Verifying Safety Policies with Size Properties and Alias Controls

Wei-Ngan Chin[1,2], Siau-Cheng Khoo[1], Shengchao Qin[1,2], Corneliu Popeea[1], and Huu Hai Nguyen[2]
[1]Department of Computer Science, National University of Singapore
[2]Computer Science Programme, Singapore-MIT Alliance
{chinwn|khoosc|qinsc|popeeaco|nguyenh2}@comp.nus.edu.sg

Abstract

Many software properties can be analysed through a relational size analysis on each function's inputs and outputs. Such relational analysis (through a form of dependent typing) has been successfully applied to declarative programs, and to restricted imperative programs; but it has been elusive for object-based programs. The main challenge is that objects may mutate *and they may be* aliased*.*

In this paper, we show how safety policies of programs can be analysed by tracking size properties of objects and be enforced by objects' invariants and the preconditions of methods. We propose several new ideas to allow both mutability *and* sharing *of objects, whilst aiming for* precision *in our analysis. We introduce the concept of* size-immutability *to facilitate sharing, and also a set of* alias controls *to track unaliased objects whose size properties may change.*

We formalise our results through a set of advance type checking rules for an object-based imperative language. We re-affirm the utility of the proposed type system by showing how a variety of software properties can be automatically verified according to size-inspired safety policies.

1 Introduction

Size properties are extremely useful in program analysis to support program optimization and verification. Each piece of data structure typically has one or more associated size information. For simple types, such as `int`, `bool` or `file`, size property can capture their values (or abstract states). For aggregate types, size may refer to the length/height of linked structures, cardinality of collections or bounds of arrays. We may also ascribe size property to memory utilization in order to estimate space required for execution. For functions, size property can capture relations between input and output sizes, input preconditions, and also size invariance for parameters in recursion[7].

Past research into size properties has been most explored for functional [18, 17, 26, 7] and logic programming languages [25, 3]. Such properties have been applied to a range of important problems, including: (i) array bound check elimination [28, 29], (ii) memory space estimates for both region model [17] and heap-with-reuse model [16], (iii) complexity analysis [13], (iv) termination analysis[3, 20], and (v) support for safer programming where size invariants (e.g. balanced AVL trees) are guaranteed at compile time[26]. These applications are important towards supporting the goal of more reliable (and rigorous) software engineering processes. They constitute a realization of the automated verification for size-based software contracts.

To track size properties accurately, a *relational analysis* is required. This is expressible through a form of *dependent typing*[26], known as *sized type*[18, 7], that has been applied to declarative paradigms. To the best of our knowledge, no one has successfully designed a relational size analysis (or corresponding dependent type system) for static reasoning of object-based imperative programs.

Two key problems with objects are *mutability* and *aliasing*. Objects may mutate and their property changes be tracked. Objects may also alias with each other, and this must be analysed and/or be controlled. To highlight the problem of aliasing in the presence of mutation, consider the following `Int` object type expressed as an abstract data type (ADT), with an increment method.

```
adt Int { int val ;
void incr(Int x, int v) { x.val = x.val + v } }
```

Each object of the `Int` type has a primitive integer `val` field, whose size we may wish to track. To do so, we introduce an annotated object type, say $\texttt{Int}\langle\texttt{n}\rangle$, where `n` is a size variable for capturing the value of the `val` field via `n=m` and with `true` as the object's invariant, as follows:

```
adt Int⟨n⟩ where n=m ; true { int⟨m⟩ val ; ··· }
```

Consider two variables `a` and `b` of the `Int` type. We can provide an annotated type for these variables using $\texttt{a} :: \texttt{Int}\langle\texttt{x}\rangle, \texttt{b} :: \texttt{Int}\langle\texttt{y}\rangle$, where `x` and `y` denote the sizes of current objects referenced by `a` and `b`, respectively. Consider now a program fragment with aliasing and updates.

```
Int a = new Int(5);     // x' = 5
Int b = a;              // y' = x ∧ x' = x
incr(a, 1);             // x' = x + 1
```

Ideally, we would like to perform size analysis, locally for each statement, as shown on the RHS. This would result in $\mathtt{x}' = 6 \wedge \mathtt{y}' = 5$ after composing the effects of the three statements. However, this reasoning is unsound, unless we take alias relationships into account. Note the use of the prime notation (e.g. $\mathtt{x}', \mathtt{y}'$) to denote new states of the size variables after each statement.

Accurately tracking the sizes of mutable objects require the aliasing problem to either be *carefully analysed* or be *tightly controlled*. In this paper, we chose the latter path, as analysing global aliases could quickly become unmanageable. Our main contributions are:

- We have successfully designed a **sized type** system for objects, which is able to track size information accurately. To the best of our knowledge, no one has applied such an advanced relational analysis to object-based imperative programs before.
- Each object declaration is expressed as an ADT. Our motivation for using ADTs is to allow each set of software properties to be specified via a **safety policy** (or **protocol**), whose correct usage can be automatically verified against the *size invariant* of each ADT and the *preconditions* on its library of operations.
- Our approach provides an elegant balance between *mutability* and *sharing*, while working towards accurate size tracking. We achieve this using a novel property, called **size-mutability**, obtained with a set of alias controls. We use *read-only* annotation for size-immutability, and *uniqueness* and *lent* annotations to limit aliasing for mutable size properties.
- We formalise a set of advanced type checking rules with alias controls and size constraints. We have implemented a **prototype** to validate our proposal and have also proven key **safety theorems** to confirm the desired properties of well-typed programs.

Sec 2 proposes a core imperative language with objects, and highlights how size properties are captured via an annotated type system. Sec 3 introduces the basic concept of size-mutability and how to track it using alias controls. Sec 4 shows how safety protocols may be specified through size invariants and methods' preconditions in ADTs. Sec 5 presents a set of advanced size-checking type rules, followed by key correctness properties. Sec 8 describes related work, followed by a short conclusion.

2 Kernel Language

We design a simple kernel language, called OIMP, with alias and size annotations. We present an object-based rather than an object-oriented language to focus on the key issues for the size analysis of objects. We intend for OIMP to be used as a backend to existing programming languages, such as Java or C. Its syntax is given in Fig. 1. Note that the suffix notation y^* denotes a list of zero or more distinct syntactic terms that are suitably separated. For example, $(t\ v)^*$ denotes $t_1\ v_1, \ldots, t_n\ v_n$ where $n \geq 0$.

$$
\begin{array}{rl}
P & ::= \mathit{prim}^*\ \mathit{user}^*\ \mathit{meth}^* \\
\mathit{prim} & ::= \mathtt{adt}\ \mathit{pn}\langle n_1, \ldots, n_p\rangle\ \mathtt{where}\ \{n^*\};\ \phi_I\ \{\mathit{pmeth}^*\} \\
\mathit{user} & ::= \mathtt{adt}\ \mathit{cn}\langle n_1, \ldots, n_p\rangle\ \mathtt{where}\ \phi\ ;\ \phi_I\ \{\mathit{field}^*\ \mathit{meth}^*\} \\
\mathit{pn} & ::= \mathtt{void} \mid \mathtt{bool} \mid \mathtt{int} \mid \mathtt{float} \mid \cdots \\
\mathit{pmeth} & ::= t\ \mathit{mn}\ ((t\ v)^*)\ \mathtt{where}\ \phi_{pr}\ ;\ \phi_{po} \\
\mathit{meth} & ::= t\ \mathit{mn}\ ((t\ v)^*)\ \mathtt{where}\ \phi_{pr}\ ;\ \phi_{po}\ \{e\} \\
\tau & ::= \mathit{cn} \mid \mathit{pn} \qquad t ::= \tau\langle n^*\rangle\mathtt{@}A \\
w & ::= v \mid v.f \qquad \mathit{field} ::= t\ f \qquad A ::= \mathtt{U} \mid \mathtt{L} \mid \mathtt{S} \mid \mathtt{R} \\
e & ::= (\mathit{cn})\ \mathtt{null} \mid k \mid w \mid w = e \mid \mathtt{if}\ v\ \mathtt{then}\ e_1\ \mathtt{else}\ e_2 \\
& \quad \mid \mathtt{new}\ \mathit{cn}(v^*) \mid \mathit{mn}\ (v^*) \mid t\ v = e_1\ ;\ e_2 \\
& \quad \mid \mathtt{bind}\ v\ \mathtt{of}\ (v_1, .., v_n)\ \mathtt{in}\ e \\
\phi \in & \ \mathbf{F} \qquad \textit{(Presburger Size Constraint)} \\
& ::= b \mid \phi_1 \wedge \phi_2 \mid \phi_1 \vee \phi_2 \mid \neg\phi \mid \exists n \cdot \phi \mid \forall n \cdot \phi \\
b \in & \ \mathbf{BExp} \qquad \textit{(Boolean Expression)} \\
& ::= \mathtt{true} \mid \mathtt{false} \mid \alpha_1 = \alpha_2 \mid \alpha_1 < \alpha_2 \mid \alpha_1 \leq \alpha_2 \\
\alpha \in & \ \mathbf{AExp} \qquad \textit{(Arithmetic Expression)} \\
& ::= c \mid n \mid c * \alpha \mid \alpha_1 + \alpha_2 \mid -\alpha \mid \mathit{max}(\alpha_1, \alpha_2) \mid \mathit{min}(\alpha_1, \alpha_2)
\end{array}
$$

where $c \in Z$ is an integer constant ; $n \in SV$ is a size variable
$f \in Fd$ is a field name ; $v \in Var$ is an object variable

Figure 1. Syntax for the OIMP language

Local variable declaration is supported by block structure of the form: $(t\ v = e_1; e_2)$ with e_2 denoting its result. For convenience, we may also use τ instead of t and rely on fresh size variables and alias defaults for annotations. We assume a call-by-value mechanism for OIMP, where values (primitives or references) are passed as arguments to parameters of methods. For simplicity, we do not allow the parameters to be updated (or re-assigned) with different values. There is no loss of generality, as we can always copy parameters to local variables for updating, without altering the observable behaviour of method calls. Each user-defined method declaration *meth* contains a precondition ϕ_{pr} and a postcondition ϕ_{po}. The former must be satisfied for each caller, while the latter must be ensured by the method's body. For size constraints used (e.g ϕ_{pr} and ϕ_{po}), we restrict to Presburger arithmetic, as decidable (and practical) constraint solvers exist, e.g. [23]. For each reference (variable, parameter or field) of an object type, we add an alias control, named A, as follows: $\tau\langle \mathtt{n}_1, \ldots, \mathtt{n}_m\rangle\mathtt{@A}$. Each alias control could either be R, U, S or L to denote *read-only*, *unique*, *shared* or *lent-once*, respectively. They are primarily used to track the mutability of size properties for objects and uniqueness of their references, and will be elaborated later in Sec 3.

The OIMP language has been kept simple to make eas-

ier the formulation of static and dynamic semantics. Typical language constructs, such as multi-declaration block, sequence, calls with complex arguments, *etc.* can be automatically translated to constructs in OIMP. Also, loops can be viewed as syntactic abbreviations for tail-recursive methods, and are supported by our analysis.

2.1 Primitive ADTs

We provide two kinds of object type, namely primitive (*prim*) type or user-definable (*user*) type. Both are expressed using the ADT mechanism to allow safety policies on objects to be specified. For each primitive ADT, we specify a set of size variables $n_1, \ldots, n_p$ with a size invariant ϕ_I that holds for each object, and $\{n^*\}$ to capture a subset of the size variables that are mutable. We also specify a set of primitive method declarations that are trusted to meet its postcondition ϕ_{po}, whenever its precondition ϕ_{pr} is satisfied. An example is the `bool` primitive ADT, with no mutable size variables and $0 \leq n \leq 1$ as its size invariant.

```
adt bool⟨n⟩ where {} ; 0≤n≤1
{ bool⟨r⟩@S not(bool⟨m⟩@S) where
      true; m′=m∧(m=0∧r=1∨m=1∧r=0); ··· }
```

Another example is the array ADT, with a safety policy that no array bound violations ever occur, as given below:

```
adt Array⟨s⟩ where {}; s>0
{Array⟨s⟩@U newArray(int⟨n⟩@S a, int⟨v⟩@S val)
   where n>0 ; s=n∧naX{n, v};
 int⟨r⟩@S sub(Array⟨s⟩@S a, int⟨i⟩@S idx)
   where 0 ≤ i<s ; naX{s, i} ;
 void⟨⟩@S update(Array⟨s⟩@S a, int⟨i⟩@S idx,
   int⟨v⟩@S val) where 0 ≤ i<s ; naX{s, i, v} ;
 int⟨r⟩@S len(Array⟨s⟩@S a) where true ; r=s∧s′=s; }
```

Safe array bound is guaranteed by the precondition $0 \leq i < s$ for the `sub` method (to retrieve an element of the array) and `update` method (to modify a location of the array), while $n > 0$ is for the `newArray` method (to create a new array of positive size). Furthermore, the array's size is never changed by any of the operations. This may be specified through $na\mathcal{X}$ which returns a formula where the original and prime variables are made equal, e.g. $na\mathcal{X}\{s, i\} \equiv s'=s \wedge i'=i$. As every call's precondition is checked, no array bound violation occurs. The binary search method below adheres to this safe array policy, with the help of the precondition $0 \leq i \wedge j < s$.

```
int⟨r⟩@S bs(Array⟨s⟩@S a, int⟨i⟩@S i, int⟨j⟩@S j, int⟨v⟩@S
   val) where 0≤i∧j<s ; naX{s, i, j}∧(r=−1∨i≤r≤j) {
 if j>i then −1 else { int m=(i+j)/2; int p=sub(a, m);
     if p<val then bs(a, i, m−1, val)
     else if p>val then bs(a, m+1, j, val) else m } }
```

For some problems, e.g. sparse matrix, the indexes may come from another array, called an indirection array. To handle such programs, we can provide:

```
adt ArrayI⟨s, mn, mx⟩ where {mn, mx}; s>0∧mn≤mx {
void⟨⟩@S update(ArrayI⟨s, mn, mx⟩@L a,
   int⟨i⟩@S idx, int⟨v⟩@S val) where 0 ≤ i<s ;
        mx′=max(mx, v)∧mn′=min(mn, v)∧naX{s, i, v}
int⟨r⟩@S sub(ArrayI⟨s, mn, mx⟩@L a, int⟨i⟩@S idx)
   where 0 ≤ i<s ; mn ≤ r ≤ mx∧naX{s, mn, mx, i}; ··· }
```

Apart from `s` which captures the size of the array, we also have `mn` and `mx` to capture the min/max values for the elements in the array. The last two size variables are *mutable* and may be changed by the `update` method, but is guaranteed to satisfy the size invariant $mn \leq mx$. Furthermore, we can bound the result of `sub` method whenever indexes are retrieved from an indirection array.

2.2 User-Defined ADTs

Each user-defined ADT is specified with a set of fields that are visible only to its accompanying methods. For each such declaration, we specify a set of size variables $n_1, \ldots, n_p$ that can be tracked via the constraint ϕ, that is limited to $\bigwedge_{i=1}^{p} n_i = \alpha_i$ whereby $V(\alpha_i) \cap \{n_1, \ldots, n_p\} = \emptyset$. This defines size properties that depend solely on the components of the objects. Another constraint ϕ_I is used to capture the size invariant of each instance of this object type. For example, a user-defined `Bool` object type can be specified in terms of a primitive `bool` field, as follows. Note that this primitive field may be updated.

```
adt Bool⟨n⟩ where n=m ; 0≤n≤1 { bool⟨m⟩@S val;
 void⟨⟩@S Not(Bool⟨m⟩@L v) where true;
     m=0∧m′=1∨m=1∧m′=0 {b.val=not(b.val)}; ··· }
```

A more complex example is the binary tree with a size variable `s` to track the number of nodes, defined as:

```
adt Tree⟨s⟩ where s=1+s1+s2 ; s≥0
{ int val ; Tree⟨s1⟩@U left ; Tree⟨s2⟩@U right; ··· }
```

Alternatively, if we desire AVL trees, we could add a size variable `h` for height and another invariant $-1 \leq h1 - h2 \leq 1$ to enforce a balanced height requirement, as shown below. This invariant can be viewed as a safety policy for our type system to ensure that only proper AVL trees (with this invariant) are constructed.

```
adt AVLTree⟨h, s⟩ where h=1+max(h1, h2)∧s=1+s1+s2
    ; h≥0∧s≥0∧(−1≤h1−h2≤1)
{ int val ; AVLTree⟨h1, s1⟩@U left ;
  AVLTree⟨h2, s2⟩@U right; ··· }
```

Size invariant is safely verified within each ADT, while a method's precondition is verified for each of its callers.

