cons}\,(H, R, \Gamma'')$ follows from Lemmas 15 and 16.

Case R-Let: $\vdash^D_{conf} \left\langle H, R, \vec{F}, E[\mathbf{let}\, x = y\, \mathbf{in}\, e]\right\rangle$
$\left\langle H, R, \vec{F}, E[\mathbf{let}\, x = y\, \mathbf{in}\, e]\right\rangle \longrightarrow_D \left\langle H, R\{x' \mapsto R(y)\}, \vec{F}, E[[x'/x]e]\right\rangle$
$x' \notin dom(R) \qquad R' = R\{x' \mapsto R(y)\}$

By inversion (see R-Deref) we have that for some Γ that:

$$\begin{array}{l}\vec{\ell} \vdash_{WF} \Gamma \\ \Gamma(y) = \tau_1 + \tau_2 \\ \Theta \mid \vec{\ell} \mid \Gamma[y \hookleftarrow \tau_1 \wedge_y y =_{\tau_1} x], x : (\tau_2 \wedge_x x =_{\tau_2} y) \vdash e : \tau \Rightarrow \Gamma' \\ \mathbf{Cons}\,(H, R, \Gamma) \qquad x \notin dom(\Gamma')\end{array}$$

We give $\Gamma'' = \Gamma[y \hookleftarrow \tau_1 \wedge_y y =_{\tau_1} x'], x' : (\tau_2 \wedge_{x'} x' =_{\tau_2} y)$.

It suffices to show (see R-Deref) that $\Theta \mid \vec{\ell} \mid \Gamma'' \vdash [x'/x]e : \tau \Rightarrow \Gamma'$ and $\mathbf{Cons}\,(H, R\{x' \mapsto R(y)\}, \Gamma'')$. The first is easy to obtain using reasoning as in the R-Deref case. It therefore remains to show $\mathbf{Cons}\,(H, R', \Gamma'')$.

To show that the output environment is consistent, we must show that $\mathbf{SATv}(H, R', R'(x'), \tau_2 \wedge_{x'} x' = y)$ and $\mathbf{SATv}(H, R', R'(y), \tau_1 \wedge_y y = x')$. By reasoning similar to that in R-Deref, it suffices to show that $\mathbf{SATv}(H, R', R'(x'), \tau_2)$ and $\mathbf{SATv}(H, R', R'(y), \tau_1)$. We know that $\mathbf{Cons}\,(H, R, \Gamma)$, from which me have $\mathbf{SAT}\,(H, R, \Gamma)$, in particular $y \in dom(R)$ and $\mathbf{SATv}(H, R, R(y), \Gamma(y))$. As $R \sqsubseteq R'$ and $\vec{\ell} \vdash_{WF} \Gamma$, from Lemmas 16, 17 and 23, we obtain from $\mathbf{SATv}(H, R, v, \tau_1 + \tau_2)$ that $\mathbf{SATv}(H, R', v, \tau_1)$ and $\mathbf{SATv}(H, R', v, \tau_2)$ where $v = R(y)$. We then have $\mathbf{SATv}(H, R', R'(x'), \tau_2)$ and $\mathbf{SATv}(H, R', R'(y), \tau_1)$ are satisfied.

We must also show that the ownership invariant is preserved. Then, it's to show $\forall\, a \in dom(H).\mathbf{Own}\,(H, R', \Gamma'')(a) \leq 1$. Define O'_0 and O'_1 as follows:

$$\begin{array}{rl}\mathbf{Own}\,(H, R, \Gamma) &= O'_0 + \mathbf{own}\,(H, R(y), \Gamma(y)) \\ \mathbf{Own}\,(H, R', \Gamma'') &= O'_1 + \mathbf{own}\,(H, R'(y), \Gamma''(y)) + \mathbf{own}\,(H, R'(x'), \Gamma''(x')) \\ O'_0 &= \Sigma_{z \in dom(\Gamma) \setminus \{y\}}\, \mathbf{own}\,(H, R(z), \Gamma(z)) \\ O'_1 &= \Sigma_{z \in dom(\Gamma'') \setminus \{y, x'\}}\, \mathbf{own}\,(H, R'(z'), \Gamma''(z'))\end{array}$$

By Lemma 19, $O'_0 = O'_1$ holds. That $\mathbf{own}\,(H, R'(x'), \tau_2) + \mathbf{own}\,(H, R'(y), \tau_1) = \mathbf{own}\,(H, R(y), \tau_1 + \tau_2)$ follows immediately from Lemma 15 and the condition $R(y) = R'(x') = R'(y)$. We therefore conclude that $\mathbf{Own}\,(H, R, \Gamma) = \mathbf{Own}\,(H, R', \Gamma'')$.

