

A Type System With Containers

by

Michael Thode

A thesis
presented to the University of Waterloo
in fulfillment of the
thesis requirement for the degree of
Master of Mathematics
in
Computer Science

Waterloo, Ontario, Canada, 2021

© Michael Thode 2021

Author's Declaration

I hereby declare that I am the sole author of this thesis. This is a true copy of the thesis, including any required final revisions, as accepted by my examiners.

I understand that my thesis may be made electronically available to the public.

Abstract

In this thesis, we will introduce the concept of containers as they apply to programming languages. Encapsulation is a common topic in programming languages with well understood benefits. Here, we will investigate its converse, namely containment. This includes a demonstration of how containers can be integrated into a programming language and what benefits they can bring.

To support containment, a dependent type system is developed to enforce container rules. We add the notion of a container label to our types to indicate the container of the referred object. Around this type system we develop a language enhanced with container syntax. We use this language to show how containers can enable pass-by-value semantics, copying of complex objects and object serialization. An interpreter is implemented for this language to demonstrate its capabilities. Included is a container inferencing algorithm intended to minimize the extra syntax needed for container specification.

A second formal system is also defined. This includes type rules, operational semantics and a proof of soundness. We show that correctly-typed programs will obey all container restrictions at run-time. We fully type the configuration used by the semantics; this includes concrete containers as run-time constructs which allow us to verify correct containment. Mappings are maintained from the container labels of the language to physical run-time containers. We show that as container labels are translated across scopes (e.g. a function call), the physical containers remain consistent.

We conclude with a discussion on ways this system can be enhanced in the future to make containers easier to use, as well as describe additional capabilities such as version control of objects.

Acknowledgements

A big thanks to my supervisor Ondřej Lhoták who challenged me throughout the process of this thesis. It's a mystery to me how Ondřej can be so consistently critical and yet equally positive and encouraging. The constructive criticism is the foundation of what I've learned and the encouragement kept me moving forward.

To Gregor Richards and Yizhou Zhang, thank you for agreeing to review this thesis and the work you invested into making it better.

Thank you to Werner Dietl, for identifying key related works in the area of ownership types.

To all members of the PLG group, the informal daily conversations prior to the Covid-19 lock-down were often amusing, but just as frequently full of random insights into complex topics. I miss those days and hope all of you are doing well.

To my family who put up with me through this process. Lockdown has affected everyone, being stuck at home, being cut off from friends, it all takes a heavy toll. This is a time when we need extra support from each other. Unfortunately a thesis also demands time and balance is difficult to achieve. We have fortunately persevered and its time for me to pay back some debts.

Table of Contents

List of Figures	viii
1 Introduction	1
1.1 Containers	3
1.2 Ownership Types	4
1.3 Outline	5
2 A Language with Containers	7
2.1 Grammar	9
2.2 Container Labels	12
2.2.1 Paths	13
2.3 Methods and Constructors	15
2.3.1 Generic Containers	16
2.3.2 Container Inferencing Algorithm	18
2.4 Self Containment	21
2.4.1 Object Copying	22
2.4.2 Equality Comparison	24
2.4.3 Serialization	25
2.5 Demonstration System	25

3	Containers Formalized	27
3.1	Types	28
3.2	Language	32
3.3	Typing Environments	33
3.4	Type Rules	33
3.4.1	Type Rule Helper Functions	35
3.4.2	Type Validation	39
3.4.3	Language Typing	43
4	Operational Semantics	48
4.1	Overview	48
4.2	Heap	49
4.3	Frames	50
4.4	Configuration	53
4.4.1	Configuration Typing Environment	54
4.5	Supporting Functions	57
4.6	Frame Typing	59
4.6.1	Program Type Annotations	59
4.7	Special Configurations	62
4.8	Transition Rules	63
4.8.1	Machinery	63
4.8.2	Statements	63
4.8.3	Expressions	64
4.8.4	Parameter Passing and Symbol Initialization	67
4.8.5	Field Initialization	67
4.8.6	Object Copy Rules	70
4.8.7	Null De-Reference Guards	72
4.8.8	Decomposition Rules	74
4.9	Configuration Typing	75