2.3 A Special Construct

Our kernel language includes a special construct, called **bind**, which is used to associate the fields $f_1, \ldots, f_n$ of an object v to a set of local variables $v_1, \ldots, v_n$. This binding is valid over the scope of an expression e and results in local variables being used as synonyms to the corresponding fields. Unlike the usual pattern-matching construct(see [27]), **bind** is based on a pass-by-name (instead of pass-by-value) mechanism. Furthermore, it is used *only* for the purpose of analysis and do not cause runtime overheads. This construct is particularly crucial for tracking the size properties of objects that depend on multiple components. For example, consider a method to insert an element into a binary search tree.

```
void⟨⟩@S insert(Tree⟨s⟩@L t, int⟨c⟩@S i)
    where s>0 ; c′=c
  { if i<t.val then if t.left=null then
                      t.left = new Tree(i, null, null)
                    else insert(t.left, i)
    else if t.right=null then
                      t.right = new Tree(i, null, null)
         else insert(t.right, i) }
```

With the field access construct, we can only capture the size relation between each object and *a single field*. With the **bind** construct, we can capture the size relation between each object and *all its fields*. Hence, we can capture a more precise size mutation for size variables that depend on multiple fields, as illustrated in the following alternative definition of `insert`. The local variables `w, l, r` are given in the same order as the corresponding field declarations.

```
void⟨⟩@S insert(Tree⟨s⟩@L t, int⟨c⟩@S i)
    where s>0 ; c′=c∧s′=s+1
  { bind t of (w, l, r) in
    if i<w then if l=null then l = new Tree(i, null, null)
                else insert(l, i)
    else if r=null then r = new Tree(i, null, null)
         else insert(r, i) }
```

Despite being special, we can automatically insert **bind** into our program as follows. For each field access $v.f$ of a variable v, we identify the largest expression scope that anticipates a field access but does not access v itself (unless it is in shared mode), before introducing a binding for v and its fields $v.f_1, \ldots, v.f_n$ through a fresh set of local variables.

3 Alias Controls for Size Tracking

We adopt four alias controls from [4, 1], namely read-only (`R`), lent (`L`), shared (`S`) and unique (`U`). Each reference in a location (field, parameter or local variable) can be marked with any of the annotations, except that `L`-mode applies to only parameters. While the alias controls are adapted from [1, 4], there are differences in how the type rules are being checked, as described later in Sec 5. Formal definitions of these alias properties are given below.

Definition 1 (Read-Only) *A location that is marked as* **read-only** *(denoted by* `R`*) can be initialized with a reference, but cannot be changed thereafter via the location.*

Definition 2 (Uniqueness) *A reference (from a location) to an object is* **unique** *(denoted by* `U`*) if it is the* **only** *reference to the object at the current scope.*

Definition 3 (Shared) *A reference (from a location) to an object is* **shared** *(denoted by* `S`*) if it may be globally aliased with other reference(s).*

In classical alias annotation work[1], the `L`-mode is used to describe variables whose references do not escape their declaration blocks. Such locations are candidates for lending, as the uniqueness of source objects are not compromised. As a result, each object can be lent out multiple times. The uniqueness property is always restored once all locations that are lent out (including transitive lending) are dead. However, the policy of allowing many lendings for a single object, and of allowing lents to be reassigned to other lent locations, may cause aliasing problems to re-surface. To avoid this problem, we propose a refined policy, called *lent-once*[1]:

Definition 4 (Lent-Once) *A parameter that is marked as* **lent-once** *(denoted by* `L`*) has the following properties: (i) its value does not escape the method, (ii) the lent-once parameter is never re-assigned, (iii) it has exclusive access to the object in its method's scope.*

The last two conditions prevent aliases from occurring in the scope of each lent-once location. As we exclude local variable from the `L`-mode, call sites are the *only* places where lendings can occur. While this lent-once policy may appear restrictive, its shortcoming can be addressed by *preprocessing* (converting lent to lent-once, where possible) and *intersection types*[22] (allowing different parameter aliasing for each call site).

3.1 Alias Subtyping

As we deal with the types of expressions, our subtype relation (shown below) need not include `R`. Each value retrieved from an `R`-field is automatically casted to the `S`-mode. Note that `U`-mode references act as universal sources, while `L`-mode locations serve as sinks only. This subtyping effectively *prevents* an `S`-reference from being copied into

[1] Subsequent to our formulation, we found out that a similar policy, called *limited unique*, has already been proposed in [6], though their focus has been on the specification of such properties and not on a formal mechanism, such as the type-based one proposed here.

a U-location. Unlike [1], we disallow $\mathtt{S} \leq_a \mathtt{L}$ to ensure non-aliasing of the L-mode parameters.

$$A \leq_a A \qquad \mathtt{U} \leq_a \mathtt{L} \qquad \mathtt{U} \leq_a \mathtt{S}$$

3.2 Size Mutability

Our main idea for tracking size properties is the concept of *size-mutability*. To help identify size properties that are immutable for objects, we shall use the *read-only* mode to protect certain fields from being changed. We classify each *field* and *size variable*, as follows:

Definition 5 (Size-Mutability for Field) *A field is classified as* size-immutable*, if it is marked with the* R*-mode. Otherwise, it is classified as* size-mutable*.*

Definition 6 (Size-Mutability for Size Variable) *A* size variable *of an object type is* size-immutable *if it depends on only size-immutable variable(s) that come from size-immutable field(s). Otherwise, it is* size-mutable*.*

For some primitive type, like $\mathtt{int}\langle\mathtt{n}\rangle$, its size variable is automatically classified as *size-immutable*, as its values cannot be changed after they have been created. For user-defined object type, like $\mathtt{Int}\langle\mathtt{n}\rangle$, its field (size property) may change after its objects have been created and is therefore *size-mutable*, in general. Nevertheless, we can obtain more size-immutable fields/variables with the help of the R-mode. Consider:

```
adt Pair⟨a, b, c⟩ where a=p ∧ b=q ∧ c=p+q ; true
  { int⟨p⟩@S x; int⟨q⟩@R y; ··· }
adt List⟨n⟩ where n=m+1 ; n≥0
  { Int⟨v⟩@S val; List⟨m⟩@R next; ··· }
```

In Pair, the first field is size-mutable, while the second field is size-immutable. Of its three size variables $\{\mathtt{a}, \mathtt{b}, \mathtt{c}\}$, only b is size-immutable. This is because b depends on only size-immutable q which came from a read-only field. In List, we use n to track the length of the linked list, and this is made size-immutable by protecting the next field with the R-mode.

Objects may be freely shared without losing track of the properties of their size-immutable variables. However, the same cannot be said for size-mutable variables, as global aliasing can quickly invalidate local analysis. One simple way to tame the aliasing problem, for objects with size-mutable fields to track, is to make their references *unique*. Uniqueness can guarantee that the references of these fields are never aliased.

For example, we can design another list structure whose next field can be updated but must be unique, as shown below. Note that uniqueness allows the length of this linked list to be tracked, even with mutation on its tail.

```
adt List2⟨n⟩ where n=m+1 ; n≥0
  { Int⟨v⟩@S val; List2⟨m⟩@U next; ··· }
```

3.3 Classification of Size Variables

Thus far, we have classified size variables into two main groups, namely size-mutable and size-immutable. However, the size-mutable group can be further classified into either *trackable* or *non-trackable* size variables, due to the *absence* or *presence* of global aliasing, respectively. This classification of size variables can be inferred from the objects' type declarations. We provide a pair of mutual recursive definitions, named $\mathbf{V}_{obj}$ and $\mathbf{V}_{field}$, to classify the size variables of annotated *object type* and *fields/parameters* (with suffixes I, T, and N to denote *immutable, trackable and non-trackable* respectively), as follows:

$$\frac{\mathbf{V}_{obj}(\tau\langle s^*\rangle) = (d_I, d_T, d_N)}{\mathbf{V}_{field}(\tau\langle s^*\rangle@\mathtt{R}) =_{df} (d_I, \emptyset, d_T \cup d_N)}$$

$$\frac{\mathbf{V}_{obj}(\tau\langle s^*\rangle) = (d_I, d_T, d_N)}{\mathbf{V}_{field}(\tau\langle s^*\rangle@\mathtt{S}) =_{df} (\emptyset, d_I, d_T \cup d_N)}$$

$$\frac{A \in \{\mathtt{U}, \mathtt{L}\} \qquad \mathbf{V}_{obj}(\tau\langle s^*\rangle) = (d_I, d_T, d_N)}{\mathbf{V}_{field}(\tau\langle s^*\rangle@A) =_{df} (\emptyset, d_I \cup d_T, d_N)}$$

$$\frac{\mathtt{adt}\ pn\langle n^*\rangle\ \mathtt{where}\ \mathit{MVar};\ \phi_I\ \{\mathit{pmeth}^*\} \in P}{\mathbf{V}_{obj}(pn\langle n^*\rangle) =_{df} (n^* - \mathit{MVar}, \mathit{MVar}, \emptyset)}$$

$$\frac{\begin{array}{c}\mathtt{adt}\ cn\langle n^*\rangle\ \mathtt{where}\ \phi\ ;\ \phi_I\ \{(t_i\ f_i)_{i:1..p}\ \mathit{meth}^*\} \in P \\ \mathbf{V}_{field}(t_i) = (d_i^I, d_i^T, d_i^N), \quad i \in \{1..m\} \\ n_I = \mathit{depends}(n^*, \bigcup_{i=1}^m d_i^I, \phi) \\ n_T = \mathit{depends}(n^*, \bigcup_{i=1}^m d_i^T, \phi) \\ n_N = \mathit{depends}(n^*, \bigcup_{i=1}^m d_i^N, \phi)\end{array}}{\mathbf{V}_{obj}(cn\langle n^*\rangle) =_{df} (n_I - (n_T \cup n_N), n_T - n_N, n_N)}$$

For each field with size variables s^*, we extract from n^* those that depend on some s^* via ϕ. This dependency is defined below. Each primitive constraint from ϕ is of the restricted form $n = \alpha$. We define V to return all free size variables in a formula. For example, $V(\mathtt{x'}=\mathtt{z}+1 \wedge \mathtt{y}=2) = \{\mathtt{x'}, \mathtt{y}, \mathtt{z}\}$. We extend it to annotated type and type environment, e.g. $V(\tau\langle n^*\rangle@A) = \{n^*\}$.

$$\mathit{depends}(n^*, s^*, \phi) =_{df} \bigcup_{(n=\alpha)\in\phi} \{n \mid (V(\alpha) \cap s^*) \neq \emptyset\}$$

Given $\phi_P = (\mathtt{a}=\mathtt{p} \wedge \mathtt{b}=\mathtt{q} \wedge \mathtt{c}=\mathtt{p}+\mathtt{q})$ from the Pair object definition, we can infer that $\{\mathtt{b}\}$ is size-immutable, $\{\mathtt{a}, \mathtt{c}\}$ are trackable size-mutable and $\emptyset$ denotes an empty set of non-trackable size-mutable variables.

$$\begin{array}{ll} n_I = \mathit{depends}([\mathtt{a}, \mathtt{b}, \mathtt{c}], [\mathtt{q}], \phi_P) & = \{\mathtt{b}, \mathtt{c}\} \\ n_T = \mathit{depends}([\mathtt{a}, \mathtt{b}, \mathtt{c}], [\mathtt{p}], \phi_P) & = \{\mathtt{a}, \mathtt{c}\} \\ \mathbf{V}_{obj}(\mathtt{Pair}\langle \mathtt{a}, \mathtt{b}, \mathtt{c}\rangle) & = (\{\mathtt{b}\}, \{\mathtt{a}, \mathtt{c}\}, \emptyset) \end{array}$$

For each object type declaration that is recursive, we may obtain definitions for $\mathbf{V}_{obj}$ and $\mathbf{V}_{field}$ that are circular, but they can still be computed via fixed-point analysis.

4 Specification of Safety Protocols

With the help of size (and alias) annotations, we can specify suitable safety policies for software through the use of *size invariance* for each object type and the *precondition* for each method declaration. Apart from enforcing data structure invariant (e.g. AVL tree) and avoiding bound violations (e.g. safe array), we may also check if operations of an ADT are invoked according to a temporal protocol.

Let us highlight how we can specify the safety policy for a file protocol (defined below) that allows read/write from a file only after it has been opened. The state of each file is encoded by $\mathtt{s}\in 1..3$ with $1, 2, 3$ to denote *initial*, *open*, and *close* states, respectively. (While the states of protocol are encoded using integer values, we stress that it is easy to use syntactic sugar to make protocol specification more natural to programmers.) Once closed, no more read/write operations can be performed. These temporal requirements are enforced by the sized preconditions of each method. In general, not all the safety preconditions can be enforced directly, but we allow programmer to insert runtime tests to ensure conformance. For example, opening a file may fail if it does not exist and this situation can be co-related with the returned flag (of `open`) that can be tested at runtime.

```
adt File⟨s⟩ where {s}; 1≤s≤3 {
  File⟨s⟩@U newFile() where true; s=1
  bool⟨b⟩@S open(File⟨s⟩@L f, string⟨⟩@S name)
       where s=1; (b=1∧s′=2)∨(b=0∧s′=1)
  string⟨⟩@S read(File⟨s⟩@L f) where s=2; s′=2
  void⟨⟩@S write(File⟨s⟩@L f, string⟨⟩@S buf)
      where s=2; s′=2
  void⟨⟩@S close(File⟨s⟩@L f) where s=2; s′=3
  int⟨r⟩@S getState(File⟨s⟩@L f) where true; r=s∧s=s′ }
```

Capturing such a protocol as a primitive ADT relies on the trust we place on primitive methods.