Case R-LetInt: $\vdash^D_{conf} \left\langle H, R, \vec{F}, E[\mathbf{let}\, x = n\, \mathbf{in}\, e]\right\rangle$
$\left\langle H, R, \vec{F}, E[\mathbf{let}\, x = n\, \mathbf{in}\, e]\right\rangle \longrightarrow_D \left\langle H, R\{x' \mapsto n\}, \vec{F}, E[[x'/x]e]\right\rangle$

By inversion (see R-DEREF) we have that, for some Γ:

$$\Theta \mid \vec{\ell} \mid \Gamma, x : \{\nu : \mathbf{int} \mid \nu = n\} \vdash e : \tau \Rightarrow \Gamma'$$
$$\mathbf{Cons}\,(H, R, \Gamma) \qquad x \notin dom(\Gamma')$$

We give that $\Gamma'' = \Gamma, x' : \{\nu : \mathbf{int} \mid \nu = n\}$, and it thus suffices to show that $\Theta \mid \vec{\ell} \mid \Gamma'' \vdash [x'/x]e : \tau \Rightarrow \Gamma'$ and $\mathbf{Cons}\,(H, R\{x' \mapsto n\}, \Gamma'')$. The first one is easy to obtain using the Lemma 10 (see R-DEREF) and the latter is trivial by similar reasoning to the T-LET and T-DEREF cases.

Case R-IFTRUE: $\vdash^{D}_{conf} \left\langle H, R, \vec{F}, E[\mathbf{ifz}\ y\ \mathbf{then}\ e_1\ \mathbf{else}\ e_2] \right\rangle$
$\left\langle H, R, \vec{F}, E[\mathbf{ifz}\ y\ \mathbf{then}\ e_1\ \mathbf{else}\ e_2] \right\rangle \longrightarrow_D \left\langle H, R, \vec{F}, E[e_1] \right\rangle$

By inversion (see R-DEREF) we have that for some Γ:

$$\Gamma(x) = \{\nu : \mathbf{int} \mid \varphi\}$$
$$\Theta \mid \vec{\ell} \mid \Gamma[x \hookleftarrow \{\nu : \mathbf{int} \mid \varphi \wedge \nu = 0\}] \vdash e_1 : \tau \Rightarrow \Gamma'$$
$$\mathbf{Cons}\,(H, R, \Gamma)$$

We take $\Gamma'' = \Gamma[x \hookleftarrow \{\nu : \mathbf{int} \mid \varphi \wedge \nu = 0\}]$, and want to show that $\mathbf{Cons}\,(H, R, \Gamma'')$ (that $\Theta \mid \vec{\ell} \mid \Gamma'' \vdash e_1 : \tau \Rightarrow \Gamma'$ is immediate).

By definition, from $\mathbf{Cons}\,(H, R, \Gamma)$ we have $\mathbf{SAT}\,(H, R, \Gamma)$, in particular $x \in dom(R)$, $R(x) \in \mathbb{Z}$ and $[R]\,[R(x)/\nu]\varphi$. The refinement predicates φ still holds in the output environment, since nothing changes in the register after transition. Also from precondition of R-IFTRUE, we have $R(x) = 0$, thus x satisfies the refinement that $\nu = 0$. Thus $[R]\,[R(x)/\nu](\varphi \wedge \nu = 0)$ is trivially satisfied.

Case R-IFFALSE**:**

Similar to the case for R-IFTRUE.

Case R-MKREF: $\vdash^{D}_{conf} \left\langle H, R, \vec{F}, E[\mathbf{let}\ x = \mathbf{mkref}\ y\ \mathbf{in}\ e] \right\rangle$
$\left\langle H, R, \vec{F}, E[\mathbf{let}\ x = \mathbf{mkref}\ y\ \mathbf{in}\ e] \right\rangle \longrightarrow_D \left\langle H', R', \vec{F}, E[[x'/x]e] \right\rangle$
$a \notin dom(H) \qquad x' \notin dom(R)$
$H' = H\{a \mapsto R(y)\} \qquad R' = R\{x' \mapsto a\}$

By inversion (see R-DEREF) we have that for some Γ:

$$\vec{\ell} \vdash_{WF} \Gamma$$
$$\Gamma(y) = \tau_1 + \tau_2$$
$$\Theta \mid \vec{\ell} \mid \Gamma[y \hookleftarrow \tau_1], x : (\tau_2 \wedge_x x = y)\,\mathbf{ref}^1 \vdash e : \tau \Rightarrow \Gamma'$$
$$\mathbf{Cons}\,(H, R, \Gamma) \qquad x \notin dom(\Gamma')$$

We give $\Gamma'' = \Gamma[y \hookleftarrow \tau_1], x' : (\tau_2 \wedge_{x'} x' = y)\,\mathbf{ref}^1$, and must show that $\Theta \mid \vec{\ell} \mid \Gamma'' \vdash [x'/x]e : \tau \Rightarrow \Gamma'$ and $\mathbf{Cons}\,(H', R', \Gamma'')$. The first follows from Lemma 10 and the reasoning found in R-DEREF, and the second from the assumed well-formedness of $\tau_1 + \tau_2$.