5	Type Safety	83
5.1	Progress and Preservation	83
5.2	Heap Validation Lemmas	85
5.3	Frame-Stack Lemmas	89
5.4	Logical to Physical Consistency	97
5.5	Preservation Lemmas	99
6	Related Work	104
6.1	Ownership Types	104
6.2	Serialization	106
6.3	Operational Semantics	107
7	Future Work	108
7.1	Sub-Containers	108
7.2	Roles	109
7.3	Data Flow Analysis	111
8	Conclusion	113
	References	115

List of Figures

1.1	Container Diagram	3
1.2	Containment vs. Encapsulation	6
2.1	Expression, Statement and Container Label Grammar	10
2.2	Declaration Grammar	11
2.3	Container Labels	13
2.4	Variables as Container Labels	14
2.5	Paths as Container Labels	15
2.6	Method Calls with Containers	16
2.7	Generic Container Labels	18
2.8	Pseudo-code for Container Label Inferencing	20
2.9	Deep Copy On Assignment	22
2.10	Pseudo-code for Deep-Copying a Self-Contained Tuple	23
2.11	Pseudo-code for Deep Object Comparison	24
3.1	Representation of Container Types	29
3.2	Formal Type Examples	30
3.3	Abstract Language Definition	32
3.4	Global and Local Typing Environments	33
3.5	call () Related Typing Functions	34
3.6	Field Access Related Typing Functions (1)	35

3.7	Field Access Related Typing Functions (2)	36
3.8	Container Labels for Field Access Expressions	38
3.9	Mutability and Mobility for Field Access Expressions	38
3.10	Path Validation	39
3.11	Container Label Validation	41
3.12	Type Validation	42
3.13	Class Validation	43
3.14	Label Map Validation	43
3.15	Assignment and Initialization	44
3.16	Expression Typing	45
3.17	Statement Typing	46
3.18	Function Validation	47
4.1	Configuration Overview	49
4.2	Heap Structure	50
4.3	Closed Frames	51
4.4	Open Frames	52
4.5	Configuration	53
4.6	Configuration Typing Environment	55
4.7	Container Mappings Example	56
4.8	Operational Rules Support Functions (1)	57
4.9	Operational Rules Support Functions (2)	58
4.10	Frame Typing (1)	60
4.11	Frame Typing (2)	61
4.12	Initial and Terminal Configurations	62
4.13	Transition Rules Machinery	64
4.14	Transition Rules Statements	65
4.15	Transition Rules Expressions	66

4.16	Transition Rules	Symbol Initialization	68
4.17	Transition Rules	Field Initialization (1)	69
4.18	Transition Rules	Field Initialization (2)	70
4.19	Transition Rules	Copy Tuple (1)	71
4.20	Transition Rules	Copy Tuple (2)	72
4.21	Transition Rules	Null Pointer Violations	73
4.22	Transition Rules	Decompositions (1)	74
4.23	Transition Rules	Decompositions (2)	75
4.24	Configuration Typing Rules	(1)	76
4.25	Configuration Typing Rules	(2)	78
4.26	Local Configuration Typing		79
4.27	Frame-Stack Typing		81
7.1	Sub-Container Notation		109
7.2	Data Flow Example		111

Chapter 1

Introduction

A container's essential property is the ability to prevent its contents from escaping. In the context of object oriented-programming languages, a container would prevent objects within the container from referencing an object outside the container. Within a container, there could exist an arbitrarily complex graph of objects, but there is no escape. Objects themselves can be containers, and we'll call an object which is a container a self-contained object.

With a container mechanism, there is an opportunity to simplify reasoning about complex programs. In addition to a container mechanism, if there were also restrictions on IO, then when calling a method of a self-contained object with read-only parameters, there would be absolute certainty that all side effects will be limited to inside that object. In complex applications, the relationships and interactions amongst objects may not be readily apparent, and by using containers, operations to an object can be typed such that there is no possibility of unexpected side effects. Complexity will still exist, but if the simple self-contained objects are typed as such, it will discourage software developers from creating undesirable entanglements.

Containment can be thought of as the converse of encapsulation. Encapsulation isolates private fields of an object from external meddling. It becomes easier to reason about an object's internal implementation when you know that there are no external readers or writers of the fields you are manipulating. The inside is protected from the outside with no incoming references. Conversely, with containment you can invoke a method on a contained object free from concern that this method could affect another critical object. The outside is protected from the inside with no outgoing references.