We can also rely on user-defined ADTs for the specification of safety protocols. An example is the following bounded buffer ADT which is guaranteed to *neither underflow nor overflow* each buffer of a given capacity.

```
adt Buffer⟨s, c⟩ where s=n∧c=a; 0≤s≤c∧0≤i<c∧c = e {
  Array⟨a⟩@R arr; int⟨e⟩@R cap; int⟨n⟩@S cur; int⟨i⟩@S i;
  Buffer⟨s, c⟩@U newBuffer(int⟨n⟩@S n)
      where n>0; s=0∧c=n∧n′=n
    { newBuffer(newArray(n, 0), n, 0, 0) }
  int⟨n⟩@S get(Buffer⟨s, c⟩@L b) where s>0; s′=s−1∧c′=c
    { bind b of (a, c, d, i) in {int j=mod(i+c−d, c);
                d=d−1; sub(a, j)} }
  void⟨⟩@S add(Buffer⟨s, c⟩@L b, int⟨n⟩@S n)
      where s<c; s′=s+1∧naX{c, n}
    { bind b of (a, c, d, i) in {update(a, i, n);
                i=mod(i+1, c); d=d+1} } ··· }
```

Note the use of a cyclic array to store the contents of the buffer. This example illustrates how we can safely build a more sophisticated ADT from simpler ones. As an example of the use of this buffer ADT, we have the following method which can safely add `n` elements, provided that the precondition $\mathtt{s+n}\leq\mathtt{c}$ is satisfied.

```
void⟨⟩@S addMany(Buffer⟨s, c⟩@L b, int⟨n⟩@S n,
    int⟨v⟩@S val) where s+n≤c
        ; s′=s+max(0, n)∧naX{c, n, v}
  { if n≤0 then ()
    else add(b, val); addMany(b, n−1, val) }
```

5 Verification via Sized Typing

We present type judgements for *expressions*, *method declarations*, *object declarations* and *programs* to check for their well-typedness, using relations of the form:

$$\Gamma; C; \Delta; \Theta \vdash e :: t, \Delta_1, \Theta_1 \qquad C \vdash_{meth} meth \qquad \vdash_{pmeth} pmeth$$
$$\vdash_{prim} prim \qquad \vdash_{user} user \qquad \vdash_{P} prim^* \; user^* \; meth^*$$

Note that Γ is a type environment mapping program variables to their annotated types; $\Delta(\Delta_1)$ denotes the size constraint that holds for the size variables associated with Γ (Γ and t) for expression e before (after) its evaluation; t is an annotated type. C captures a set of user-defined ADTs whose fields are visible. Also, $\Theta(\Theta_1)$ is used to hold the variables whose uniqueness have been consumed before (after) the evaluation of e. We shall refer to such references as *dead*, as they are not to be accessed. Take note that both $\Delta(\Delta_1)$ and $\Theta(\Theta_1)$ are tracked flow-sensitively.

The syntax-directed type rules for these judgements are given in Fig 2. When compared to [1], we achieved two key innovations. Firstly, we do not require a separate *last-use analysis* to determine when a unique reference (from a variable) becomes dead, as we keep track of consumed uniqueness via the dead set (denoted by Θ). Secondly, we do not require *destructive read* for field access (which nullifies the field after each read). Instead, we track the liveness of each unique field with the help of the `bind` construct and the dead set. These changes improve both the accuracy and usability of our alias analysis.

In the rest of this section, we highlight important aspects of our sized type system via examples. Before that, we introduce some preliminary definitions. The function *prime* takes a set of size variables and returns their primed version, e.g. $prime(\{s_1, \ldots, s_n\}) = \{s'_1, \ldots, s'_n\}$. Note that prime operation is idempotent, namely $v'' = v'$. We extend this to annotated type, type environment and even substitution. For example, $prime(\tau\langle n_1, \ldots, n_k\rangle) = \tau\langle n'_1, \ldots, n'_k\rangle$, and $prime\,[x \mapsto a, y \mapsto b] = [x' \mapsto a', y' \mapsto b']$. We use

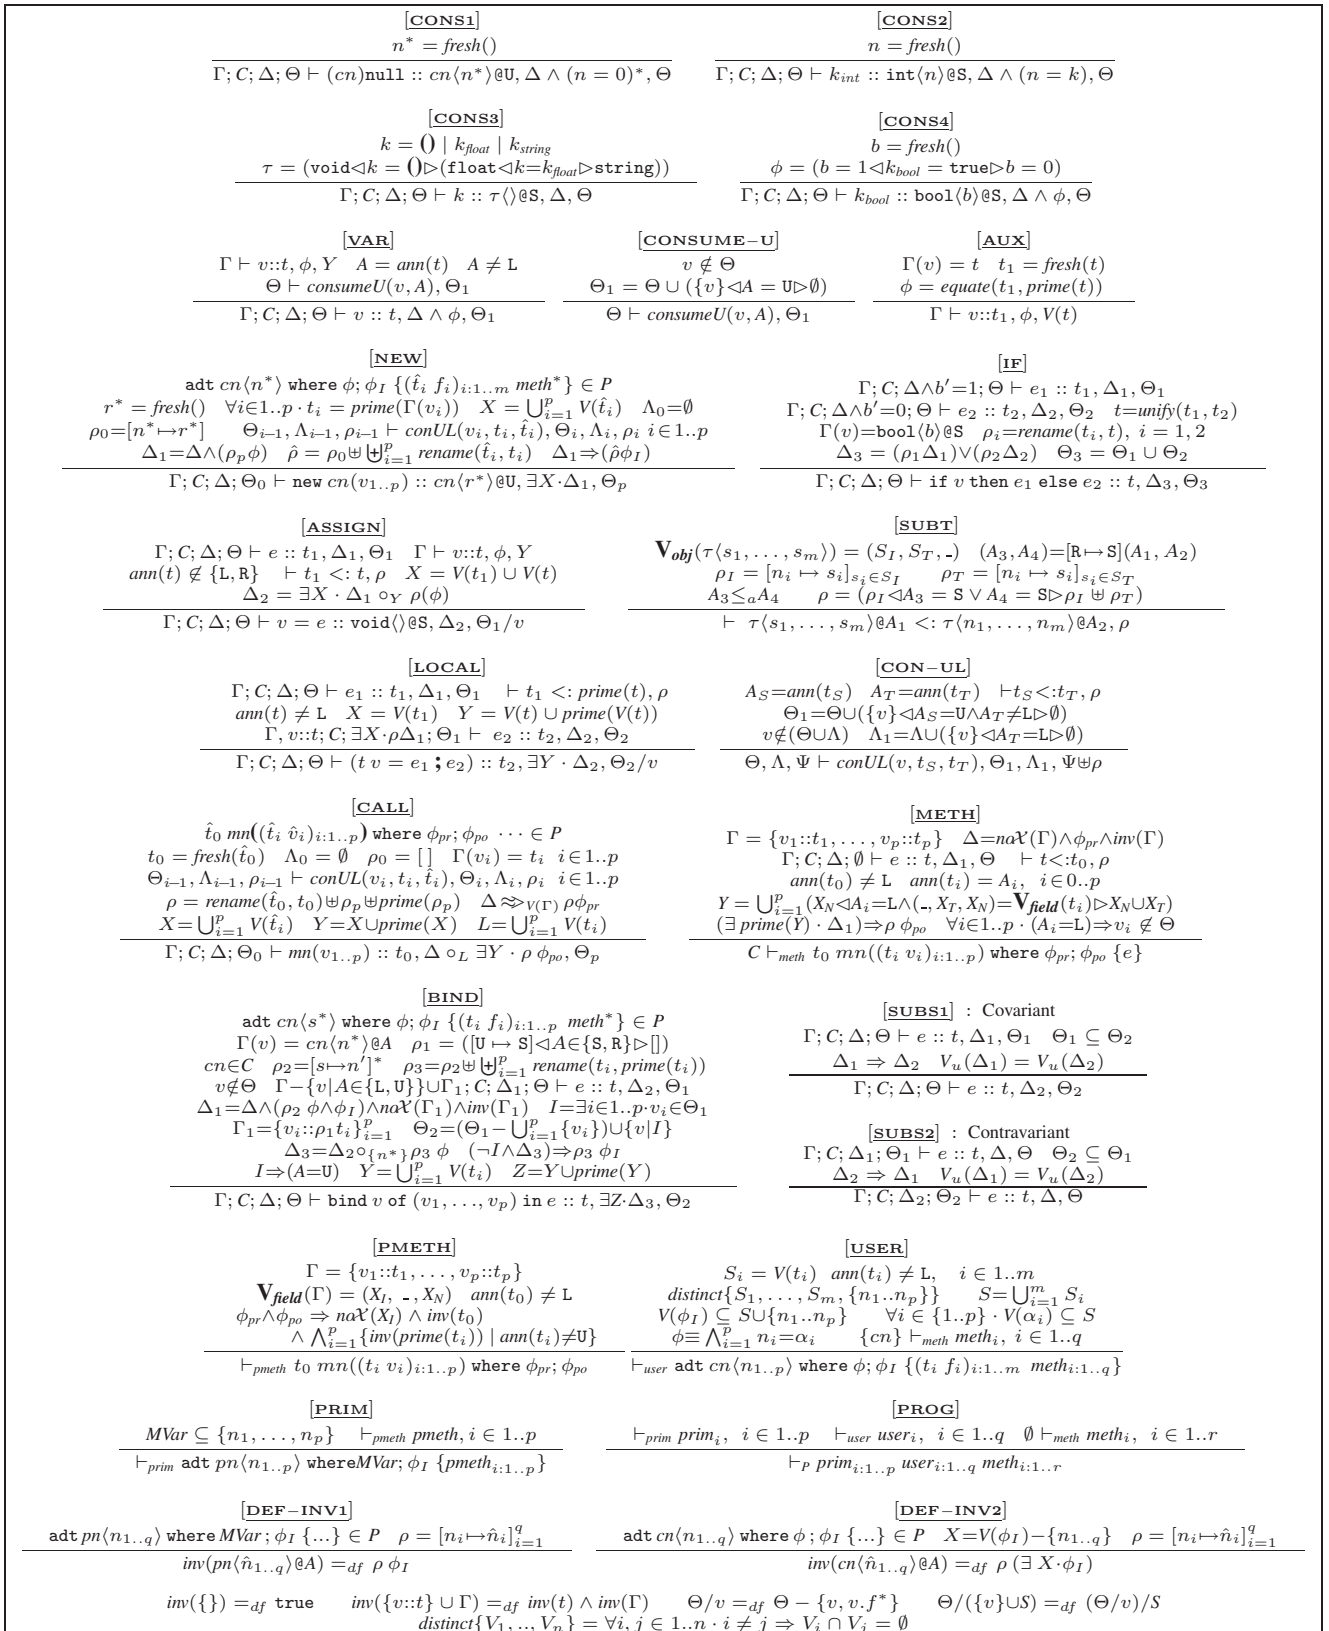

Figure 2. Size and Alias Type Rules

$n^* = \mathit{fresh}()$ to generate new size variables n^*. We extend it to annotated type, so that $\hat{t} = \mathit{fresh}(t)$ will return a new type $\hat{t}$ with the same underlying type as t but with fresh size variables instead. The function $\mathit{equate}(t_1, t_2)$ generates equality constraints for the corresponding size variables of its two arguments, assuming that both arguments share the same underlying type. For example, we have $\mathit{equate}(\texttt{Int}\langle r\rangle, \texttt{Int}\langle s'\rangle) = (r = s')$. To extract the alias of an annotated type, we use $\mathit{ann}(\tau\langle v^*\rangle\texttt{@}A) = A$. The function $\mathit{inv}(t)$ returns the size invariant of t (after quantifying the size variables of its fields), while $\mathit{inv}(\Gamma)$ extends to type environment. The function $\mathit{unify}(t_1, t_2)$ returns a new fresh type whose alias annotation is an upper bound of t_1 and t_2. Also, conditional is expressed as $\xi_1 \lhd b \rhd \xi_2 =_{df}$ *if* b *then* ξ_1 *else* ξ_2.

Updates of size constraint due to change of program state are effected by a sequential composition operator, $\circ_X$, with X denoting the set of size variables that are being updated. For example, if $\Delta{=}(a'{=}a{+}b \wedge b'{=}3 \wedge i{=}b'{+}a')$, then:

$$\begin{array}{rl}(\Delta \circ_{\{a\}} i{=}a') & = \exists a_0 \cdot a_0{=}a{+}b \wedge b'{=}3 \wedge i{=}b'{+}a_0 \wedge i{=}a' \\ & = b'{=}3 \wedge i{=}b'{+}a{+}b \wedge i{=}a'\end{array}$$

More formally, sequential composition is defined as:

$$\phi_1 \circ_X \phi_2 =_{df} \exists R \cdot \rho_1(\phi_1) \wedge \rho_2(\phi_2)$$

where $X = \{s_1, \ldots, s_n\}$ are size variables being updated
$R = \{r_1, \ldots, r_n\}$ are fresh size variables
$\rho_1 = \{s'_i \mapsto r_i\}_{i=1}^n \qquad \rho_2 = \{s_i \mapsto r_i\}_{i=1}^n$

5.1 Variable Read

The [VAR] rule is used to return an object via a reference captured by variable v. Here, a unique reference loses its uniqueness and must not be referenced again (unless it has been re-assigned). It is thus added to a *dead set* using $\Theta_1{=}\Theta \cup (\{v\} \lhd A{=}\texttt{U} \rhd \emptyset)$ where A is the alias annotation for the type of v. This rule also returns a freshly annotated type $t_1{=}\mathit{fresh}(t)$ and a size constraint $\phi{=}\mathit{equate}(t_1, \mathit{prime}(t))$ to relate the resulting annotated type with that of variable v. Assuming $\Gamma{=}\{\texttt{v::Int}\langle\texttt{d}\rangle\texttt{@U}\}$, $\Delta{=}n\alpha\mathcal{X}(\{\texttt{d}\})$, we have:

$$\frac{\phi = (\texttt{r}_1{=}\texttt{d}') \qquad \texttt{v} \notin \Theta \qquad \texttt{U} \neq \texttt{L}}{\Gamma; C; \Delta; \Theta \vdash \texttt{v} :: \texttt{Int}\langle\texttt{r}_1\rangle\texttt{@U}, \Delta{\wedge}\phi,\ \Theta \cup \{\texttt{v}\}}$$

Note how the dead set has changed, and how the size property of the result (namely $\texttt{r}_1$) is associated with the latest size variables (namely $\texttt{d}'$) from the type environment. L-mode references cannot escape via this rule.

5.2 Conditional

For conditional, our rule is able to track the size constraint path-sensitively, but the dead set is tracked flow-sensitively. Given $\Gamma_1 = \Gamma + \{\texttt{x} :: \texttt{bool}\langle\texttt{h}\rangle\texttt{@S}\}$ and $\Delta = (\texttt{e}'{=}0)$, we have:

$$\frac{\begin{array}{c}\Gamma_1; C; \Delta{\wedge}\texttt{h}'{=}1; \{\} \vdash \texttt{v::Int}\langle\texttt{r}_1\rangle\texttt{@U}, \texttt{false}, \{\texttt{v}\} \\ \Gamma_1; C; \Delta{\wedge}\texttt{h}'{=}0; \{\} \vdash \texttt{new Int(5)::Int}\langle\texttt{r}_2\rangle\texttt{@U}, \Delta{\wedge}\texttt{r}_2{=}5, \{\}\end{array}}{\begin{array}{r}\Gamma_1; C; \Delta; \{\} \vdash \texttt{if x then v else (new Int(5))} \\ :: \texttt{Int}\langle\texttt{r}_3\rangle\texttt{@U}, \Delta{\wedge}\texttt{r}_3{=}5,\ \{\texttt{v}\}\end{array}}$$

Note that $\texttt{h}'{=}1$ and $\texttt{h}'{=}0$ captures path-sensitivity, but the uniqueness of v is consumed even though it occurs in an infeasible branch. However, the final size constraint is able to ignore the infeasible branch. (A branch is infeasible if its corresponding size constraint yields false.)