It remains to show $\mathbf{Cons}\,(H', R', \Gamma'')$. To show that the output environment is consistent, we must show that $\mathbf{SATv}(H', R', R'(x'), (\tau_2 \wedge_{x'} x' = y)\,\mathbf{ref}^1)$ and $\mathbf{SATv}(H', R', R'(y), \tau_1)$. By reasoning similar to that in R-Deref, it suffices to show that $\mathbf{SATv}(H', R', R'(x'), \tau_2\,\mathbf{ref}^1)$ and $\mathbf{SATv}(H', R', R'(y), \tau_1)$. We know that $\mathbf{Cons}\,(H, R, \Gamma)$, from which we have $\mathbf{SAT}\,(H, R, \Gamma)$, in particular $y \in dom(R)$ and $\mathbf{SATv}(H, R, R(y), \Gamma(y))$. As $R \sqsubseteq R'$ and $\vec{\ell} \vdash_{WF} \Gamma$, from Lemmas 17 and 23, we have $\mathbf{SATv}(H, R, R(y), \tau_1 + \tau_2)$ implies $\mathbf{SATv}(H, R', R'(y), \tau_1 + \tau_2)$. By Lemma 24, we then have $\mathbf{SATv}(H', R', R(y), \tau_1 + \tau_2)$. Then by Lemma 16, we have $\mathbf{SATv}(H', R', v, \tau_1)$ and $\mathbf{SATv}(H', R', v, \tau_2)$ where $v = R(y)$. We then have $\mathbf{SATv}(H', R', R'(x'), \tau_2\,\mathbf{ref}^1)$ and $\mathbf{SATv}(H', R', R'(y), \tau_1)$ are satisfied.

We must also show that the ownership invariant is preserved. Then, it's to show $\forall\, a' \in dom(H).\mathbf{Own}\,(H', R', \Gamma'')(a') \leq 1$. From $\mathbf{Cons}\,(H, R, \Gamma)$ and Lemmas 15 and 25 we have:

$$\begin{aligned}\mathbf{Own}\,(H', R', \Gamma'') &= \Sigma_{z \in dom(\Gamma'')}\,\mathbf{own}\,(H', R'(z), \Gamma''(z)) \\ &= \Sigma_{z \in dom(\Gamma)}\,\mathbf{own}\,(H, R(z), \Gamma(z)) + \{a \mapsto 1\} \\ &= \mathbf{Own}\,(H, R, \Gamma) + \{a \mapsto 1\}\end{aligned}$$

Since $a \notin dom(H)$ and $\forall\, a' \in dom(H).\mathbf{Own}\,(H, R, \Gamma)(a') \leq 1$, we have $\forall\, a' \in dom(H).\mathbf{Own}\,(H', R', \Gamma'')(a') \leq 1$.

Case R-Assign: $\quad \vdash^D_{conf} \left\langle H, R, \vec{F}, E[y := x\,;\, e] \right\rangle$

$$\left\langle H, R, \vec{F}, E[y := x\,;\, e] \right\rangle \longrightarrow_D \left\langle H', R', \vec{F}, E[e] \right\rangle$$

$$a = R(y) \qquad H' = H\{a \hookleftarrow R(x)\} \qquad R' = R$$

By inversion (see the R-Deref case) we have that

$$\begin{aligned}&\Theta \mid \vec{\ell} \mid \Gamma[x : \tau_1 + \tau_2][y : \tau'\,\mathbf{ref}^1] \vdash y := x\,;\, e : \tau \Rightarrow \Gamma' \\ &\Theta \mid \vec{\ell} \mid \Gamma[x \hookleftarrow \tau_1][y \hookleftarrow (\tau_2 \wedge_y y =_{\tau_2} x)\,\mathbf{ref}^1] \vdash e : \tau \Rightarrow \Gamma' \\ &\mathbf{Cons}\,(H, R, \Gamma)\end{aligned}$$

We give $\Gamma'' = \Gamma[x \hookleftarrow \tau_1][y \hookleftarrow (\tau_2 \wedge_y y =_{\tau_2} x)\,\mathbf{ref}^1]$. That $\Theta \mid \vec{\ell} \mid \Gamma'' \vdash e : \tau \Rightarrow \Gamma'$ is immediate.

We must therefore show that $\mathbf{Cons}\,(H', R', \Gamma'')$. To show that the output environment is consistent, we must show that $\mathbf{SATv}(H', R, R(y), (\tau_2 \wedge_y y = x)\,\mathbf{ref}^1)$ and $\mathbf{SATv}(H', R, R(x), \tau_1)$. By reasoning similar to that in R-Deref, it suffices to show that $\mathbf{SATv}(H', R, R(y), \tau_2\,\mathbf{ref}^1)$ and $\mathbf{SATv}(H', R, R(x), \tau_1)$.

From $\mathbf{Cons}\,(H, R, \Gamma)$, we know that $\mathbf{SAT}\,(H, R, \Gamma)$, in particular, $\mathbf{SATv}(H, R, R(x), \Gamma(x))$. If we show that $\mathbf{own}\,(H, R(x), \tau_1 + \tau_2)(a) = 0$ and $H \approx_a H'$, then, by Lemma 22, we will obtain $\mathbf{SATv}(H', R, R(x), \tau_1 + \tau_2)$, from which, by Lemma 16, $\mathbf{SATv}(H', R, R(x), \tau_1)$ and $\mathbf{SATv}(H', R, R(x), \tau_2)$ follow. We then have $\mathbf{SATv}(H', R, R(y), \tau_2\,\mathbf{ref}^1)$ and $\mathbf{SATv}(H', R, R(x), \tau_1)$ as $H'(R(y)) = R(x)$. (That any other variables z is consistent will follow from $\mathbf{own}\,(H, R(z), \Gamma(z))(a) = 0$ as proved below and lemma 22.)