Although encapsulation features are far more common in programming languages than

containment, containment is heavily used in computer systems. Perhaps the most well-known use of containers is the virtual machine. A single physical computer can run multiple virtual machines with complete isolation, as if they were running on separate hardware. If an administrator were to attempt to run many services without this separation, there would be significant risk of downtime due to simple things like a patch to one application requiring that the system reboot. Processes can also have issues with contention for operating system resources, e.g., attempting to listen on the same network port number. These issues disappear when services can be isolated into their own virtual machines.

However, virtual machines have other benefits beyond just isolation. With a defined container, it becomes easy to clone a virtual machine. It is common practice to take snapshots of virtual machines at regular intervals. If something catastrophic occurs, the machine can quickly be restored. Physical hardware, on the other hand, is more difficult to back up with the same precision. The virtual machine clearly defines what the machine is. For example, hardware BIOS settings are part of the virtual machine's definition and this is something that is not typically saved when a backup is done for a physical machine. Now, thinking in terms of a programming language, our goal is to show that the simple act of defining an enforced container boundary makes it easy to copy, serialize and restore objects.

Drawing once more from our analogy with virtual machines, we can see the role virtual machines play with scalable concurrency. Elastic scalability is the ability to quickly add and remove resources for an application depending on present demand. Here, the virtual machine is a unit of deployment. Because of its self-contained nature, virtual machines can be replicated on-demand to add servers and likewise instances can be shutdown when needed. Again, we can see parallels in programming languages with significant research efforts working with concurrency models that do not use shared-mutable memory. Isolating threads from each other can greatly simplify a concurrent program.

Some examples of concurrency models that avoid shared memory are message passing models, actor models and map-reduce frameworks. One specific example is the use of channels and go-routines in the Go programming language. These features of Go provide a convenient way to coordinate threads. However, nothing in this system prevents the sharing of mutable memory across a channel. Of course a good programmer will know not to do that, but similarly a good programmer also knows to pass parameters with the correct types. The latter mistake will be caught by static typing. This can help to catch the errors that humans inevitably make. But, if we had a language that could enforce container boundaries, then we could use the type system to prevent a message from containing external references. The type system could prevent the sharing of memory across threads.

Figure 1.1: Containers: The diagram shows objects within containers (circles). The dotted red lines represent illegal references, which escape their container. Inward references (in yellow) are legal.

The goal of this thesis is to demonstrate how a container concept can be added to an imperative object-oriented language and to further show how containers can help solve issues such as the deep-vs-shallow-copy problem, serialization and deep comparison. We can also leverage containers to encourage pass-by-value semantics, reducing the unwanted aliasing that commonly appears in languages such as Java and Python where pass-by-reference is the norm.

1.1 Containers

For an object to be within a container c means that any field of the object must only reference objects which are also contained within c . Typically, a container is an object, but we also propose that a container could be a scope. For example, the local scope of a

function invocation could be a useful container for temporary objects. When an object is constructed, it is placed into a specific container that can never change.

In addition to objects existing in containers, every reference will be typed with a container constraint which we call its container label. This indicates the container of the object being referenced. For example, a container label could specify that a reference is constrained to only point to objects contained within local variables. We can see that container labels lead to a dependent type system. Every label requires its environment in order to be properly interpreted. For references that are fields in an object, the container label must not violate the container boundary of the object. It can be narrower than the container of the object, but never outside of the container.

We differentiate two kinds of objects in our system. Data objects are self-contained, meaning that no matter which container the object is placed into, fields of that object must not reference outside of the object. Entity objects are less strict; here, the container the object is placed in is the outer bound for references. We'll show that self-contained objects have useful properties, making the distinction worthwhile.

In order to enforce container rules, the typing system must check the container labels on every assignment and passing of parameters. Passing parameters to functions requires a careful consideration of scoping, as method parameters are declared in a different scope from the invocation of the method. Typing a function call is where most of the complexity of this system exists.

In the system developed in this thesis, there is significant tracking of containers at run-time. However, this exists for the purpose of verification of the type system. The goal of the type system is to statically check all references for proper containment. In doing so, the run-time overhead of containers can be zero.