5.3 Method Declaration and Call

For method declaration, our type rule must ensure that each decoration at the method header is consistent with its body. The declared size constraint may be weaker than the analysed size constraint (from the method body), but not vice versa. Note that non-trackable size variables and mutable size variables of unique parameters must be quantified, as they cannot be tracked across method boundary

The rule for method call must ensure that U-mode parameters consume uniqueness, and that L-mode parameters adhere to the lent-once policy. The relation $\Theta, \Lambda, \Psi \vdash \mathit{conUL}(v, t_S, t_T), \Theta_1, \Lambda_1, \Psi_1$ (named [CON–UL]) helps ensure the above. It checks that each variable v is neither lent twice, nor has its uniqueness consumed twice. Note that Λ captures unique variables which are temporarily lent out. In addition, fields of such parameters must not be in the dead set. The method invocation rule also includes a safety check on the precondition of the callee. This is performed by the predicate $\approx\!\gg_X$ defined as follows:

$\Delta \approx\!\gg_X \phi =_{df} (\Delta \Rightarrow \rho\phi)$, where
$\rho = [s_1 \mapsto s'_1, .., s_n \mapsto s'_n]$ and $V_u(\phi) \cap X = \{s_1, .., s_n\}$.

where V_u returns the size variables in unprimed form, e.g. $V_u(x' = z + 1 \wedge y = 2) = \{x, y, z\}$. As an example of checking a method call, given

$$\begin{array}{l}\Gamma = \{\texttt{p} :: \texttt{Buffer}\langle\texttt{r}, \texttt{z}\rangle\texttt{@U}, \texttt{v} :: \texttt{int}\langle\texttt{m}\rangle\texttt{@S}\} \\ \Delta = (\texttt{r}'{=}3 \wedge \texttt{r}{=}0 \wedge \texttt{z}{=}10 \wedge \texttt{m}{=}30 \wedge n\alpha\mathcal{X}(\{\texttt{z}, \texttt{m}\}))\end{array}$$

we obtain $V(\Gamma) = \{\texttt{r}, \texttt{z}, \texttt{m}\}$ and the following. Note that uniqueness of p is only lent-out and not consumed.

$$\frac{\begin{array}{c}t_1 = \texttt{Buffer}\langle\texttt{r}, \texttt{z}\rangle\texttt{@U} \quad \hat{t}_1 = \texttt{Buffer}\langle\texttt{s}, \texttt{c}\rangle\texttt{@L} \\ \texttt{p} \notin \Theta \quad \rho_1 = [\texttt{s}{\mapsto}\texttt{r}, \texttt{c}{\mapsto}\texttt{z}] \\ \hline \Theta, \emptyset, [\,] \vdash \mathit{conUL}(\texttt{p}, t_1, \hat{t}_1), \Theta, \{\texttt{p}\}, \rho_1 \\ t_2 = \texttt{Int}\langle\texttt{m}\rangle\texttt{@U} \qquad \hat{t}_2 = \texttt{Int}\langle\texttt{n}\rangle\texttt{@U} \\ \texttt{v} \notin \Theta \cup \{\texttt{p}\} \quad \rho_2 = \rho_1 \uplus [\texttt{n}{\mapsto}\texttt{m}] \\ \hline \Theta, \{\texttt{p}\}, \rho_1 \vdash \mathit{conUL}(\texttt{v}, t_2, \hat{t}_2), \Theta, \{\texttt{p}\}, \rho_2 \\ \hat{t}_2 = \texttt{Int}\langle\texttt{n}\rangle\texttt{@U} \qquad \Delta \approx\!\gg_{V(\Gamma)} (\texttt{r} < \texttt{s})\end{array}}{\Gamma; C; \Delta; \Theta \vdash \texttt{add(p, v)} :: \texttt{void}\langle\rangle\texttt{@S}, \Delta_2,\ \Theta}$$

where $\Delta_2 = \Delta \circ_{\{\texttt{r},\texttt{z},\texttt{m}\}} (\texttt{r}' = \texttt{r} + 1 \wedge n\alpha\mathcal{X}(\{\texttt{z}, \texttt{m}\}))$
$= (\texttt{r}'{=}4 \wedge \texttt{r}{=}0 \wedge \texttt{z}{=}10 \wedge \texttt{m}{=}30 \wedge n\alpha\mathcal{X}(\{\texttt{z}, \texttt{m}\}))$

5.4 Binding

The **bind** expression enables us to accurately capture the size relations among object fields as well as between an object and its fields. In addition, it can replace the field access construct altogether (for type checking). When an object subjected to a **bind** operation is marked by either R or S, all its unique fields are remarked as S using [U↦S] in the scope of **bind**. On the other hand, the object will lose its uniqueness if any one of its fields loses it. For instance, given Γ=Γ_0+{v::List2⟨n⟩@U}, the following derivation demonstrates how v loses its uniqueness when its unique y field loses it (denoted by $I = (\mathtt{y} \in \Theta_1)$).

$$\frac{\mathtt{v} \notin \Theta \quad \Gamma_0 \cup \Gamma_1; C; \Delta_1; \Theta \vdash \mathtt{y} :: t_1, \Delta_2, \Theta_1 \quad \Theta_1 = \Theta \cup \{\mathtt{y}\} \quad \Gamma(\mathtt{v}){=}t \quad A{=}ann(t) \quad I{=}(\mathtt{y} \in \Theta_1) \quad I \Rightarrow A{=}\mathtt{U} \quad \neg I \wedge \Delta_3 \Rightarrow \mathtt{n}' \geq \mathtt{0}}{\Gamma; C; \Delta; \Theta \vdash \mathtt{bind\ v\ of\ (x,y)\ in\ y} :: t_1, \Delta_4, \Theta_1 - \{\mathtt{y}\} \cup \{\mathtt{v}|I\}}$$

for some values Γ_1, t_1, Δ_1,Δ_2, Δ_3 and Δ_4. As a counter-example, consider `bind v of (x,y) in x`. Here, v does not lose its uniqueness as the uniqueness of its y field remains intact. A well-typed **bind** expression also permits the size invariant ϕ_I of the bound object to be *temporarily violated* in its scope, but ensures that ϕ_I holds whenever the object escapes.

6 Dynamic Semantics

The dynamic operational semantics is described in small steps. Notations used are defined as follows.

Locations :	ι	$\in$	*Location*
Primitives :	k	$\in$	$prim = \mathtt{int} \uplus \mathtt{bool} \uplus \mathtt{float} \uplus \mathtt{null} \uplus \mathtt{void}$
Values :	δ	$\in$	$Value = prim \uplus Location$
AnnVal :	ν	$\in$	$AVal = (A^+ \times Value)$
Aliases :	A	$\in$	$A^+ = A \uplus \mathtt{U}_D$
Store :	ϖ	$\in$	$Store = Location \rightarrow_{fin} ObjVal$
Variable Env :	Π	$\in$	$VEnv = Var \rightarrow_{sfin} AVal$
Object values :	η	$\in$	$ObjVal = Type \times (Fd \rightarrow_{fin} AVal)$
Type :	$\tau\langle n^* \rangle$	$\in$	$Type = CN \times SV^*$

Note that $f : A \rightarrow_{sfin} B$ denotes a *stackable* mapping from A to B. Such a stackable mapping is defined as:

$$A \rightarrow_{sfin} B =_{df} (A^{\mathtt{int}} \rightarrow_{fin} B,\ \mathtt{int})$$

where `int` denotes the current frame number of the mapping, while $A^{\mathtt{int}}$ denotes the domain that has been marked with frame numbers. Each time a new frame is created, the current frame number is increased by 1. Frame numbers are useful in making all variables distinct in the stack without having to resort to variable renaming. They are also useful for formulating the lent-once property (see Sec 7) for stack frames. Main operators are:

Start a new frame :
$newframe(f, n) = (f, n{+}1)$
Push variables into current frame :
$(f, n) + \{(v \mapsto \nu)^*\} = (f + \{(v^n \mapsto \nu)^*\},\ n)$
Pop variables from the topmost frame :
$(f, n) - \{v^*\} =$
$\mathtt{let}\ \hat{f} = f \backslash \{(v^n)^*\}\ \mathtt{in}\ (\hat{f},\ max\ \{m \mid \exists u \cdot u^m \in dom(\hat{f})\})$

The variable environment Π is such a stackable mapping. We write $\Pi[\nu/v]$ to denote an update of the value of the latest variable v in Π to ν. We write $\Pi + \{v \mapsto \nu\}$ to denote an extension of Π to include a binding of ν to v, while $\Pi - \{v^*\}$ removes a subset of the mappings. Similar notations are used for the update and enhancement of object values and stores. In the case of store, we also provide an abbreviated notation $\varpi[\nu/\iota.f] =_{df} \mathtt{let}\ (t, \rho_f) = \varpi(\iota)\ \mathtt{in}\ \varpi[(t, \rho_f[\nu/f])/\iota]$. Given an object value, $\eta = (t, \rho_f)$, we have $Flds(\eta) =_{df} \rho_f$ and $type(\eta) =_{df} t$.

We maintain alias annotations for variables in the stack and for fields in the store at run-time. We enhance the current set of alias annotations with a new value: $\mathtt{U}_D$. A variable/field can be assigned an alias $\mathtt{U}_D$ if its original alias annotation was U, and it has since relinquished this uniqueness ownership (of its reference). By maintaining this annotation at run-time (corresponding to dead set for static semantics), our system (both dynamic and static) is able to identify all unique references which have been consumed via read operations. Note that the run-time types for objects also carry size variables. Both these instruments (alias and size) are used only for proving the soundness of the type system. They can be erased without affecting the underlying semantics.

We require some intermediate expressions for the dynamic semantics to follow through. Our syntax is thus extended from the original expression syntax as follows:

$$e ::= \cdots \mid \iota \mid \nu \mid \mathtt{ret}(v^*, \rho, e) \mid \mathtt{retb}(v^*, A, \iota, e)$$

The expression $\mathtt{ret}(v^*, \rho, e)$ is used to capture the result of evaluating a local block, or the result of a method invocation. The list of variables associated with `ret` is the local variables declared and used by the block. This set of variables is popped from the stack at the end of the block's evaluation. We also supply ρ to capture the mapping of formal parameters to actual arguments for method call. This is used by the soundness proof. The expression $\mathtt{retb}(v_{0..n}, A, \iota, e)$ is used to capture the result of evaluating a bind construct $\mathtt{bind}\ v_0\ \mathtt{of}\ (v_1, .., v_n)\ \mathtt{in}\ e$.

We extend the static semantics of the language to include these new expressions. In the process, we require introduction of a *store typing* to describe the type of each location. This ensures that objects created in the store during run-time

are type-wise consistent with that captured by the static semantics. Store typing is conventionally used to link static and dynamic semantics. In our case, it is denoted by:

$$\Sigma \in \mathit{StoreType} = \mathit{Location} \rightarrow_{\mathit{fin}} \mathit{Type}$$

Judgements in the static semantics will be extended with store typing, as follows:

$$\Gamma; C; \Sigma; \Delta; \Theta \vdash e :: t, \Delta_1, \Theta_1$$

The static semantics for the new intermediate expressions are shown in Fig 3.

$$\frac{v^* \subseteq \mathit{dom}(\Gamma) \quad \Gamma; C; \Sigma; \Delta; \Theta \vdash e :: t, \Delta_1, \Theta_1 \quad v^* = \{v_1, \ldots, v_p\} \quad X = V(\Gamma(v^*)) \quad Y = X \cup \mathit{prime}(X) \quad \rho_2 = \biguplus \{\hat{\rho}_i \mid t_i = \Gamma(v_i),\ \hat{t}_i = \Gamma(\rho\, v_i),\ \vdash \hat{t}_i <: t_i, \hat{\rho}_i\}_{i=1}^{p}}{\Gamma; C; \Sigma; \Delta; \Theta \vdash \mathtt{ret}(v^*, \rho, e) :: t, \exists Y \cdot (\mathit{prime}(\rho_2)\Delta_1), \Theta_1/\{v^*\}} \ [\textbf{ELF}]$$

$$\frac{v_{1..n} \subseteq \mathit{dom}(\Gamma) \quad v \notin \mathit{dom}(\Gamma) \lhd A \in \{\mathtt{U}, \mathtt{L}\} \rhd v \in \mathit{dom}(\Gamma) \quad \Gamma; C; \Sigma; \Delta; \Theta \vdash e :: t, \Delta_1, \Theta_1 \quad X = V(\Gamma(v^*)) \quad \Sigma(\iota) = \mathit{cn}\langle m^* \rangle @ A_1 \quad \mathtt{adt}\ \mathit{cn}\langle s^* \rangle\ \mathtt{where}\ \phi; \phi_I\ \{\ldots\} \in P \quad \rho = [s \mapsto m]^* \quad \Delta_2 = \Delta_1 \circ_{\{m^*\}} \rho\, \phi \quad Y = X \cup \mathit{prime}(X) \quad v \notin \Theta_1 \quad I = \exists i \in 1..n \cdot v_i \in \Theta_1 \quad \Theta_2 = (\Theta_1 - \bigcup_{i=1}^{n} \{v_i\}) \cup \{v \mid I\}}{\Gamma; C; \Sigma; \Delta; \Theta \vdash \mathtt{retb}(v, v_{1..n}, A, \iota, e) :: t, \exists Y \cdot \Delta_2, \Theta_2} \ [\textbf{RETB}]$$

$$\frac{}{\Gamma; C; \Sigma; \Delta; \Theta \vdash \iota :: \Sigma(\iota) @ \mathtt{U}, \Delta, \Theta} \ [\textbf{LOC}]$$

$$\frac{\Gamma; C; \Sigma; \Delta; \Theta \vdash \delta :: t, \Delta_1, \Theta_1 \quad A_t = \mathit{ann}(t) \quad A_t \leq_a A}{\Gamma; C; \Sigma; \Delta; \Theta \vdash (A, \delta) :: [A_t \mapsto A]t, \Delta_1, \Theta_1} \ [\textbf{A-DATA}]$$

Figure 3. Type Rules for Intermediates

The dynamic evaluation rules are of the following form.

$$\langle \Pi, \varpi \rangle\ [e] \hookrightarrow \langle \Pi', \varpi' \rangle\ [e']$$

We shall formulate the rules using an exception-style semantics with three possible errors, namely **E** = **Error-Alias** | **Error-Prim** | **Error-Null**. Whenever one such error is raised, the evaluation aborts. This error occurrence can be stated using $\langle \Pi, \varpi \rangle\ [e] \hookrightarrow \mathbf{E}$.

The small-step dynamic semantics is formalised in Fig 4, together with some auxiliary functions. The function $\mathit{read}(\Pi, v)$ reads a value from the stack. It performs alias and null checks to determine if errors should be raised. The function $\mathit{upd}(\Pi, v, \nu)$ performs a corresponding write operation to the stack. It also performs appropriate alias and null checks. The function $\mathit{updVE}(\Pi, v_t, A_t, v)$ performs alias checks for each argument (v) that is passed to corresponding formal parameter (v_t). The uniqueness of the argument is consumed, if needed. The function *callPrim* call the primitive method mn. Take note that the actual primitive operation *primOp* includes checks which ensure that the size-related pre-condition of primitive method (denoted by ϕ_{pr}) is satisfied, and that the necessary changes have been made to store, yielding ϖ'. If the primitive's precondition is not satisfied, *primOp* sets a return flag to `false`. Correspondingly, the reduction will get stuck and the exception **Error-Prim** is raised.

7 Soundness of Type System

The soundness of our static semantics relies on the following consistency relationship between the static and dynamic semantics.