To show $\mathbf{own}\,(H, R(x), \tau_1 + \tau_2)(a) = 0$, we define O'_0, O''_0, O'_1 and O''_1 as below:

$$\begin{aligned}
\mathbf{Own}\,(H, R, \Gamma) &= O'_0 + O''_0 \\
\mathbf{Own}\,(H', R, \Gamma'') &= O'_1 + O''_1 \\
O'_0 &= \Sigma_{z \in dom(\Gamma) \setminus \{y,x\}}\, \mathbf{own}\,(H, R(z), \Gamma(z)) \\
O''_0 &= \mathbf{own}\,(H, R(y), \Gamma(y)) + \mathbf{own}\,(H, R(x), \Gamma(x)) \\
O'_1 &= \Sigma_{z \in dom(\Gamma'') \setminus \{y,x\}}\, \mathbf{own}\,(H', R(z), \Gamma''(z)) \\
O''_1 &= \mathbf{own}\,(H', R(y), \Gamma''(y)) + \mathbf{own}\,(H', R(x), \Gamma''(x))
\end{aligned}$$

By the definition of the ownership function, $\Gamma(y) = \tau' \,\mathbf{ref}^1$ and $\Gamma(x) = \tau_1 + \tau_2$, we have:

$$\begin{aligned}
O''_0 &= \mathbf{own}\,(H, H(R(y)), \tau') + \{a \mapsto 1\} + \mathbf{own}\,(H, R(x), \tau_1 + \tau_2) \\
O''_1 &= \mathbf{own}\,(H', H'(R(y)), \tau_2) + \{a \mapsto 1\} + \mathbf{own}\,(H', R(x), \tau_1)
\end{aligned}$$

As $\mathbf{Own}\,(H, R, \Gamma)(a) \leq 1$ (from $\mathbf{Cons}\,(H, R, \Gamma)$) and from

$$\begin{aligned}
\mathbf{Own}\,(H, R, \Gamma)(a) &= O'_0(a) + O''_0(a) \\
&= O'_0(a) + \mathbf{own}\,(H, H(R(y)), \tau')(a) + 1 + \mathbf{own}\,(H, R(x), \tau_1 + \tau_2)(a) \\
&= 1
\end{aligned}$$

we have that:

$$\begin{aligned}
\mathbf{own}\,(H, H(R(y)), \tau')(a) &= \mathbf{own}\,(H, R(x), \tau_1 + \tau_2)(a) \\
&= O'_0(a) = 0
\end{aligned}$$

We now show that $H \approx_a H'$. The first two conditions are clear, so it remains to show that, for any n, $H \vdash a \Downarrow n$ iff $H' \vdash a \Downarrow n$. From Lemma 20, we have $H \vdash a \Downarrow |\tau' \,\mathbf{ref}^1\,|$, and a proof by contradiction gives that $|\tau' \,\mathbf{ref}^1\,|$ is the only such n for which $H \vdash a \Downarrow n$. We now argue the forward case for the bi-implication, the backwards case follows similar reasoning.

Given $H \vdash a \Downarrow |\tau' \,\mathbf{ref}^1\,|$, we must show $H\{a \hookleftarrow R(x)\} \vdash a \Downarrow |\tau_2 \,\mathbf{ref}^1\,|$, for which it suffices to show $H\{a \hookleftarrow R(x)\} \vdash R(x) \Downarrow |\tau_2|$. From our requirement that τ' and τ_2 (and therefore $\tau_1 + \tau_2$) have similar shapes, we have $|\tau'| = |\tau_2| = |\tau_1 + \tau_2|$. By inverting the well-typing of the input configuration, we must have $\mathbf{SATv}(H, R, R(x), \tau_1 + \tau_2)$, thus by Lemma 20 we must have $H \vdash R(x) \Downarrow |\tau_2|$. As $|\tau_2| = |\tau'| < |\tau' \,\mathbf{ref}^1\,|$, a cannot be reachable from $R(x)$ in H (otherwise we would have a reaches an integer along multiple heap paths of differing lengths, a clear contradiction). Then the value of a in H is irrelevant to the derivation of $H \vdash R(x) \Downarrow |\tau_2|$, whereby $H\{a \hookleftarrow R(x)\} \vdash R(x) \Downarrow |\tau_2|$ must hold.