1.2 Ownership Types

Earlier, we compared containment with encapsulation. Continuing this comparison, we highlight an area of research called ownership types. Chapter 6 discusses specific related work, but for now, we will outline key similarities and differences. With ownership types, objects are assigned an owner, which directly corresponds with our placing an object into a container. Where these systems differ is in their purpose. Ownership types have the concept of a dominator, where the owning object has exclusive control over the owned object. In our work with containers, we do not have this concept. Containers permit

outside aliasing of contained objects. We lose the encapsulation benefits of ownership types, but we are free to focus on the containment concept in isolation.

```
class Demo {
  // The ':: self' notation indicates that reference r refers to an object
  // that is 'within' the instance.
  // In an ownership type system, we would say that r is owned by the object
  // In this thesis, we say r is contained by the object
  ref r : Node :: self;
}

method Main.demo() {
  ref demo = Demo(); // Construct a new instance of Demo

  ref inner = demo.r; // This would be a compile error in an ownership-type system.
                    // The field r is owned by demo, which encapsulates it.
                    // However, in our containment system this line is legal.
                    // We do not enforce encapsulation.

  ref demo2 = Demo(); // Construct another new instance of Demo

  demo.r =^ demo2.r; // Reference assignment of the field r from demo2 to demo
                   // This line is illegal in our containment system, but
                   // encapsulation is not the reason.
                   // This is a containment violation.
                   // The field r in the demo container is not permitted to
                   // reference an object in the demo2 container.
}
```

Both ownership and containment use dependent type systems. In terms of type system mechanics, these two systems are very similar. Having a type be dependant on an owner versus a container makes little difference from a dependent-types perspective. The primary distinction between these two systems are the goals. We allow encapsulation violations in order to examine the potential of containment without the restrictions of ownership types.

1.3 Outline

We begin in chapter 2 by describing a language with container types. The grammar is introduced and code examples help explain the various behaviors. The container labels are fully specified and explained. Special attention is given to the handling of method calls and the use of generic container labels. Also detailed is a container label inferencing system which allows explicit specification of labels to be omitted in many cases. The language described has a working interpreter written in Haskell. The chapter ends with notes on the usage of this interpreter.

Figure 1.2: Containment vs. Encapsulation

Chapter 3 describes a formalized version of the language. The formal version is reduced in scope, but retains the key complication of managing container labels across scope boundaries. A full description of how types are composed is given, including specifying immutability and the container label. The grammar for this language is abstract and there is no container inferencing in this version. Finally, a typing environment is defined as well as a number of type rules to type the language in general, but also many specifics of working with container labels.

Chapter 4 defines the operational semantics of the formal container language. A system of small-step operational semantics is defined running in a configuration that tracks container properties. The semantics extend the language with a number of additional frames needed to execute more complex operations such as tuple initialization, parameter passing and copying objects. Type rules are provided for the additional frames. Additionally, type rules are provided to fully type the entire configuration. This includes the heap, frame stack and variable stack.

Chapter 5 takes the typings defined in chapter 4 and proves their soundness with the traditional progress and preservation theorems. There are a set of lemmas establishing that in this dependent type system, the elements that are depended upon are immutable, therefore typing will be preserved. The most important lemmas work with passing symbols across scopes to make sure that two distinct but corresponding container labels refer to the same physical object in the configuration.

Chapter 6 outlines how this thesis relates to other areas of research. And finally, chapter 7 describes areas where this work can be improved and extended.

Chapter 2

A Language with Containers

This chapter describes a demonstration language that implements a simple object oriented language which uses containers. We'll begin introducing the language by example. In the first example, we declare a reference to an instance of a class called demo. The extra syntax "`:: local`" constrains the references to only point at objects inside the container indicated by `local`, which refers to the scope of the current function. We call the container specification a container label

```
ref r : DemoLL :: local =^ null; // Operator =^ means assign reference
```

Objects can be containers as well and in most cases are more useful containers than the local scope. A container groups a set of objects that can work together (reference each other) and if you need to temporarily create such a group, then a `local` container label is useful. In our second code snippet, we define a class that we'll use to show how objects can be containers.

```
entity class DemoLL( id : Int, ref next : DemoLL :: container )
constructor DemoLL( id : Int ) {
  self.id = id;
  self.next =^ null; // For clarity, redundant since null is default.
}
```