Definition 7 (Consistency Relation) *The consistency relation between the static and dynamic environments is defined as follows:*

$$\frac{\mathit{dom}(\Pi) = \mathit{dom}(\Gamma) \quad \mathit{dom}(\varpi) = \mathit{dom}(\Sigma) \quad \mathcal{C}_{\mathbf{size}}(\Gamma, \Sigma, \Delta, \Theta, \Pi, \varpi) \quad \mathcal{C}_{\mathbf{alias}}(\Gamma, \Sigma, \Theta, \Pi, \varpi)}{\Gamma; \Sigma; \Delta; \Theta \models \langle \Pi, \varpi \rangle}$$

The relations $\mathcal{C}_{\mathbf{alias}}$ and $\mathcal{C}_{\mathbf{size}}$ denote *alias consistency* and *size consistency* respectively. They are defined as follows:

$$\frac{\forall v \in \mathit{dom}(\Gamma) \cdot (v \notin \Theta \Rightarrow \mathit{ann}(\Pi(v)) \leq_a \mathit{ann}(\Gamma(v))) \quad \forall \iota \in \mathit{dom}(\varpi) \cdot \rho = \mathit{Flds}(\varpi(\iota)) \wedge \forall f \in \mathit{dom}(\rho) \cdot (\mathit{live}(\iota, f, \Pi, \Theta) \Rightarrow \mathit{ann}(\rho(f)) \leq_a \mathit{ann\text{-}fd}(\Sigma(\iota), f))}{\mathcal{C}_{\mathbf{alias}}(\Gamma, \Sigma, \Theta, \Pi, \varpi)}$$

Here, $\mathit{live}(\iota, f, \Pi, \Theta) =_{df} \neg(\exists v \in \mathit{dom}(\Pi) \cdot \Pi(v) = (_, \iota) \wedge v.f \in \Theta)$. We define $\mathit{ann\text{-}fd}(\mathit{cn}\langle u^* \rangle, f) =_{df} A$, where $\mathtt{adt}\ \mathit{cn}\langle n^* \rangle\ \mathtt{where}\ \ldots\{..\tau\langle m^* \rangle @ A\ f..\} \in P$.

$$\frac{D_\Theta = \mathit{prime}(\bigcup_{w \in \Theta} \{Z \mid \Gamma \vdash w :: t, \phi, Z\}) \quad \forall v \in \mathit{dom}(\Gamma) \cdot \Delta_v = \mathit{size}(\Gamma(v), \Pi(v), \varpi) \quad X = (V(\Delta) - \mathit{prime}(V(\Gamma))) \cup D_\Theta \quad \Delta_\Pi = \wedge_{v \in \mathit{dom}(\Gamma)} \Delta_v \quad \Delta_\Pi \Rightarrow \exists X \cdot \Delta}{\mathcal{C}_{\mathbf{size}}(\Gamma, \Sigma, \Delta, \Theta, \Pi, \varpi)}$$

In the above, *size* computes the size of a run-time value, and presents it in constraint form. It is defined as follows:

$$\begin{array}{l} \mathit{size}(\mathtt{int}\langle r \rangle, (A, k), \varpi) =_{df} (r' = k) \\ \mathit{size}(\mathtt{bool}\langle r \rangle, (A, k), \varpi) =_{df} (r' = (1 \lhd k = \mathtt{true} \rhd 0)) \\ \mathit{size}(\mathit{cn}\langle r^* \rangle, (A, \iota), \varpi) =_{df} \mathit{let}\ \ (\mathit{cn}\langle n_{1..p} \rangle, \rho_f) = \varpi(\iota) \\ \qquad\qquad\qquad\qquad\qquad\qquad \rho = \mathit{getSize}(\iota, \varpi) \\ \qquad\qquad\qquad\qquad\qquad\quad \mathit{in}\ \wedge_{i=1}^{p} \{r'_i = (\rho\, n_i)\} \\ \mathit{size}(t, (A, _), \varpi) =_{df} \mathit{True} \end{array}$$

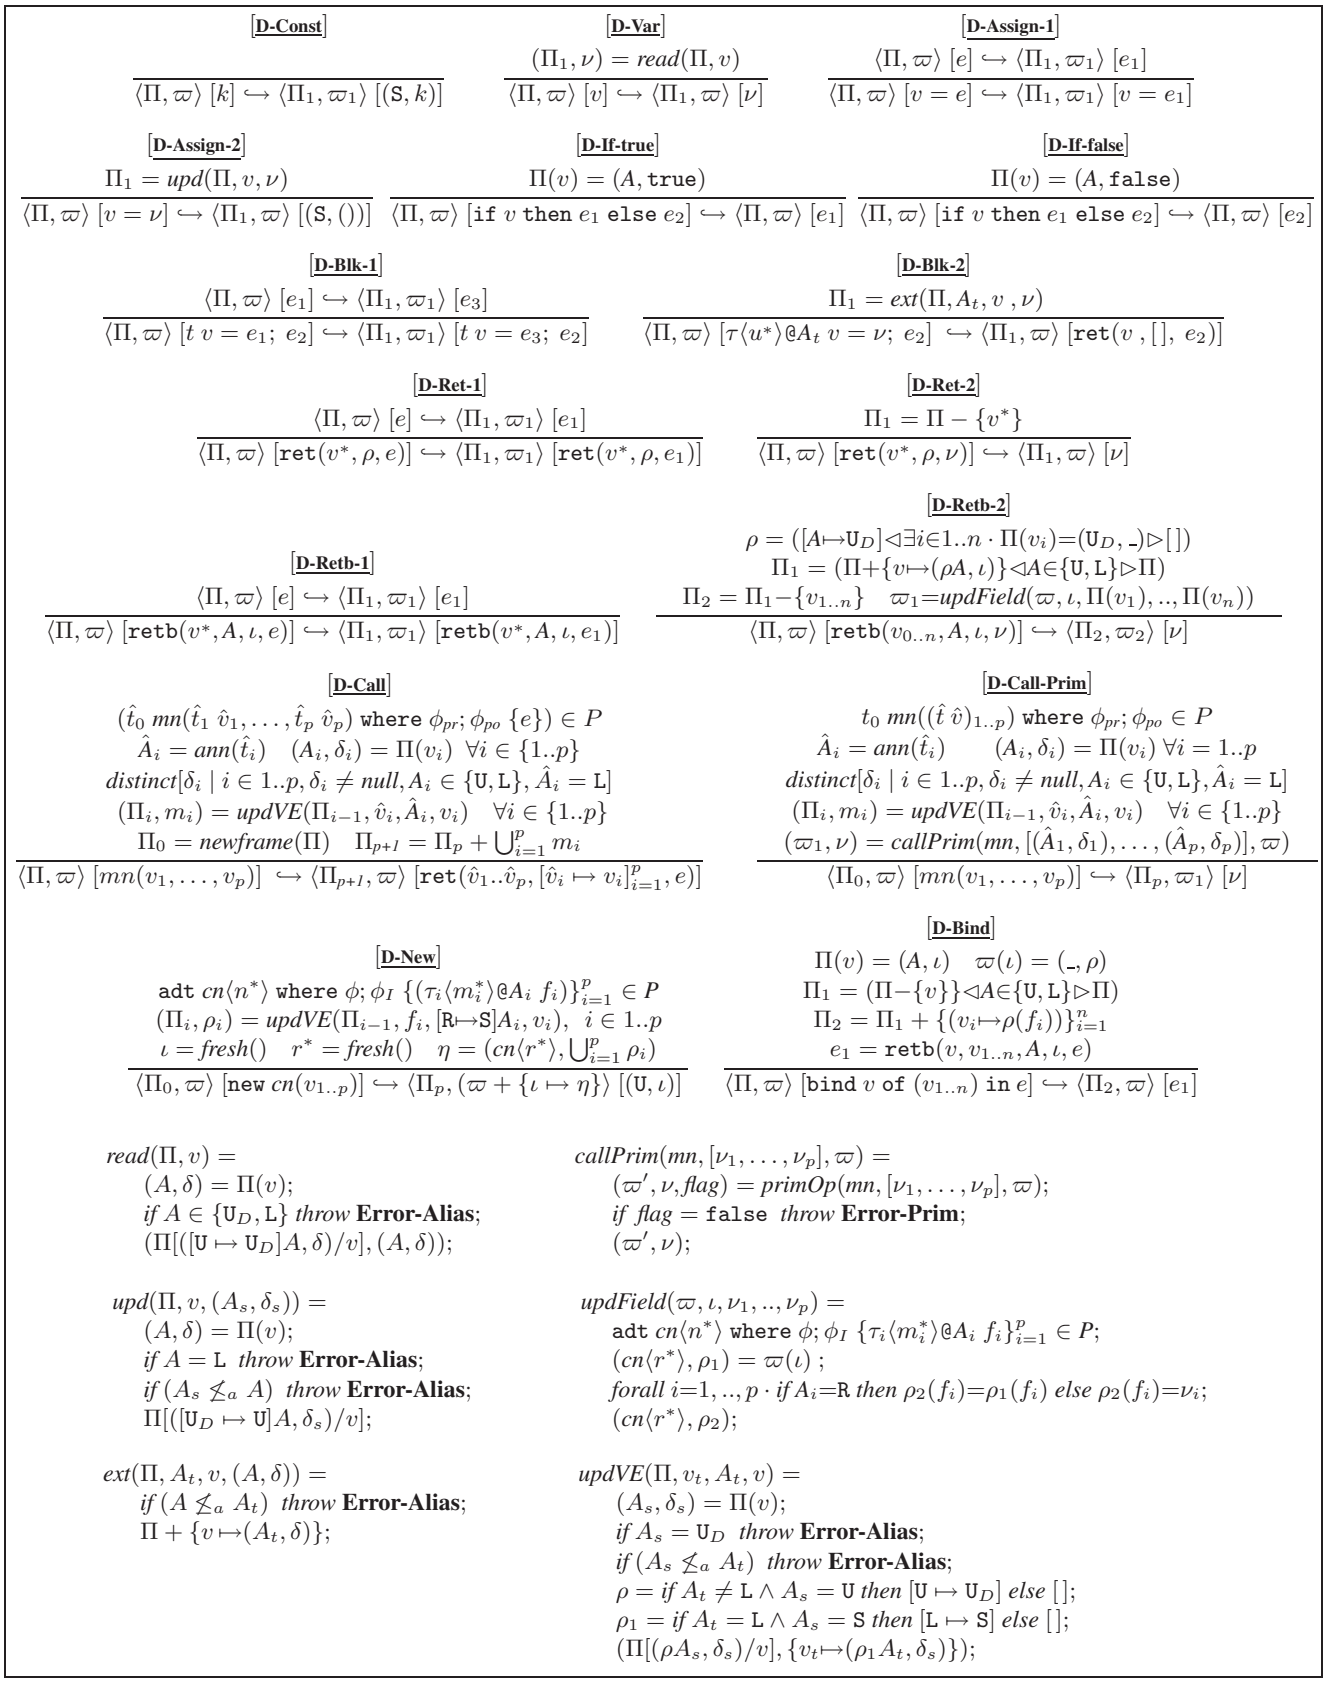

Figure 4. Dynamic Semantics

$$
\begin{array}{l}
\mathit{getSize} \;::\; (\mathit{Locations} \times \mathit{Store}) \rightarrow (\mathit{SV} \rightarrow_{\mathit{fin}} \mathtt{int}) \\
\mathit{getSize}(\iota, \varpi) = \mathit{getSize1}(\iota, \varpi, [\iota]) \\
\mathit{getSize1}(\iota, \varpi, \mathit{activeLst}) = \\
\quad \mathtt{adt}\ cn\langle n_{1..q}\rangle\ \mathtt{where}\ \phi; \phi_I\ \{\tau_i\langle m_i^*\rangle @ \hat{A}_i\ f_i\}_{i=1}^{p} \in P \\
\quad (cn\langle r_{1..q}\rangle,\ \{f_i \mapsto (A_i, \delta_i)\}_{i=1}^{p}) = \varpi(\iota); \\
\quad \rho = []; \\
\quad \rho_\iota = \{n_i \mapsto r_i\}_{i=1..q}; \\
\quad \mathit{forall}\ i = 1, .., p \quad \{ \\
\quad\quad \mathit{if}\ (\tau_i = \mathtt{int}) \quad \rho = \{m_i \mapsto \delta_i\} \uplus \rho; \\
\quad\quad \mathit{if}\ (\tau_i = \mathtt{bool}) \quad \rho = \{m_i \mapsto 1 \lhd \delta_i = \mathtt{true} \rhd 0\} \uplus \rho; \\
\quad\quad \mathit{if}\ (\tau_i = \mathtt{float} \mid \mathtt{void}) \quad \rho = \{m_i \mapsto 0\} \uplus \rho; \\
\quad\quad \mathit{if}\ (\mathit{member}\ (\delta_i, \mathit{activeLst})) \quad \rho = \{m_i \mapsto 0\} \uplus \rho \\
\quad\quad \mathit{else}\ \{\ \ \rho_1 = \mathit{getSize1}(\delta_i, \varpi, \mathit{add}(\delta_i, \mathit{activeLst})); \\
\quad\quad\quad\quad \rho = \rho_1 \uplus \rho\} \\
\quad \} \\
\quad \uplus\{(\rho_\iota n_i) \mapsto (\rho\ \alpha_i) \mid \phi = \wedge_{i=1..q}(n_i = \alpha_i)\}
\end{array}
$$

In what follows we shall state several desired properties that well-typed expressions possess. Before presenting the first theorem with respect to alias properties, we need several formal definitions, which are given as follows.

Definition 8 (Bag of References) *Function RefB applied to* Π,ϖ, *or* e *returns a bag of references in triple form:*

$$
\begin{array}{l}
\mathit{RefB}(\Pi) \;=_{df}\; \{(v^s, A, \delta) | (v^s, (A, \delta)) \in \Pi, \delta \in \mathit{dom}(\varpi)\}_b \\
\mathit{RefB}(\varpi) \;=_{df} \\
\quad \{(\iota.f, A, \delta) | (\iota, \eta) \in \varpi, (f \mapsto (A, \delta)) \in \mathit{Flds}(\eta), \delta \in \mathit{dom}(\varpi)\}_b \\
\mathit{RefB}(e) \;=_{df}\; \mathit{case}\ e\ \mathit{of} \\
\quad \begin{array}{lcl}
(A, \iota) & \rightarrow & \{(e, A, \iota)\}_b \\
\mathtt{ret}(v^*, \rho, e) \mid \mathtt{retb}(..., e) \mid w = e & \rightarrow & \mathit{RefB}(e) \\
(t\ v = e_1; e_2) & \rightarrow & \mathit{RefB}(e_1) \uplus_b \mathit{RefB}(e_2) \\
\mathtt{if}\ v\ \mathtt{then}\ e_1\ \mathtt{else}\ e_2 & \rightarrow & \mathit{RefB}(e_1) \cup_b \mathit{RefB}(e_2) \\
\mid w \mid \mathtt{new}\ \mathit{cn}(v^*) \mid \mathit{mn}(v^*) \mid (A, k) & \rightarrow & \emptyset
\end{array}
\end{array}
$$

Note that $\{...\}_b$ denotes a bag, while $\uplus_b$ (summative) and $\cup_b$ are bag union operators.