Then, it's to show $\forall\, a' \in dom(H').\mathbf{Own}\,(H', R, \Gamma'')(a') = (O'_1 + O''_1)(a') = O'_1(a') + O''_1(a') \leq 1$. For every $z \in dom(\Gamma) \setminus \{y, x\}$ (and similarly for Γ''), we have $\Gamma(z) = \Gamma''(z)$. Further, from $O'_0(a) = 0$ above, we must have $\mathbf{own}\,(H, R(z), \Gamma(z))(a) = 0$ for all such z. As $H \approx_a H'$, by Lemma 19, we have that $O'_0 = O'_1$. Then, from $\forall\, a' \in dom(H).\mathbf{Own}\,(H, R, \Gamma)(a') = O'_0(a') + O''_0(a') \leq 1$, it suffices to show that $\forall\, a' \in dom(H).O''_1(a') \leq O''_0(a')$.

We first consider the case for a:

$$O_1''(a) = \mathbf{own}\,(H', H'(R(y)), \tau_2)(a) + \mathbf{own}\,(H', R(x), \tau_1)(a) + 1$$
$$O_0''(a) = \mathbf{own}\,(H, R(x), \tau_1 + \tau_2)(a) + \mathbf{own}\,(H, H(R(y)), \tau')(a) + 1$$

From above, we have $\mathbf{own}\,(H, R(x), \tau_2 + \tau_1)(a) = \mathbf{own}\,(H, H(R(y)), \tau')(a) = 0$. By Lemma 19 and $H \approx_a H'$, we have $\mathbf{own}\,(H, R(x), \tau_2{+}\tau_1) = \mathbf{own}\,(H', R(x), \tau_2{+}\tau_1)$. Also by Lemma 15, we have $\mathbf{own}\,(H', R(x), \tau_2 + \tau_1) = \mathbf{own}\,(H', R(x), \tau_1) + \mathbf{own}\,(H', R(x), \tau_2)$. From $H'(R(y)) = R(x)$, we therefore have $\mathbf{own}\,(H', H'(R(y)), \tau_2)(a) = \mathbf{own}\,(H', R(x), \tau_1)(a) = 0$, and thus:

$$O_1''(a) = \mathbf{own}\,(H', H'(R(y)), \tau_2)(a) + \mathbf{own}\,(H', R(x), \tau_1)(a) + 1 = 1 = O_0''(a)$$

Next, consider some $a \neq a'$;

$$O_1''(a') = \mathbf{own}\,(H', H'(R(y)), \tau_2)(a') + \mathbf{own}\,(H', R(x), \tau_1)(a')$$
$$O_0''(a') = \mathbf{own}\,(H, R(x), \tau_2 + \tau_1)(a') + \mathbf{own}\,(H, H(R(y)), \tau')(a')$$

By reasoning similar to the case for $a = a'$, we have $O_1''(a') \leq \mathbf{own}\,(H, R(x), \tau_2 + \tau_1)(a') \leq O_0''(a')$. We therefore conclude that $\forall\, a' \in dom(H').\mathbf{Own}\,(H', R, \Gamma'')(a') \leq 1$.

Case R-Alias**:**
$$\vdash^{D}_{conf} \left\langle H, R, \vec{F}, E[\mathbf{alias}(x = y)\,; e] \right\rangle$$
$$\left\langle H, R, \vec{F}, E[\mathbf{alias}(x = y)\,; e] \right\rangle \longrightarrow_D \left\langle H, R, \vec{F}, E[e] \right\rangle$$
$$R(x) = R(y)$$

By inversion (see R-Deref) we have for some Γ that:

$$\Theta \mid \vec{\ell} \mid \Gamma[x : \tau_1\,\mathbf{ref}^{r_1}][y : \tau_2\,\mathbf{ref}^{r_2}] \vdash \mathbf{alias}(x = y)\,; e : \tau \Rightarrow \Gamma'$$
$$\Theta \mid \vec{\ell} \mid \Gamma[x \hookleftarrow \tau_1'\,\mathbf{ref}^{r_1'}][y \hookleftarrow \tau_2'\,\mathbf{ref}^{r_2'}] \vdash e : \tau \Rightarrow \Gamma'$$
$$\tau_1\,\mathbf{ref}^{r_1} + \tau_2\,\mathbf{ref}^{r_2} \approx \tau_1'\,\mathbf{ref}^{r_1'} + \tau_2'\,\mathbf{ref}^{r_2'}$$
$$\mathbf{Cons}\,(H, R, \Gamma)$$

We give $\Gamma'' = \Gamma[x \hookleftarrow \tau_1'\,\mathbf{ref}^{r_1'}][y \hookleftarrow \tau_2'\,\mathbf{ref}^{r_2'}]$, and must show $\Theta \mid \vec{\ell} \mid \Gamma'' \vdash e : \tau \Rightarrow \Gamma'$ and $\mathbf{Cons}\,(H, R, \Gamma'')$. The first is immediate.

To show $\mathbf{Cons}\,(H, R, \Gamma'')$ we first define:

$$\tau_{p\,1} = \tau_1\,\mathbf{ref}^{r_1}$$
$$\tau_{p\,2} = \tau_2\,\mathbf{ref}^{r_2}$$
$$\tau_{q\,1} = \tau_1'\,\mathbf{ref}^{r_1'}$$
$$\tau_{q\,2} = \tau_2'\,\mathbf{ref}^{r_2'}$$
$$\tau_q = \tau_{q\,1} + \tau_{q\,2}$$
$$\tau_p = \tau_{p\,1} + \tau_{p\,2}$$

We thus have $\tau_q \approx \tau_p$.