Here, we define the class `DemoLL` which is a simple node of a linked list with two fields. The `next` field is declared with the label `container`. So far, we've seen a container label associated with a reference, but when an instance `DemoLL` is created, it is also placed in a container. The `next` field is given the container label `container` which refers to the container that the object was placed into. This means that the `next` field can only reference objects in the same container as the object. Also, note the `entity` keyword in

the class declaration. This modifier indicates that the class is permitted to have references to other objects within a common container. By default, a class is a data class, which means it is self-contained. The label `self` container is illegal in a data class because a reference with that label would violate self-containment. Now, we'll continue the example and make use of the `DemoLL` class in the next code snippet.

```
class Main ( fixed ref list : DemoLL :: self )
constructor Main() {
  self.list =^ DemoLL(1);
  self.start.next =^ DemoLL(2);
}
```

We have a program which begins by instantiating the class `Main`, invoking its constructor. `Main` is not declared with entity `self`, so it is a self-contained class. Accordingly, its field `list` is declared with the container label `self`. Two nodes are created for our list and both are placed in the container `self`, which is the instance of class `Main`. When the constructor call `DemoLL(1)` was made, how did it know which container to place the new object into?

Container label inferencing allows the container label on the right-hand side of an assignment to be inferred based on the expectations of the left-hand side. In the example, `self.list` is declared with `:: self` and with this label on the left-hand side, there is only one correct label choice for placing `DemoLL(1)` into a container. Similarly, on the next line, `self.start.next` also resolves to the label `self`. The field `next` is typed with label container, which in this context refers to the container of `self.start`.

For containers to be a useful feature, the type checking must catch all container violations at compile time. We'll look at a simple error example next.

```
class Main ( ref list : DemoLL :: self )
constructor Main() {
  self.list =^ DemoLL(1);
  ref otherList :: local =^ DemoLL(2);

  self.list.next = otherList;           // COMPILER ERROR!
}
```

By placing the second instance of `DemoLL` into container `:: local`, it is segregated from objects in `:: self`. The attempt to assign to `self.list.next` is detected as a container violation. This type checking can prevent accidental mixing of objects that have the same class, but aren't intended to be used together. As an example, if you were assigning a partner to a police officer, you would want them both to be based in the same precinct. All data for a precinct could be in its own container object and the type system would ensure that any partner assignment operation respects the container boundaries. If a police management application were structured this way, it would be impossible to assign an officer a partner from a different precinct.

So far, we have been focusing on references, but there are also value types and indirect reference types. Indirect references are denoted `iref`.

```
class Main ()
constructor Main() {
    var x = 1;
    iref r =^ x;    // r's type is inferred from its initializer
                  // Assign through the indirect reference
    r = 2;         // Outputs 2, x has been modified
    print x;
    x = 3;
    print r;      // Outputs 3

    var y = x;    // Non-references (val & var) are always assign-by-value
    x = 4;
    print y;     // Outputs 3, y is fully independent from x
}
```

Value types are declared using `val` or `var`. Only self-contained data classes are permitted to be values. Entity classes are entangled with their environment and considered unsuitable for pass-by-value semantics, so they must be handled through references. Although the above example used the primitive type `int`, the assign-by-value rules are the same for any self-contained class. When the variable is initialized, a full deep copy of the initializer object is made.

With that quick introduction to the language, we will now take a deeper look at the language specification.

2.1 Grammar

Figures 2.1 and 2.2 detail the grammar of the language. Starting with figure 2.1, a number of binary operators are implemented. Most are self explanatory, but `or=` and `==^`, the former tests value equality and the latter tests reference equality. Instances of any self-contained classes can be compared for equality, no matter how complex. The algorithm for this comparison is explained in section 2.4.

Expressions include field accesses using a dot as in Java and other languages. However, we have both value and reference types in our language. In general, the semantics of the language is to automatically de-reference when needed. Field access, passing a value to a method and value assignment are all value contexts. Any time a value is expected, if the input is a reference, then an automatic de-reference is done. In contrast with C++, we do away with the `->` and `*` operators, but add `=^` to distinguish reference assignment from value assignment. Note that if the left-hand side of an assignment is `iref`, then

bin-op ϕ + | - | * | / | or | and | xor | > | == | ==^

ϵ Bool | Integer | String

ϵ e bin-op e

ϵ e : eld-name

ϵ e : method-name(e ; ::: ; e)