Definition 9 (Unique Reference) *A reference to a location* $\iota \in \mathit{dom}(\varpi)$ *is said to be unique with respect to a bag of triple references* $G = [\{(a, \mathtt{A}, \iota)^*\}_b$, *if the location is exclusively referenced, as defined below.*

$$
\frac{\mathit{count}(\{(_, A, \iota) \in G \mid A \in \{\mathtt{U}, \mathtt{S}, \mathtt{R}\}\}_b) \leq 1}{\mathit{unique}(G, \iota)}
$$

Definition 10 (Uniqueness Property) *The uniqueness property is defined for a runtime configuration* $\langle \Pi, \varpi \rangle [e]$ *to check that each unique reference is unaliased, as follows:*

$$
\frac{\begin{array}{c}
G_1 = \mathit{RefB}(\Pi) \uplus_b \mathit{RefB}(e) \quad G_2 = \mathit{RefB}(\varpi) \quad G = G_1 \uplus_b G_2 \\
\forall (_, A, \iota) \in G_1 \cdot A = \mathtt{U} \Rightarrow \mathit{unique}(G, \iota) \\
\forall (\iota_1, A, \iota_2) \in G_2 \cdot A = \mathtt{U} \wedge \mathit{unique}(G, \iota_1) \Rightarrow \mathit{unique}(G, \iota_2)
\end{array}}{\mathit{UniqueInv}(\Pi, \varpi, e)}
$$

Definition 11 (Lent-Once Property) *The lent-once property is defined for a runtime environment* $\langle \Pi, \varpi \rangle$ *to check for unique references within each stack frame, as follows:*

$$
\frac{\begin{array}{c}
\forall (a^{s1}, A_a, \iota_a), (b^{s2}, A_b, \iota_b) \in \mathit{RefB}(\Pi) \cdot \ (a \neq b \wedge \\
s1 = s2 \ \wedge \ A_a = \mathtt{L} \ \wedge \ A_b \in \{\mathtt{U}, \mathtt{L}, \mathtt{S}\}) \ \Rightarrow \iota_a \neq \iota_b
\end{array}}{\mathit{LentInv}(\Pi, \varpi)}
$$

Definition 12 (Read-Only Property) *We define as follows that two stores,* ϖ_1 *and* ϖ_2, *satisfy the read-only property during evaluation, where* R*-mode fields never change.*

$$
\frac{\begin{array}{c}
\mathit{dom}(\varpi_1) \subseteq \mathit{dom}(\varpi_2) \\
\forall \{\iota \mapsto (_, \rho_1)\} \in \varpi_1 \cdot \{\iota \mapsto (_, \rho_2)\} \in \varpi_2 \ \Rightarrow \\
(\forall (f \mapsto (A, \delta)) \in \rho_1 \cdot (A = \mathtt{R} \Rightarrow (f \mapsto (A, \delta)) \in \rho_2))
\end{array}}{\mathit{ReadOnlyInv}(\varpi_1, \varpi_2)}
$$

The following theorem states that well-typed expressions (during evaluation) are subject to the uniqueness property, lent-once property, and read-only property.

Theorem 1 (Alias Properties) *If*

$$
\begin{array}{c}
\Gamma; C; \Sigma; \Delta; \Theta \vdash e :: t, \Delta_1, \Theta_1 \\
\Gamma; \Sigma; \Delta; \Theta \models \langle \Pi, \varpi \rangle \\
\langle \Pi, \varpi \rangle\ [e] \hookrightarrow \langle \Pi', \varpi' \rangle\ [e'] \\
\mathit{UniqueInv}(\Pi, \varpi, e) \wedge \mathit{LentInv}(\Pi, \varpi)
\end{array}
$$

then the following holds:

$$
\mathit{UniqueInv}(\Pi', \varpi', e') \wedge \mathit{LentInv}(\Pi', \varpi') \wedge \mathit{ReadOnlyInv}(\varpi, \varpi')
$$

Proof: *By induction over the depth of type derivation for expression* e. *Details are given in the appendix (Sec 10.2).*

The following main theorem states that each well-typed expression preserves its type under reduction with a runtime environment and a store which are consistent with the compile-time counterparts. Furthermore, the final size constraint is consistent with the value obtained on termination.

Theorem 2 (Preservation)

(a) (Expression) If

$$
\begin{array}{c}
\Gamma; C; \Sigma; \Delta; \Theta \vdash e :: t, \Delta_1, \Theta_1 \\
\Gamma; \Sigma; \Delta; \Theta \models \langle \Pi, \varpi \rangle \\
\langle \Pi, \varpi \rangle\ [e] \hookrightarrow \langle \Pi', \varpi' \rangle\ [e']
\end{array}
$$

then there exists Σ', Γ', Δ', *and* Θ', *such that*

$$
\begin{array}{c}
\Gamma - \mathit{diff}(e, e') = \Gamma' - \mathit{diff}(e', e) \\
\Sigma' \supseteq \Sigma \\
\Gamma'; C; \Sigma'; \Delta'; \Theta' \vdash e' :: t, \Delta_1, \Theta_1 \\
\Gamma'; \Sigma'; \Delta'; \Theta' \models \langle \Pi', \varpi' \rangle\ .
\end{array}
$$

(b) (Value) If

$$\Gamma; C; \Sigma; \Delta; \Theta \vdash (A, \delta) :: t, \Delta_1, \Theta_1$$
$$\Gamma; \Sigma; \Delta; \Theta \models \langle \Pi, \varpi \rangle$$

then the following holds:

$$\Theta = \Theta_1$$
$$\Gamma + \{x :: t\}; C; \Sigma; \Delta_2; \Theta_1 \models \langle \Pi + \{x \mapsto (A, \delta)\}, \varpi \rangle$$

where $x = \mathit{fresh}()$; $\Delta_2 = [v \mapsto v']_{v \in V(t)} \Delta_1$.

Function $\mathit{diff}(e, e')$ returns a list of local variables that appears in e but not e' :

$$\begin{array}{lllll} \mathit{diff}(e, e') =_{df} & \mathit{let} & lst & = & \mathit{local}(e) \\ & & lst1 & = & \mathit{local}(e') \\ & & n & = & \mathit{length}(lst) - \mathit{length}(lst1) \\ & \mathit{in} & (\mathit{take}(n, lst) \lhd n \geq 0 \rhd [\,]) & & \end{array}$$
$$\mathit{take}(n, lst) =_{df} ([] \lhd n \leq 0 \rhd [\mathit{hd}(lst)] +\!\!+ \mathit{take}(n-1, \mathit{tl}(lst)))$$

Function $\mathit{local}(e)$ returns a list of sets of local variables. It is defined as follows:

$$\begin{array}{lll} \mathit{local}(e) =_{df} \mathit{case}\ e\ \mathit{of} & & \\ \quad \mathtt{ret}(v^*, \rho, e) \mid \mathtt{retb}(v, v^*, A, \iota, e) & \rightarrow & \mathit{local}(e) +\!\!+ [\{v^*\}] \\ \quad w = e & \rightarrow & \mathit{local}(e) \\ \quad (t\ v = e_1; e_2) & \rightarrow & \mathit{local}(e_1) \\ \quad \mathtt{if}\ v\ \mathtt{then}\ e_1\ \mathtt{else}\ e_2 & & \\ \quad \delta \mid w \mid \mathtt{new}\ \mathit{cn}(v^*) \mid \mathit{mn}(v^*) & \rightarrow & \emptyset \end{array}$$

Note that $\Gamma - [\,] =_{df} \Gamma$, $\Gamma - ([s] +\!\!+ S) =_{df} (\Gamma - s) - S$.
Proof: *By induction over the depth of type derivation for expression* e*. Details are given in the appendix (Sec 10.2).*

Theorem 3 (Progress) *If* $\Gamma; C; \Sigma; \Delta; \Theta \vdash e :: t, \Delta_1, \Theta_1$ *and* $\Gamma; \Sigma; \Delta; \Theta \models \langle \Pi, \varpi \rangle$*, then either* e *is a value, or* $\langle \Pi, \varpi \rangle\ [e] \hookrightarrow$ **Error-Null***, or there exist* Π', ϖ', e' *such that* $\langle \Pi, \varpi \rangle\ [e] \hookrightarrow \langle \Pi', \varpi' \rangle\ [e']$.

Proof: *By induction over the depth of type derivation for expression* e*. Details are given in the appendix (Sec 10.2).*

8 Related Work

Sized type was first proposed by Hughes, Pareto and Sabry [18, 17] as a formal mechanism for proving the correctness of certain size-related properties for embedded functional programs. Their proposal captures an upper or lower bound for each data, depending on whether it is declared as data or co-data. This catered to independent attribute analysis, rather than relational analysis. Separately, dependent type was proposed by Xi and Pfenning [26, 28] to capture size information and has been applied to a variety of applications. These early works are based on the type checking framework. Later, Chin and Khoo [7] enriched sized typing by devising fixed-point methods for safely inferring the size relations for functional programs.

Programs	Size (lines)		Safety Checking
	Source	Annotation	(seconds)
bisort	340	7	0.68
em3d	462	19	0.30
health	562	22	6.08
mst	473	31	0.20
power	765	24	1.32
treeadd	195	6	0.12
tsp	545	10	0.42
perimeter	745	12	8.48
n-body	1128	31	0.37
Voronoi	1000	45	1.90
insert(AVL)	64	1	2.88
bs(array)	141	1	0.08
buffer	65	6	0.12

Figure 5. Performance of Verification

Recently, some advance has been achieved in extending dependent types to imperative programming through a language, called Xanadu [27]. However, Xanadu supports only a limited set of types, including read-only lists; but not general objects. This restriction avoided the aliasing problem, as only *size-immutable* objects are effectively used.

Sharing objects through aliases is both a powerful feature and a weakness of object-oriented languages [1, 4, 14]. As each object may be updated via any alias, object being changed may not even be aware of it; causing such programs to be harder to understand and reason. To regain control over aliasing, various researchers have proposed alias annotations to limit the capability of pointer variables so that local reasoning is possible. A wide range of alias annotations have been proposed, which have been systematically organised into a general capability system[1, 4]. Our work is built on top of such alias controls.

Instead of confining aliasing[14], there have been a number of attempts at analysing possible aliasing structures. Systems such as TVLA [24], PALE [21], and Roles [19], enable the shape of local object graphs to be analysed and captured. Most of these systems employ heavier machinery to capture state change and to enforce the expected aliasing structures. For example, Sagiv, Rep and Wilhelm [24] employed three-value logic to capture points-to relation between objects as well as imperative changes on shape graph. These proposals are complimentary to our goal of size analysis in relational form, and may open up new classes of mutable objects that could be precisely analysed.

In recent years, several approaches have been advocated for verifying software to ensure that it conforms to stated safety policy. One popular approach [15, 2, 9, 5] is based on the concept of model checking in which program states are approximated through boolean predicates. Choosing a suitable predicate abstraction remains one of the main challenges, and various symbolic techniques have been pro-

posed to facilitate this task. With a different approach, several type based techniques [10, 12, 8] have also been advocated for checking if user programs conform to some stated safety protocols. As an approach between the above two, Flanagan et.al.[11] use verification-condition generation and automatic theorem-proving techniques in their extended static checking system for Java. However, they incorporate a trade-off between soundness and usefulness. Morever, most of these proposals are based on finite state approximation of the program states, and are typically formulated in a path-insensitive manner. Through the use of relational size analysis (with disjunctive formulae), our proposal has the potential to give more precise verification. Furthermore, our use of alias controls and size-immutability (often omitted in other works) has provided a clearer path towards analysing object-based imperative programs.

9 Conclusion

We have proposed an advanced type system to capture size relations for objects. We make use of size-immutability to facilitate sharing, and alias controls to identify unaliased objects. This mechanism permits size relation in constraint form to be accurately analysed for object-based programs. Our type system can be viewed as a sound and decidable lightweight verification method. We use the ADT mechanism augmented with size invariants together with pre/post conditions on methods, to allow us to express a variety of size-inspired safety policies - ranging from safe array bounds, data structure invariant to protocol verification.

We have implemented a prototype system which converts annotated Java programs (without concurrency and exceptions) to our kernel language, before type checking is applied to verify the specified safety policies (demo at `http://loris-4.ddns.comp.nus.edu.sg/~nguyenh2/mi`). Our prototype was implemented using the Glasgow Haskell compiler. To evaluate this prototype, we hand-annotated a set of programs from the Java Olden benchmark together with programs described in this paper, before subjecting them to sized type checking. The performance statistics (on Redhat Linux 9.0 for Pentium 2.4 GHz with 768MB) is summarised in Fig 5. Due to modular type checking, our system verifies each medium-sized program under 10 seconds, despite the use of Presburger constraint solving[23]. Size and alias annotations are required for the headers of object and method declarations which account for less than 5% of the source code. Our recent investigations into automatic inference of size constraints for an imperative language without objects [29] (for array bound check elimination) is expected to further reduce the burden of annotations.

References

[1] J. Aldrich, V. Kostadinov, and C. Chambers. Alias Annotation for Program Understanding. In *ACM OOPSLA*, Seattle, Washington, November 2002.

[2] T. Ball, R. Majumdar, T. Millstein, and S. K. Rajamani. Automatic Predicate Abstraction of C Programs. In *ACM PLDI*, June 2001.

[3] F. Benoy and A. King. Inferring argument size relationships with CLP(R). In *Logic Programming Synthesis and Transformation*, Springer-Verlag, August 1997.

[4] J. Boyland, J. Noble, and W. Retert. Capabilities for Sharing: A Generalization of Uniqueness and Read-Only. In *ECOOP*, Budapest, Hungary, June 2001.

[5] S. Chaki, E. M. Clarke, A. Groce, S. Jha, and H. Veith. Modular Verification of Software Components in C. In *Proceedings of the International Conference on Software Engineering (ICSE'03)*, pages 385–395, 2003.

[6] E. C. Chan, J. Boyland, and W. L. Scherlis. Promises: Limited Specifications for Analysis and Manipulation. In *Proceedings of the International Conference on Software Engineering (ICSE'98)*, Kyoto, Japan, April 1998.

[7] W.N. Chin and S.C. Khoo. Calculating sized types. In *ACM SIGPLAN PEPM*, pages 62–72, Boston, Massachusetts, United States, January 2000.

[8] T. Colcombet and P. Fradet. Enforcing Trace Properties by Program Transformation. In *ACM POPL*, January 2000.

[9] J. C. Corbett, M. B. Dwyer, J. Hatcliff, S. Laubach, C. S. Pasareanu, Robby, and H. Zheng. Bandera: extracting finite-state models from Java source code. In *Proceedings of the International Conference on Software Engineering (ICSE'00)*, pages 439–448, Limerick, Ireland, June 2000.

[10] Manuel Fahndrich and Robert DeLine. Adoption and Focus: Practical linear types for imperative programming. In *ACM PLDI*, June 2002.

[11] C. Flanagan, K.R.M. Leino, M. Lillibridge, G. Nelson, J.B. Saxe, and R. Stata. Extended Static Checking for Java. In *ACM PLDI*, Berlin, Germany, June 2002.

[12] J.S. Foster, T. Terauchi, and A. Aiken. Flow-Sensitive Type Qualifiers. In *ACM PLDI*, Berlin, Germany, June 2002.

[13] B. Grobauer. Cost Recurrences for DML Programs. In *Proceedings of the International Conference on Functional Programming (ICFP '01)*, pages 253–277, Florence, Italy, September 2001.

[14] C. Grothoff, J. Palsberg, and J. Vitek. Encapsulating Objects with Confined Types. In *ACM OOPSLA*, October 2001.

[15] T. A. Henzinger, R. Jhala, R. Majumdar, and G. Sutre. Lazy Abstraction. In *ACM POPL*, pages 58–70, Portland, Oregon, January 2002.

[16] M. Hofmann and S. Jost. Static Prediction of Heap Space Usage for First Order Functional Programs. In *ACM POPL*, New Orleans, Louisiana, January 2003.

[17] J. Hughes and L. Pareto. Recursion and Dynamic Data-Structures in Bounded Space: Towards Embedded ML Programming. In *Proceedings of the International Conference on Functional Programming (ICFP '99)*, September 1999.

[18] J. Hughes, L. Pareto, and A. Sabry. Proving the Correctness of Reactive Systems Using Sized Types. In *ACM POPL*, pages 410–423. ACM Press, January 1996.

[19] V. Kuncak, P. Lam, and M. Rinard. Role analysis. In *ACM POPL*, Portland, Oregon, January 2002.

[20] C. S. Lee, N. D. Jones, and A. M. Ben-Amram. The size-change principle for program termination. In *ACM POPL*, pages 81–92. ACM Press, 2001.

[21] A. Moeller and M. I. Schwartzbach. The Pointer Assertion Logic Engine. In *ACM PLDI*, June 2001.

[22] B. C. Pierce. Intersection types and bounded polymorphism. *Mathematical Structures in Computer Science*, 7(2):129–193, 1997.

[23] W. Pugh. The Omega Test: A fast practical integer programming algorithm for dependence analysis. *Communications of the ACM*, 8:102–114, 1992.