We know that $\mathbf{Cons}\,(H, R, \Gamma)$, from which we have $\mathbf{SAT}\,(H, R, \Gamma)$, in particular $\mathbf{SATv}(H, R, R(y), \Gamma(y) = \tau_2\,\mathbf{ref}^{r_2})$ and $\mathbf{SATv}(H, R, R(x), \Gamma(x) = \tau_1\,\mathbf{ref}^{r_1})$. From $\tau_{p\,1}+\tau_{p\,2} = \tau_p$ and Lemma 16, we have $\mathbf{SATv}(H, R, v, \tau_{p\,1})$ and $\mathbf{SATv}(H, R, v, \tau_{p\,2})$ imply $\mathbf{SATv}(H, R, v, \tau_p)$, where $v = H(R(y)) = H(R(x))$. From $\tau_q \approx \tau_p$ and Lemma 13, we have that $\mathbf{SATv}(H, R, v, \tau_p)$ implies $\mathbf{SATv}(H, R, v, \tau_q)$. From Lemma 16 we also have that $\mathbf{SATv}(H, R, v, \tau_q)$ implies $\mathbf{SATv}(H, R, v, \tau_{q\,1})$ and $\mathbf{SATv}(H, R, v, \tau_{q\,2})$, where again $v = H(R(y)) = H(R(x))$.

Then from the reasoning above, the refinements of $\tau_{q\,1}$ and $\tau_{q\,2}$ are valid and $\mathbf{Cons}\,(H, R, \Gamma'')$ holds.

Then, it's to show $\forall\, a \in \mathit{dom}(H).\mathbf{Own}\,(H, R, \Gamma'')(a) \leq 1$. To prove that $\mathbf{Own}\,(H, R, \Gamma) = \mathbf{Own}\,(H, R, \Gamma'')$ follows from:

$$\begin{aligned}\mathbf{own}\,(H, R(x), \tau_1\,\mathbf{ref}^{r_1}) + \mathbf{own}\,(H, R(y), \tau_2\,\mathbf{ref}^{r_2}) =\\ \mathbf{own}\,(H, R(x), \tau_1'\,\mathbf{ref}^{r_1'}) + \mathbf{own}\,(H, R(y), \tau_2'\,\mathbf{ref}^{r_2'})\end{aligned}$$

which follows immediately from the conditions $\tau_1\,\mathbf{ref}^{r_1} + \tau_2\,\mathbf{ref}^{r_2} \approx \tau_1'\,\mathbf{ref}^{r_1'} + \tau_2'\,\mathbf{ref}^{r_2'}$, $R(x) = R(y)$, and Lemmas 14 and 15.

Case R-AliasPtr**:**

By reasoning similar to the R-Alias case.

Case R-AliasFail,R-AliasPtrFail**:**

The result configuration **AliasFail** is trivially well-typed.

Case R-Assert**:** $\vdash^D_{conf} \left\langle H, R, \vec{F}, E[\mathbf{assert}(\varphi)\,;\, e]\right\rangle \qquad \Gamma \models [R]\,\varphi$
$\left\langle H, R, \vec{F}, E[\mathbf{assert}(\varphi)\,;\, e]\right\rangle \longrightarrow_D \left\langle H, R, \vec{F}, E[e]\right\rangle$

By inversion (see R-Deref) we can obtain $\Theta \mid \vec{\ell} \mid \Gamma \vdash \mathbf{assert}(\varphi)\,;\, e : \tau \Rightarrow \Gamma'$ and $\Theta \mid \vec{\ell} \mid \Gamma \vdash e : \tau \Rightarrow \Gamma'$, and the result follows immediately by taking $\Gamma'' = \Gamma$.

Case R-AssertFail**:** $\vdash^D_{conf} \left\langle H, R, \vec{F}, E[\mathbf{assert}(\varphi)\,;\, e]\right\rangle$
$\left\langle H, R, \vec{F}, E[\mathbf{assert}(\varphi)\,;\, e]\right\rangle \longrightarrow_D \mathbf{AssertFail}$
$\Theta \mid \vec{\ell} \mid \Gamma \vdash \mathbf{assert}(\varphi)\,;\, e : \tau \Rightarrow \Gamma'$

By inversion (see the R-Deref case) we have that $\Gamma \models \varphi$, i.e., $\models [\![\Gamma]\!] \implies \varphi$, for some Γ such that $\mathbf{Cons}\,(H, R, \Gamma)$. From Lemma 11 we therefore have $\models [R]\,[\![\Gamma]\!]$. From the precondition of R-AssertFail we have that $\not\models [R]\,\varphi$. But from $\models [\![\Gamma]\!] \implies \varphi$ and $\models [R]\,[\![\Gamma]\!]$ we can conclude that $\models [R]\,\varphi$, yielding a contradiction. We therefore conclude that this case is impossible.