ϵ class-name(e ; ::: ; e)

ϵ var-name

ϵ self

ϵ null

$\$$ if e f s g ;

$\$$ while e f s g ;

$\$$ return e ;

$\$$ e = e ;

$\$$ e =^ e ;

$\$$ val var-name = e ;

$\$$ var var-name = e ;

$\$$ [read-only][fixed] ref var-name [:: container-label] = e ;

$\$$ [read-only][fixed] iref var-name [:: container-label] = e ;

$\$$ save e e ;

$\$$ load e e ;

$\$$ print e ;

path self

path var-name

path path . eld-name

container-label path

container-label container

container-label local

container-label ` generic-name

container-label unknown

Figure 2.1: Expression, Statement and Container Label Grammar

```

program decl
decl [entity ] class class-name( eld-decl; :::; eld-decl )
decl [read-only ] methodclass-name method-name( parm-decl :::; parm-decl) ret-decl
decl constructor class-name( parm-decl :::; parm-decl)

parm-decl [val ] eld-name : class-name
parm-decl var eld-name : class-name
parm-decl [fixed ] [read-only ] ref eld-name : class-name[: container-label]
parm-decl [fixed ] [read-only ] iref eld-name : class-name[: container-label]

eld-decl val eld-name : class-name
eld-decl [var] eld-name : class-name
eld-decl [fixed ] [read-only ] ref eld-name : class-name[: container-label]
eld-decl [fixed ] [read-only ] iref eld-name : class-name[: container-label]

ret-decl : class-name
ret-decl [read-only ] ref : class-name[: container-label]
ret-decl [fixed ] [read-only ] iref : class-name[: container-label]

```

Figure 2.2: Declaration Grammar

both reference assignment and value assignment are possible. With this extra operator, it is always clear whether a value or reference is needed.

Method/constructor calls and return statements behave as one would expect, but with additional consideration of containers. Section 2.3 will detail the handling of containers as parameters are passed and return values are received. The expression `self` functions like the `this` keyword in familiar languages like C++ and Java. We use the term self-contained extensively and it was natural for the language to match.

Like many other languages, statements include variable declarations, assignment, and flow control statements like `if`, `while`, `return`. In the examples, we've seen variable declarations of the various kinds. Also included is `print` statement to display an object's contents and `save/load` statements which can save a self-contained object to a file and then load it back into a variable that will compare as equivalent to the original. These three IO statements are discussed further in section 2.4.

Also in figure 2.1 is the full definition of the container labels. We have already seen examples of `local`, `self`, `container`, and the remaining labels will be discussed in the next section 2.2.

In figure 2.2, the remainder of the grammar is defined. A program is a series of dec-

larations which can be classes, methods or constructors. Nested within the top-level declarations are individual parameter declarations and in the case of methods a return type declaration.

Classes are data classes by default, meaning they must be self-contained. The optional `entity` keyword enables the declaration of references that point outside of the object. When `elds` are declared, the default kind is `var` and that keyword can be omitted. The container label of references can be omitted as well with the default being `self`. Reference `elds` can also be typed as `fixed` and/or `read-only`. The meaning of `fixed` is that the identity of the referred object cannot change and reference assignment is not allowed with this symbol. For a `read-only` reference, you are not permitted to modify the referred object and value-assignment statements are not allowed. Note that `read-only` doesn't guarantee immutability, as another alias could be able to modify the object.

When declaring parameters of methods and constructor calls there are also defaults, but they are different. If a container label is omitted from a parameter or return value then the default is `unknown`, which is useful when the method doesn't care what container the object is in. If no kind is specified for a parameter, then `val` is the default.

When values are returned, they are always immutable. Since value assignment requires a copy, the mutability of an expression result is of no consequence. Attempting to reference-assign an expression result value (r-value) to a reference is illegal.

2.2 Container Labels

In this section, we explain in more detail the full set of container labels. The grammar is repeated in [figure 2.3](#) for convenience. The simplest form of path container label is a local variable or the special keyword `self`. These symbols directly indicate a container object. Since the container is an object and we know its type, a path can be extended through `eld` accesses to indicate a more precise container. [Section 2.2.1](#) will explain paths further.