[24] M. Sagiv, T. Reps, and R. Wilhelm. Parametric Shape Analysis via 3-valued Logic. In *ACM POPL*, January 1999.

[25] K. Sohn and A. Van Gelder. Termination detection in logic programs using argument sizes (extended abstract). In *Proceedings of the 10th ACM Symposium on Principles of Database Systems*, pages 216–226, 1991.

[26] H. Xi. *Dependent Types in Practical Programming*. PhD thesis, Carnegie Mellon University, 1998.

[27] H. Xi. Imperative Programming with Dependent Types. In *Proceedings of 15th Symposium on Logic in Computer Science (LICS '00)*, Santa Barbara, June 2000.

[28] H. Xi and F. Pfenning. Eliminating array bound checking through dependent types. In *ACM PLDI*, pages 249–257. ACM Press, June 1998.

[29] D.N. Xu, C. Popeea, S.C. Khoo, and W.N. Chin. Modular Inference for Array Checks Optimization. Technical report, SoC, Natl Univ. of Singapore, November 2003. avail. at http://www.comp.nus.edu.sg/~xun/research/array_tech.ps.

APPENDIX

10 Proofs

10.1 Theorem 1

Proof: By induction over the depth of type derivation for expression e. In the following, let $G_1 = \mathit{RefB}(\Pi) \uplus_b \mathit{RefB}(e)$, $G_2 = \mathit{RefB}(\varpi)$, $G = G_1 \uplus_b G_2$, $G'_1 = \mathit{RefB}(\Pi') \uplus_b \mathit{RefB}(e')$, $G'_2 = \mathit{RefB}(\varpi')$, $G' = G'_1 \uplus_b G'_2$.

Case [CONS1–4]. Trivial.

Case [VAR]. We deal with the expression v. $\mathit{LentInv}(\Pi', \varpi')$ holds as v can not be in lent mode. $\mathit{ReadOnlyInv}(\varpi, \varpi')$ follows as the evaluation does not update any read-only fields. During the evaluation ([**D-Var**]), v will lose its uniqueness (from U to $\mathtt{U}_D$) if it is unique, or remain unchanged otherwise. Take note that $G'_1 = G_1$, and $G'_2 = G_2$. The uniqueness property is preserved.

Case [ASSIGN]. We deal with expression $v = e$. There are two rules by which one step evaluation can be conducted:

Subcase [D-Assign-1] We have $\Gamma; C; \Sigma; \Delta; \Theta \vdash e :: t_1, \Delta_1, \Theta_1$, and $\langle \Pi, \varpi \rangle\ [e] \hookrightarrow \langle \Pi', \varpi' \rangle\ [e']$. By induction hypothesis, we conclude.

Subcase [D-Assign-2] $\mathit{LentInv}(\Pi', \varpi')$ holds immediately as the evaluation does not involve any lent references. Let $\Pi(v) = (A, \delta)$, then $\varpi' = \varpi$, $\mathit{ReadOnlyInv}(\varpi, \varpi')$ follows immediately. If $A = \mathtt{S}$, the uniqueness property is preserved as the evaluation does not involve unique references. If $A = \mathtt{U}$ or $A = \mathtt{U}_D$, $\Pi(v)$ is updated with a unique location $(\mathtt{U}, \delta_s)$ due to the alias subtyping (note that $\mathit{unique}(G, \delta_s)$ holds from the assumption, where $(_, \mathtt{U}, \delta_s) \in \mathit{RefB}(e)$). During the evaluation, $(v, \mathtt{U}, \delta_s)$ is added into $\mathit{RefB}(\Pi')$, while $(_, \mathtt{U}, \delta_s)$ is removed from e yielding $e' = (\mathtt{S}, ())$. $G' = G$. The uniqueness property is still preserved.

Case [NEW]. We deal with expression **new** $cn(v_1, .., v_p)$. The property $\mathit{ReadOnlyInv}(\varpi, \varpi')$ holds, as all read-only field references from ϖ remain unchanged in ϖ' during the evaluation ([**D-New**]). As field references are not allowed to be L, there are no L parameters involved due to subtyping ([**D-AUX3**]), therefore, $\mathit{LentInv}(\Pi', \varpi')$ follows. Whenever a unique reference v_i is passed to field f_i, its uniqueness is lost immediately, as guaranteed by function *updVE*. That is, $(v_i, \mathtt{U}, \iota_i)$ becomes $(v_i, \mathtt{U}_D, \iota_i)$. Alias subtyping also ensures that unique fields be assigned unique locations. That is, $(\iota, \mathtt{U}, \iota_i)$ is added into G_2, yielding G'_2. By applying the function to $v_1, \ldots, v_p$ sequentially, it ensures that a unique location cannot be passed to two unique fields during the evaluation. Thus, we have $G' = G \uplus_b \{(_, \mathtt{U}, \iota)\}_b$. Moreover, ι is a fresh location in the store, this implies $\mathit{unique}(G', \iota)$. Therefore, the uniqueness property is preserved.

Case [CALL]. We deal with expression $mn(v_1, .., v_p)$. We only consider the subcase [**D-Call**], the proof for [**D-Call-prim**] is similar. $\mathit{ReadOnlyInv}(\varpi, \varpi')$ follows immediately from $\varpi' = \varpi$. The condition *distinct* in [**D-Call**] ensures that all L variables in the new stack frame point to distinct locations. This, together with $\mathit{LentInv}(\Pi, \varpi)$ and $\varpi' = \varpi$, ensures $\mathit{LentInv}(\Pi', \varpi')$. During the parameter passing, unique actual parame-

ters (v_i) lose their uniqueness, while unique formal parameters ($\hat{v}_i$) gain the uniqueness. Updating the variable environment sequentially ensures that no unique v_i can pass its uniqueness to two different parameters. Moreover, the evaluation does not introduce any U-locations into e'. Also note that $\varpi' = \varpi$. Thus, the uniqueness property *UniqueInv*(Π', ϖ', e') still holds.

Case [LOCAL]. We deal with expression $t\ v = e_1; e_2$. There are two rules by which one step evaluation can be conducted:

Subcase [D-Blk-1] We have $\Gamma; C; \Sigma; \Delta; \Theta \vdash e_1 :: t_1, \Delta_1, \Theta_1$, and $\langle \Pi, \varpi \rangle\ [e_1] \hookrightarrow \langle \Pi', \varpi' \rangle\ [e_1']$. By induction hypothesis, we conclude.

Subcase [D-Blk-2] *ReadOnlyInv*(ϖ, ϖ') follows immediately from $\varpi' = \varpi$. *LentInv*(Π', ϖ') also holds as no new L references are introduced in Π'. Suppose $\nu = (A, \iota)$. Due to the alias subtyping, if $A = \mathtt{U}$, $(e, \mathtt{U}, \iota)$ is removed from G_1, while $(v, \mathtt{U}, \iota)$ is added into G_1'. That is, $G' = G$. *unique*(G, ι) still holds after the initialisation. If $(A = \mathtt{S})$, we still have $G' = G$. Therefore, the uniqueness property is preserved.

Case [IF]. We deal with expression **if** v **then** e_1 **else** e_2. There are two rules by which one step evaluation can be conducted: [**D-If-true**] and [**D-If-false**]. In either case, the variable environment and the store remain unchanged, and *RefB*$(e_i) \subseteq$ *RefB*(e), thus the properties hold.

Case [ELF].We deal with expression $\mathtt{ret}(v^*, \rho, e)$. There are two rules by which one step evaluation can be conducted:

Subcase [D-ret-1] We have $\Gamma; C; \Sigma; \Delta; \Theta \vdash e :: t, \Delta_1, \Theta_1$, and $\langle \Pi, \varpi \rangle\ [e] \hookrightarrow \langle \Pi', \varpi' \rangle\ [e']$. By induction hypothesis, we conclude.

Subcase [D-ret-2] Local variables (v^*) are removed from the latest stack frame, while the store remains unchanged. This preserves the Lent-Once property, Read-Only property. In case that the result $\nu = (\mathtt{U}, \iota)$, from the assumption, we have *unique*(G, ι). The uniqueness property holds as $G_1' = G_1, G_2' = G_2$.

Case [BIND]. We deal with expression $\mathtt{bind}\ v\ \mathtt{of}\ (v_1, .., v_n)\ \mathtt{in}\ e$. We only need to check the uniqueness and lent-once property in the case where $\Pi(v) = (\mathtt{U}, \iota)$. It is concluded from the hypothesis and the dynamic rule **D-Bind**.

Case [RETB]. We deal with expression $\mathtt{retb}(v, v_{1..n}, A, \iota, e)$.

Subcase [D-Retb-1] By induction hypothesis.

Subcase [D-Retb-2] For the case where $\Pi_1(v) = (\mathtt{U}, \iota)$, it can be concluded from the fact that v_0 is the only reference to ι after the evaluation.

□

10.2 Theorem 2

The proof for Theorem 2 requires a lemma, called *assumption weakening lemma*, that states that the static judgment remains valid despite a variation of its assumption. This assumes the store type Σ to have unbounded mapping of locations to types. However, the type environment Γ takes the form of stackable mapping between variables and types, and is allowed to grow (by pushing in new mappings) and shrink (by popping out mappings from stack). The lemma states that such change to type environment preserves the type judgment, if the change are properly constrained.

Lemma 1 (Assumption Weakening) *Given that the following judgment holds:*

$$\Gamma; C; \Sigma; \Delta; \Theta \vdash e :: t, \Delta_1, \Theta_1$$

Let Γ' *and* Σ' *be such that:*

$$\begin{gathered} V_u(\Delta) \subseteq V(\Gamma') \qquad V(t) \cap V(\Gamma') = \emptyset \\ (vars(e) \cup \{v \mid v \in \Theta\}) \subseteq dom(\Gamma') \\ \exists v^* \cdot (\Gamma - \{v^*\} = \Gamma') \vee (\Gamma' - \{v^*\} = \Gamma) \\ \Sigma' \supseteq \Sigma \end{gathered}$$

Then,

$$\Gamma'; C; \Sigma'; \Delta; \Theta \vdash e :: t, \Delta_1, \Theta_1$$

The call $vars(e)$ returns all variables occurring in e.

Proof for Lemma 1: By structural induction on the static semantics of the form $\Gamma; C; \Sigma; \Delta; \Theta \vdash e :: t, \Delta_1, \Theta_1$. For any Γ' and Σ', we say that the pair satisfy the premises of the Lemma if the following holds:

$$\begin{gathered} V_u(\Delta) \subseteq V(\Gamma') \qquad V(t) \cap V(\Gamma') = \emptyset \\ (vars(e) \cup \{v \mid v \in \Theta\}) \subseteq dom(\Gamma') \\ \exists v^* \cdot (\Gamma - \{v^*\} = \Gamma') \vee (\Gamma' - \{v^*\} = \Gamma) \\ \Sigma' \supseteq \Sigma \end{gathered}$$

Case [CONS1–4]. Trivial.

Case [VAR]. We deal with the expression v. For any Γ' satisfying the premises of the lemma, we see that $\Gamma'(v) = \Gamma(v)$, and thus $\Gamma' \vdash v :: t, \phi, Y$. Thus, $\Gamma'; C; \Sigma'; \Delta; \Theta \vdash v :: t, \Delta \wedge \phi, \Theta_1$.

Case [ASSIGN]. We deal with expression $v = e$. For any Γ' and Σ' satisfying the premises of the Lemma, we have $\Gamma'; C; \Sigma'; \Delta; \Theta \vdash e :: t_1, \Delta_1, \Theta_1$ by the induction hypothesis. Furthermore, *prime*(Γ') $\vdash$ $v :: t, \phi, Y$ can be derived. Thus, $\Gamma'; C; \Sigma'; \Delta; \Theta \vdash v = e :: \mathtt{void}\langle\rangle@\mathtt{S}, \Delta_2, \Theta_1/v$.

Case [NEW]. We deal with expression **new** $cn(v_1, .., v_p)$. For any Γ' and Σ' satisfying the premises of the Lemma, we have $\Gamma'(v_i) = t_i$, $\forall i : 1..p$. Let $\Delta_2 = \Delta \wedge (\exists X \cdot \rho\phi)$, then $\Gamma'; C; \Sigma'; \Delta; \Theta_0 \vdash$ **new** $cn(v_{1..p}) :: cn\langle r^* \rangle@\mathtt{U}, \Delta_2, \Theta_p$.

Case [CALL]. We deal with expression $mn(v_1, .., v_p)$. This can be proven using the similar argument as that in [NEW].

Case [BIND]. Similar argument as that in [NEW].

Case [LOCAL]. We deal with expression $t\ v = e_1; e_2$. For any Γ' and Σ' satisfying the premises of the Lemma, we obtain $\Gamma'; C; \Sigma'; \Delta; \Theta \vdash e_1 :: t_1, \Delta_1, \Theta_1$ and $\Gamma', v :: t; C; \Sigma'; \exists X \cdot \rho\Delta_1; \Theta_1 \vdash e_2 :: t_2, \Delta_2, \Theta_2$ by the induction hypothesis. Hence, $\Gamma'; C; \Sigma'; \Delta; \Theta \vdash t\ v = e_1\ ; e_2 :: t_2, \exists Y \cdot \Delta_2, \Theta_2/v$.

Case [IF]. We deal with expression **if** v **then** e_1 **else** e_2. For any Γ' and Σ' satisfying the premises of the Lemma, we have $\Gamma'(v) = \Gamma(v)$, hence $\Gamma'(v) = \textbf{bool}\langle b\rangle$. Furthermore, we have $\Gamma'; C; \Sigma'; \Delta \wedge b' = 1; \Theta \vdash e_1 :: t_1, \Delta_1, \Theta_1$ and $\Gamma'; C; \Sigma'; \Delta \wedge b' = 0; \Theta \vdash e_2 :: t_2, \Delta_2, \Theta_2$ by the induction hypothesis. Hence, $\Gamma'; C; \Sigma'; \Delta; \Theta \vdash$ **if** v **then** e_1 **else** $e_2 :: t, \Delta_3, \Theta_3$.

Case [ELF]. We deal with expression $\texttt{ret}(v^*, \rho, e)$. For any Γ' and Σ' satisfying the premises of the Lemma, $v^* \subseteq dom(\Gamma')$ and $\Gamma'; C; \Sigma'; \Delta; \Theta \vdash e :: t, \Delta_1, \Theta_1$ by the induction hypothesis. Thus, $\Gamma'; \Sigma'; \Delta; \Theta \vdash \texttt{ret}(v^*, \rho, e) :: t, \Delta_2, \Theta_1/\{v^*\}$ where $\Delta_2 = \exists Y \cdot (prime(\rho_2)\Delta_1)$.

Case [RETB]. Similar to the [ELF] rule.

Case [LOC] and [A−DATA]. Trivial. □

Proof for Theorem 2:

(a) By induction over the depth of type derivation for expression e that will be evaluated. At each step of the induction, we assume that the desired property holds for each induction step, and proceed by case analysis on the next derivation step. Our induction is carried out on the expected forms of well-type expressions, as follows:

Case [VAR] Given that $\Gamma; C; \Sigma; \Delta; \Theta \vdash v :: t, \Delta \wedge \phi, \Theta_1$, and $\langle \Pi, \varpi\rangle\ [w] \hookrightarrow \langle \Pi', \varpi'\rangle\ [\nu]$ for some value ν, with $\Gamma; \Sigma; \Delta; \Theta \models \langle \Pi, \varpi\rangle$, we let $\Gamma' = \Gamma$, $\Sigma' = \Sigma$, $\Delta' = \Delta \wedge \phi$, and $\Theta' = \Theta_1$. It is then easy to show that $\Gamma'; \Sigma'; \Delta'; \Theta' \models \langle \Pi', \varpi'\rangle$, and we hence arrive at $\Gamma'; C; \Sigma'; \Delta'; \Theta' \vdash \nu :: t, \Delta', \Theta'$ via subtyping of Δ'.