Case R-Call**:** $\vdash^D_{conf} \left\langle H, R, \vec{F}, E[\mathbf{let}\, x = f^{\ell}(y_1, \ldots, y_n)\,\mathbf{in}\, e']\right\rangle$
$f \mapsto (x_1, .., x_n)e \in D$
$\left\langle H, R, \vec{F}, E[\mathbf{let}\, x = f^{\ell}(y_1, \ldots, y_n)\,\mathbf{in}\, e']\right\rangle$
$\longrightarrow_D \left\langle H, R, E[\mathbf{let}\, x = []^{\ell}\,\mathbf{in}\, e'] : \vec{F}, [y_1/x_1]\cdots[y_n/x_n]e\right\rangle$

We must show that $\vdash^D_{conf} \left\langle H, R, E[\mathbf{let}\, x = []^{\ell}\,\mathbf{in}\, e'] : \vec{F}, [y_1/x_1]\cdots[y_n/x_n]e\right\rangle$ for some Γ''.

By inversion on the configuration typing, we have that, for some Γ:

$$\Theta \mid \vec{\ell} \mid \Gamma \vdash E[\mathbf{let}\, x = f^{\ell}(y_1, \ldots, y_n)\, \mathbf{in}\, e'] : \tau_n \Rightarrow \Gamma_n.$$

By Lemma 7, we then have for some τ, and Γ' that:

$$\Theta \mid \vec{\ell} \mid \Gamma \vdash \mathbf{let}\, x = f^{\ell}(y_1, \ldots, y_n)\, \mathbf{in}\, e' : \tau \Rightarrow \Gamma'$$
$$\Theta \mid [] : \tau \Rightarrow \Gamma' \mid \vec{\ell} \vdash_{ectx} E : \tau_n \Rightarrow \Gamma_n$$

Taking $\tau_1 = \tau, \Gamma_1 = \Gamma', \Gamma_0 = \Gamma, \Gamma_2 = \Gamma_n, \tau_2 = \tau_n$, by Lemma 28 we have, for some τ''', Γ''':

$$\Theta \mid \ell : \vec{\ell} \mid \Gamma \vdash \sigma_x e : \tau''' \Rightarrow \Gamma'''$$
$$\Theta \mid [] : \tau''' \Rightarrow \Gamma''' \mid \vec{\ell} \vdash_{ectx} E[\mathbf{let}\, x = []^{\ell}\, \mathbf{in}\, e'] : \tau_n \Rightarrow \Gamma_n$$

where:

$$\sigma_x = [y_1/x_1] \cdots [y_n/x_n]$$
$$\Theta(f) = \forall \lambda.\, \langle x_1 : \tau_i, \ldots, x_n : \tau_n \rangle \rightarrow \langle x_1 : \tau'_1, \ldots, x_n : \tau'_n \mid \tau_p \rangle$$

We therefore take $\Gamma'' = \Gamma$.

We must also prove that $\forall i \in \{1..n+1\}\, .\Theta \mid [] : \tau_i \Rightarrow \Gamma_i \mid \vec{\ell}_{i-1} \vdash_{ectx} E'_{i-1} : \tau_{i-1} \Rightarrow \Gamma_{i-1}$ where $E'_n = E[\mathbf{let}\, x = []^{\ell}\, \mathbf{in}\, e']$ and $E'_i = E_i (0 \leq i < n)$, which can be divided into proving $\forall i \in \{1..n\}\, .\Theta \mid [] : \tau_i \Rightarrow \Gamma_i \mid \vec{\ell}_{i-1} \vdash_{ectx} E'_{i-1} : \tau_{i-1} \Rightarrow \Gamma_{i-1}$ and $\Theta \mid [] : \tau_{n+1} \Rightarrow \Gamma_{n+1} \mid \vec{\ell}_n \vdash_{ectx} E'_n : \tau_n \Rightarrow \Gamma_n$. The first follows by inversion on $\vdash^{D}_{conf} \left\langle H, R, \vec{F}, E[\mathbf{let}\, x = f^{\ell}(y_1, \ldots, y_n)\, \mathbf{in}\, e'] \right\rangle$. To show the latter, we define $\Gamma_{n+1} = \Gamma'''$ and $\tau_{n+1} = \tau'''$, whereby the well-typing holds from the result of applying Lemma 28 above.

Finally, $\mathbf{Cons}\,(H, R, \Gamma'')$ follows immediately from $\mathbf{Cons}\,(H, R, \Gamma)$ and $\Gamma'' = \Gamma$.

D Proof of Progress

We first state the standard decomposition lemma.