Each of the container labels we've seen so far specify a container relative to the current scope. A mechanism is needed to relate containers as they are passed to a function. The generic container labels provide the solution. A function can declare its reference parameters as generic, which enables its implementation to enforce abstract container constraints without knowing the identity of the containers. [Section 2.3](#) provides a deeper look into generic container labels.

We've seen the label `container` in our previous examples. When this label is used in the type of a `eld`, it refers to the container of the object holding the `eld`. When used

in a method, it refers to the container of the implicit self parameter. When used with parameters, the `labelcontainer` behaves just like a generic label. It can be modeled as an implicit `self :: `container-of-self` declaration. The `labelself` has a similar dual meaning depending on context. For fields, it represents the object holding the field and for parameters, it represents the `implicitself` parameter.

The `labelunknown` is typically used implicitly by omitting the `::` for a parameter. The normal use case would be a reference parameter of a method. Container labels are only needed when two or more objects are related to each other by a common container. A method may call methods on another object without knowing its container as long as the called method doesn't have a parameter typed with `container`. If the called method did have such a parameter, then our outer method would be unable to pass a parameter with the correct label.

<code>container-label</code>	<code>path</code>		
<code>container-label</code>	<code>container</code>		
<code>container-label</code>	<code>local</code>	<code>path</code>	<code>self</code>
<code>container-label</code>	<code>`generic-name</code>	<code>path</code>	<code>var-name</code>
<code>container-label</code>	<code>unknown</code>	<code>path</code>	<code>path.eld-name</code>

Figure 2.3: Container Labels

2.2.1 Paths

Figure 2.4 shows an example of a local variable as a container label. Notice that the variables `x` and `y` are declared as `fixed`, meaning that they can't be re-assigned. The code of this example would not compile otherwise and you would receive a compile error `that =^ r2`. The reason the error would be raised is that `r2` could have been re-assigned between the initialization of `r1` and `r2`. Without full data-flow analysis, the compiler doesn't know that `r1` and `r2` reference the same container. With `path` container labels, every component of the path must be `fixed`. With explicit container labels, you will receive an error if you attempt to declare an unboxed label. If you declare a reference without an explicit label and the inferred label is unboxed, then the reference is implicitly declared with `labelunknown`. This stability requirement means that in the dependent type system, all dependencies are immutable and types are preserved.

The example of figure 2.4 continues and declares a variable `z` and then attempts to mismatch containers. The compiler will not allow `r1 =^ y.v1` because `r1` can only reference objects contained with `x`.

```

class Info( id : Int, v1 : Int, v2 : Int )
constructor Info( id : Int, v1 : Int, v2 : Int ) {
  self.id = id;
  self.v1 = v1;
  self.v2 = v2;
}

class Main ()
constructor Main() {
  fixed var x = Info( 1, 10, 100 );

  iref r1 :: x =^ x.v1;      // The local variable x is the container
  iref r2 =^ x.v2;         // The container is inferred as x based on rhs
  r1 =^ r2;                // Legal because the containers match

  fixed var y = Info( 2, 20, 200 );
  r1 =^ y.v1;              // COMPILER ERROR!
}

```

Figure 2.4: Variables as Container Labels

Figure 2.5 contains a more complex example of paths using field accesses. References `r1` and `r2` have matching containers as before, but this time the path to the container has been extended to include a field access. Container labels can be declared where the container isn't directly in the current scope. In this case, we reach inside the variable to find the container.

As mentioned before, every element of a path must be immutable. The local variable `x` is declared as `fixed` as well as the fields `i1` and `i2`. This means that the expression `x.i1` always evaluates to the same container and we can rely on it as a type. Similar to the last example, Figure 2.5 also shows a compile error where there is an attempt to reference-assign an object contained in `i2` to a reference typed with label `i1`.