Case [ASSIGN] We deal with well-typed expressions of the form $v = e$. From the static semantics, we have $\Gamma; C; \Sigma; \Delta; \Theta \vdash e :: \textbf{void}\langle\rangle\texttt{@S}, \Delta_2, \Theta_1/v$. There are two rules by which one-step derivation can be derived:

Subcase [D-Assign-1] By induction hypothesis, there exists Γ', Σ', Δ' and Θ' such that $\Gamma'; \Sigma'; \Delta'; \Theta' \models \langle \Pi', \varpi'\rangle$ and $\Gamma; C; \Sigma; \Delta; \Theta \vdash e' :: t_1, \Delta_1, \Theta_1$. The relationship between Γ' and Γ for them to be consistent ensures that $prime(\Gamma') \vdash v :: t, \phi, Y$, as shown in the premise of the rule [ASSIGN]. Now, by applying the rule [ASSIGN] on $\Gamma', \Sigma', \Delta'$ and Θ', we obtain the result that $\Gamma'; C; \Sigma'; \Delta'; \Theta' \vdash v = e' :: \textbf{void}\langle\rangle\texttt{@S}, \Delta_2, \Theta_1/v$.

Subcase [D-Assign-2] Here, e is some run-time annotated value $\nu = (A_s, \delta_s)$. From the dynamic semantics, we obtain $\Pi' = upd(\Pi, v, \nu)$, and need to find $\Gamma', \Sigma', \Delta', \Theta'$ such that $\Gamma'; C; \Sigma'; \Delta'; \Theta' \vdash () :: \textbf{void}\langle\rangle\texttt{@S}, \Delta_2, \Theta_1/w$ and $\Gamma'; \Sigma'; \Delta'; \Theta' \models \langle \Pi', \varpi\rangle$. This is attained by setting $\Gamma' = \Gamma, \Sigma' = \Sigma, \Delta' = \Delta_2$ and $\Theta' = \Theta/v$, where $\Delta_2 = \exists X \cdot \Delta_1 \circ_Y \rho(\phi)$ as given in the static rule for [ASSIGN].

Case [NEW] In the premise, we have $\Gamma; C; \Sigma; \Delta; \Theta \vdash$ **new** $cn(v_{1..p}) :: cn\langle r^*\rangle\texttt{@U}, \Delta_1, \Theta_p$ as well as $\langle \Pi_0, \varpi_0\rangle\ [$**new** $cn(v_{1..p})] \hookrightarrow \langle \Pi, \varpi'\rangle\ [(\texttt{U}, \iota)]$. We choose $\Gamma' = \Gamma$, $\Sigma' = (\Sigma, \{\iota \mapsto cn\langle n^*\rangle\texttt{@U}\})$, $\Delta' = \Delta_1$, and $\Theta' = \Theta_p$, and the proof follows.

Case [CALL] Given $\Gamma; C; \Sigma; \Delta; \Theta \vdash mn(v_1, \ldots, v_p) :: t_0, \Delta_2, \Theta_p$ where $\Delta_2 = \Delta \circ_L \exists Y \cdot \rho(\phi_{po})$. There are two subcases.

Subcase [D-CALL] Let the one-step reduction be $\langle \Pi, \varpi\rangle\ [mn(v_1, \ldots, v_p)] \hookrightarrow \langle \Pi', \varpi\rangle\ [\texttt{ret}(\hat{v}_1..\hat{v}_p, \rho_m, e)]$, where ρ_m maps formals to actuals, we need to show that there exists Γ', Σ', Δ', Θ' such that $\Gamma'; \Sigma'; \Delta'; \Theta' \models \langle \Pi', \varpi\rangle$ and $\Gamma'; C; \Sigma'; \Delta'; \Theta' \vdash \texttt{ret}(\hat{v}_1..\hat{v}_p, \rho_m, e) :: t_0, \Delta_2, \Theta_p$.
Given that $\hat{t}_0\ mn((\hat{t}_i\ \hat{v}_i)_{1..p})$ **where** $\phi_{pr}; \phi_{po}\ \{e\} \in P$, we choose $\Gamma' = (\Gamma, \hat{v}_1^k :: \hat{t}_1, \ldots, \hat{v}_p^k :: \hat{t}_p)$, where k is the new stack frame number. Furthermore, we choose $\Sigma' = \Sigma$, $\Delta' = \Delta \circ_L \Delta_e$ where $\Delta_e = n\alpha\mathcal{X}(\cup_{i=1}^{p}\{V_u(\hat{t}_i)\})$, and $\Theta' = \Theta_p$. Both $\mathcal{C}_{\textbf{alias}}(\Gamma', \Sigma', \Theta', \Pi', \varpi)$ and $\mathcal{C}_{\textbf{size}}(\Gamma', \Sigma', \Delta', \Theta', \Pi', \varpi)$ are ensured by the definition of *updVE* and the respective premises. Since the judgment for *conUL* is consistent with the operation *updVE*, that $\Gamma'; \Sigma'; \Delta'; \Theta' \models \langle \Pi', \varpi\rangle$ holds can be easily proven. Since no variables/fields in Θ will be used by the method body e, checking the method body only involves size variables of formal parameters $(\hat{v}^k)^*$. Thus, we obtain $\Gamma'; C; \Sigma'; \Delta'; \Theta' \vdash e :: t, \Delta \circ_L \Delta_M, \Theta' \cup \Theta_M$, where Δ_M and Θ_M are the Δ_1 and Θ in the rule [METH]. From the rule [ELF] for local frame, we obtain $\Gamma'; C; \Sigma'; \Delta'; \Theta' \vdash \texttt{ret}(\hat{v}_1..\hat{v}_p, \rho_m, e) :: t, \Delta'', \Theta'$, where $\Delta'' = \exists Y_m \cdot \rho(\Delta \circ_L \Delta_M)$. Note that $\exists Y_m \cdot \rho(\Delta \circ_L \Delta_M)$ = $\Delta \circ_L \exists Y_m \cdot \rho(\Delta_M) \Rightarrow \Delta \circ_L \exists Y_m \cdot \rho(\phi_{po})$ = Δ_2.

Subcase [D-CALL-Prim] Given that *mn* is a primitive call, we assume that the size property ϕ_{po} associated with the primitive method is consistent with the returned result of the primitive call invocation, including modification of store. Proof of alias consistency follows that of the subcase [D-CALL]. Thus, the conclusion follows.

Case [LOCAL] From the type rule, we have $\Gamma; C; \Sigma; \Delta; \Theta \vdash (t\ v = e_1\ ; e_2) :: \texttt{void}\langle\rangle\texttt{@S}, \Delta_1, \Theta_1$.
We consider the two rules by which one-step derivation can be derived:

Subcase [D-Blk-1] By the induction hypothesis and assumption weakening lemma.

Subcase [D-Blk-2] By induction hypothesis and the rule [ELF].

Case [IF] Both the subcases [D-If-1] and [D-If-2] can be easily proven by the induction hypothesis and the covariance subsumption rule [SUB1].

Case [ELF] There are two subcases: The subcases [D-Ret-1] is proven by the induction hypothesis. For the subcase [D-Ret-2], we have in the premise that $\Gamma; C; \Sigma; \Delta; \Theta \vdash \texttt{ret}(v^*, [\,], \nu) :: t_2, \Delta_2, \Theta_2$, and $\langle \Pi, \varpi \rangle\ [\texttt{ret}(v^*, [\,], \nu)] \hookrightarrow \langle \Pi', \varpi \rangle\ [\nu]$. We choose $\Gamma' = \Gamma - \{v^*\}$, $\Sigma' = \Sigma$, $\Delta' = \Delta_2$, and $\Theta' = \Theta_2$. That $\Gamma'; \Sigma'; \Delta'; \Theta' \models \langle \Pi', \varpi \rangle$ can be easily shown from the premise. For the judgment $\Gamma'; C; \Sigma'; \Delta'; \Theta' \vdash \nu :: t_2, \Delta'', \Theta''$, we can show that $\Theta'' = \Theta'$ and $\Delta'' \Rightarrow \Delta'$ by a case analysis on the possible type of ν. □

Case [BIND] From the type rule, we have $\Gamma; C; \Sigma; \Delta; \Theta \vdash \texttt{bind}\ v\ \texttt{of}\ (v_1, \ldots, v_p)\ \texttt{in}\ e :: t_0, \Delta_4, \Theta_2$ where $\Delta_4 = \exists Z \cdot \Delta_3$. From the dynamic semantics, we have $\langle \Pi, \varpi \rangle\ [\texttt{bind}\ v\ \texttt{of}\ (v_{1..n})\ \texttt{in}\ e] \hookrightarrow \langle \Pi_1, \varpi \rangle\ [e_1]$, where $e_1 = \texttt{retb}(v, v_{1..n}, A, \iota, e)$. We choose Γ' = Γ - $\{v | A \in \{\texttt{L}, \texttt{U}\}\} \cup \Gamma_1$, where Γ_1 is as defined in the [BIND] rule. We also choose Σ' = Σ, Δ' = Δ_1 (which is defined in [BIND] rule), and Θ' = Θ. It is easy to verify that $\Gamma'; \Sigma'; \Delta'; \Theta' \models \langle \Pi_1, \varpi \rangle$. The judgment for $\Gamma; C; \Sigma; \Delta; \Theta \vdash e_1 :: t_0, \Delta_4, \Theta_2$ thus follows.

Case [RETB] There are two subcases: The subcase [D-Retb-1] is proven by the induction hypothesis. For the subcase [D-Retb-2], we have in the premise that $\Gamma; C; \Sigma; \Delta; \Theta \vdash \texttt{retb}(v_{0..n}, A, \iota, \nu) :: t_2, \Delta_2, \Theta_2$, and $\langle \Pi, \varpi \rangle\ [\texttt{retb}(v_{1..n}, A, \iota, \nu)] \hookrightarrow \langle \Pi_1, \varpi_2 \rangle\ [\nu]$. Let $I = \exists i \in 1..n \cdot \Pi(v_i) = (\texttt{U}_D, _)$. We choose Γ' = $(\Gamma \lhd I \rhd (\Gamma + \{v_0 :: t_\iota\}))$ - $\{v_{1..n}\}$, where t_ι is obtained from the judgement $\Gamma; C; \Sigma; \Delta; \Theta \vdash \iota :: t_\iota, \Delta_\iota, \Theta_\iota$. We also choose $\Sigma' = \Sigma$, $\Delta' = \Delta_2$, and $\Theta' = \Theta_2 \cup \{v_0\} \lhd I \rhd \Theta_2$. That $\Gamma'; \Sigma'; \Delta'; \Theta' \models \langle \Pi', \varpi \rangle$ can be easily shown from the premise. For the judgment $\Gamma'; C; \Sigma'; \Delta'; \Theta' \vdash \nu :: t_2, \Delta'', \Theta''$, we can show that $\Theta'' = \Theta'$ and $\Delta'' \Rightarrow \Delta'$ via subtyping.

(b) This can be proved by establishing that each well-typed constant from [CONS1] to [CONS4] does not change Θ. In addition, the resulted post-condition Δ obtained is consistent with the run-time stack and store. □

10.3 Theorem 3

Proof: By induction over the depth of type derivation for expression e.

Case [CONS1–4]. Trivial.

Case [VAR]. We deal with the expression v. As v is well-typed, from type rule [VAR], v is not in $\texttt{U}_D$ or L mode. Thus from the evaluation rule [D-Var] (the function *read*), the evaluation either reports an **Error-Null**, or advances one step yielding a value.

Case [ASSIGN]. We deal with expression $v = e$. From the type rule, we know $\Gamma; C; \Sigma; \Delta; \Theta \vdash e :: t_1, \Delta_1, \Theta_1$. By induction hypothesis, either (i) e is a value ν, or (ii) $\langle \Pi, \varpi \rangle\ [e] \hookrightarrow$ **Error-Null**, or (iii) $\langle \Pi, \varpi \rangle\ [e] \hookrightarrow \langle \Pi', \varpi' \rangle\ [e']$.
In case (i), from the well-typedness checking, v is not in L or R mode, and the RHS and LHS subsumes the alias subtyping, from the evaluation rule [D-Assign-2] (the function *upd*), either the evaluation reports an **Error-Null**, or the one-step evaluation succeeds, yielding a value $(\texttt{S}, ())$.
In case (ii), the evaluation for the assignment reports an **Error-Null**.
In case (iii), from evaluation rule [D-Assign-1], the evaluation advances with one-step.

Case [NEW]. We deal with expression **new** $cn(v_1, .., v_p)$. From the type rule [AUX3] (used in [NEW]), No arguments v_i are in the consumed set, and the type of each argument and that of its corresponding field conform to the (alias) subtyping relation. Thus from the evaluation rule [D-New] (the function *updVE*), the evaluation succeeds, yielding a newly created object in the store.

Case [CALL]. We deal with expression $mn(v_1, .., v_p)$. We only consider the subcase [D-Call], the proof for [D-Call-prim] is similar. As guaranteed by type rules [AUX3] (used in [CALL]), no unique arguments are consumed before the method call, and the type of each argument and that of its corresponding field conform to the (alias) subtyping relation. Thus from the evaluation rule [D-Call] (the function *updVE*), the evaluation succeeds, yielding the intermediate ret-expression.

Case [LOCAL]. We deal with expression $t\ v = e_1; e_2$. From the type rule, we have $\Gamma; C; \Sigma; \Delta; \Theta \vdash e_1 :: t_1, \Delta_1, \Theta_1$. By induction hypothesis, either (i) e_1 is a value ν, or (ii) $\langle \Pi, \varpi \rangle\ [e_1] \hookrightarrow$ **Error-Null**, or (iii) $\langle \Pi, \varpi \rangle\ [e_1] \hookrightarrow \langle \Pi', \varpi' \rangle\ [e_1']$.
In case (i), the type rule [LOCAL] guarantees the alias subtyping relation needed in the evaluation [D-Blk-2] (the function *ext*). Thus the evaluation succeeds. In case (ii), the evaluation reports **Error-Null** immediately. In case (iii), the proof follows from [D-Blk-1].

Case [IF]. We deal with expression **if** v **then** e_1 **else** e_2. The type rule ensures v is of type bool, thus there always exists an evaluation step to be taken in spite of the value of v, that is, either [D-If-true] or [D-If-false].

Case [ELF]. We deal with expression $\texttt{ret}(v^*, \rho, e)$. From the type rule, we have $\Gamma; C; \Sigma; \Delta; \Theta \vdash e :: t, \Delta_1, \Theta_1$. By induction hypothesis, either (i) e is a value ν, or (ii) $\langle \Pi, \varpi \rangle\ [e] \hookrightarrow$ **Error-Null**, or (iii) $\langle \Pi, \varpi \rangle\ [e] \hookrightarrow \langle \Pi', \varpi' \rangle\ [e']$. In case (i), the proof follows from [**D-Ret-2**]. In case (ii), the evaluation reports **Error-Null** immediately. In case (iii), the proof follows from [**D-Ret-1**].

Case [BIND]. We deal with expression $\texttt{bind}\ v\ \texttt{of}\ (v_1, .., v_n)\ \texttt{in}\ e$. Guaranteed by the type rule and the consistency relation, the only possible evaluation step given by rule [**D-Bind**] succeeds.

Case [RETB]. We deal with expression $\texttt{retb}(v, v_{1..n}, A, \iota, e)$. In case that e is not a value ν, it is concluded by induction hypothesis and [**D-Retb-1**]. Otherwise, it is concluded by [**D-Retb-2**]. □