Lemma 29 (Decomposition). *For any term e, either $e = x$ for some x or there exists some E and e' where $E[e'] = e$ and one of the following cases hold:*

1. $e' = \mathbf{let}\, x = \mathbf{mkref}\, y\, \mathbf{in}\, e''$
2. $e' = \mathbf{let}\, x = y\, \mathbf{in}\, e''$
3. $e' = \mathbf{let}\, x = n\, \mathbf{in}\, e''$
4. $e' = \mathbf{let}\, x = *y\, \mathbf{in}\, e''$
5. $e' = \mathbf{let}\, x = f^{\ell}(y_1, \ldots, y_n)\, \mathbf{in}\, e''$
6. $e' = x\,;\, e''$
7. $e' = \mathbf{alias}(x = y)\,;\, e''$
8. $e' = \mathbf{alias}(x = *y)\,;\, e''$

9. $e' = \mathbf{ifz}\, x\, \mathbf{then}\, e_1\, \mathbf{else}\, e_2$
10. $e' = \mathbf{assert}(\varphi)\,;\, e''$
11. $e' = x := y\,;\, e$

Proof. Straightforward induction on e.

Proof (Progress; Lemma 4). By inversion on $\vdash^{D}_{conf}$ $\mathbf{C}$, either $\mathbf{C} = \mathbf{AliasFail}$ or $\mathbf{C} = \left\langle H, R, \vec{F}, e \right\rangle$. In the former case the result is immediate. In the latter case we have that $\Theta \mid \vec{\ell} \mid \Gamma \vdash e : \tau \Rightarrow \Gamma'$ for some τ, Γ and Γ', and further from Lemma 29, we have that either $e = x$ for some x or there exists some E or e' where $e = E[e']$ and e' meets one of the cases in Lemma 29.

In the case $e = x$, we further make case analysis on the form of $\vec{F}$. The case where $\vec{F} = \epsilon$ is immediate; In the other case where $\vec{F} = F : \vec{F}'$, the configuration can step to $\left\langle H, R, \vec{F}, F[x] \right\rangle$ according to R-Var.

For the remaining cases where $e = E[e']$, by the well-typing of e with respect to Γ and Lemma 7, we have that $\Theta \mid \mathcal{L} \mid \Gamma \vdash e' : \tau_0 \Rightarrow \Gamma_0$ some τ_0 and Γ_0.

We now treat the remaining forms of e'

Case: $e' = \mathbf{let}\, x = *y\, \mathbf{in}\, e''$

By inversion (Lemma 6) and Lemma 12 we must have that for some Γ_p where $\mathbf{Cons}\,(H, R, \Gamma_p)$ that $y \in dom(\Gamma_p)$ and $\Gamma_p(y) = \tau'\, \mathbf{ref}^{r}$. From $\mathbf{Cons}\,(H, R, \Gamma_p)$ we must have $y \in dom(R)$ and further $\mathbf{SATv}(H, R, R(y), \tau'\, \mathbf{ref}^{r'})$, from which we must have $R(y) = a$ and $a \in dom(H)$. Then $\mathbf{C}$ can step according to R-Deref.

Case: $e' = \mathbf{let}\, x = y\, \mathbf{in}\, e''$

Again, by Lemmas 6 and 12 and the definition of **Cons**, we must have that $y \in dom(R)$, and the system can step according to R-LetVar.

Case: $e' = \mathbf{let}\, x = \mathbf{mkref}\, y\, \mathbf{in}\, e''$

Similar to the R-LetVar case above.

Case: $e' = \mathbf{let}\, x = n\, \mathbf{in}\, e''$
$e' = x\,;\, e''$
$e' = \mathbf{assert}(\varphi)\,;\, e''$

The first two can trivially step according to R-LetInt and R-Seq respectively. the last can step according to R-Assert or R-AssertFalse (although by Lemmas 2 and 3 the latter is impossible).

Case: $e' = \mathbf{alias}(x = y)\,;\, e''$

Again by Lemmas 6 and 12 and that $\mathbf{Cons}\,(H, R, \Gamma_p)$ implies x and y are bound to addresses in the register file, we have that the configuration can step according to R-Alias or R-AliasFail.

Case: $e' = \mathbf{alias}(x = *y)\,;\, e''$

Similar to the case above, we must have that x is bound to an address in the register file, and that y is bound to an address that is itself mapped to an address in the heap H. Then the configuration may step according to R-AliasPtr or R-AliasPtrFail

Case: $e = \mathbf{ifz}\, x\, \mathbf{then}\, e_1\, \mathbf{else}\, e_2$

As above, from the well-typing we must have that x is bound in R to some integer n. Then the configuration may step according to R-IfTrue or R-IfFalse depending on whether $n = 0$ or $n \neq 0$.

Case: $e' = x := y\,; e''$

From the well-typing of e', Lemmas 6 and 12 and the definition of **Cons**, we must have that $y \in dom(R)$, $x \in dom(R)$, $R(x) = a$, and $a \in dom(H)$. Then the configuration can step according R-Assign.

Case: $e' = \mathbf{let}\, x = f^{\ell}(y_1, \ldots, y_n)\, \mathbf{in}\, e''$

From the well-typing of the function call we must have that $f \in dom(\Theta)$. From $\Theta \vdash D$ in the precondition of $\vdash^{D}_{conf} \mathbf{C}$, we must have that $f \mapsto (x_1, .., x_j)e''' \in D$. Then from T-FunDef we must have that $j = n$ whereby the configuration can step according to R-Call.