Generally, container labels must be an exact match when doing a reference assignment. One exception is that `null` can be assigned to any reference. A second exception is when a reference is declared with label `unknown`, which allows for the reference to point within any container. There is no reflection mechanism to narrow a reference typed with `unknown` label into a specific container label and this means that you are limited in what you can do with an `unknown` container label. Fields are not permitted to use the `unknown` label because a field is always restricted to the container of its object.


```

class Inner( id : Int, v1 : Int, v2 : Int )
constructor Info( id : Int, v1 : Int, v2 : Int ) {
    self.id = id;
    self.v1 = v1;
    self.v2 = v2;
}

class Outer( fixed ref i1 : Inner, fixed ref i2 : Inner )
constructor Info( id1 : Int, id2 : Int ) {
    self.i1 =^ Inner( id1, id1 * 10, id1 * 100 );
    self.i2 =^ Inner( id2, id2 * 10, id2 * 100 );
}

class Main ()
constructor Main() {
    fixed var x = Outer( 1, 2 );

    ref r1 :: x.i1 =^ x.i1.v1;      // The field i1 of local variable x is the container
    iref r2 =^ x.i1.v2;           // The container is inferred as x.i1 based on rhs
    r1 =^ r2;                    // Legal because the containers match

    r1 =^ x.i2.v1;              // COMPILER ERROR: r2 is distinct from x.i1
}

```

Figure 2.5: Paths as Container Labels

2.3 Methods and Constructors

The checking of container labels for parameters and return values of methods requires a correspondence between symbols in two different scopes. We'll begin with an example shown in figure 2.6. We have a method called `test()` which takes a fixed reference `r1` as its first parameter and this serves as the container we'll be focusing on. The second parameter `r2` has `r1` as its container label so we know that `r2` is within `r1` and if we modify `r2`, then part of `r1` will change as well. The method updates the value referenced by `r2`, then dumps the entirety of `r1` to reveal what has changed.

Now, looking at each of the calls to `test()`, we can see that the first two correct cases pass a second parameter which is contained within the first parameter. And, to demonstrate the error case, the third example mismatches `x` and `y`.

When type checking a method call like this, a mapping is built for all of the labels in the function call's parameters. First, when processing the first parameter, the container label `r1` from the method's scope is associated with the container label `x` in the caller's scope. Then, when type checking the second parameter, we can use the mapping established by the first parameter to take the declared container label `r1` of the second parameter and

```

class Info( v1 : Int, v2 : Int )
constructor Info( v1 : Int, v2 : Int ) {
    self.v1 = v1;
    self.v2 = v2;
}

class Main ()
method Main.setTest( fixed ref r1 : Info, iref r2 : Int :: r1, v : Int ) : Int {
    r2 = v;           // Assign through the indirect reference
    print r1.v1;
    print r1.v2;

    return 0;       // Currently the language doesn't permit returning void
}

constructor Main() {
    fixed var x = Info( 1, 10 );

    self.setTest( x, x.v1, 2 );           // Outputs 2, 10
    self.setTest( x, x.v2, 20 );         // Outputs 2, 20

    fixed var y = Info( 5, 50 );
    self.setTest( x, y.v1, 6 );         // COMPILER ERROR! y is not in x
}

```

Figure 2.6: Method Calls with Containers

map it into the caller's scope. With this mapping complete, we can type check the second parameter, which must have label `x`.

The first call to `setTest()` modifies `v1` and the second call modifies `v2`. Because of the explicit specification of container, the side effect of `r1` being changed when `r2` is modified should not be a surprise. Perhaps more important than knowing what will change after an operation is certainty about what cannot change. It is impossible for the line `r2 = v` to modify anything other than the object `r1`.

When returning references, the same mapping process occurs. A method could return a reference contained within any of its parameters or the implicit `self` parameter. You can also pass and return references contained within the container of any parameter. We'll see how this can be done with generic container labels.

2.3.1 Generic Containers

In the previous example in figure 2.6, we declared a container label to be a previously declared parameter, but in many cases there is no need to pass the container itself, just the

pieces we intend to interact with. With generic container labels, we can pass a parameter with an abstract container label. Figure 2.7 demonstrates this capability. Here, we've implemented a linked list insert method twice, and each implementation has a distinct container label. Now, in `Main`, we create a linked list with 3 elements in it. First, notice that when we call the constructor of `DemoLL` we never explicitly need to place it in a container; it happens automatically. In the first construction, the container label is inferred from the left-hand side of the assignment as we've seen before. However, for the inserted elements, the containers for the newly constructed nodes are inferred based on the expected parameter types of the method calls. This is interesting because the parameter types are generic and also must be resolved.

The next section will outline the inferencing algorithm, but for now we will reason through the inferencing steps specific to this example. With the method `self.insert()`, the generic container label `c` appears three times. As long as one of those instances has a definite container label in the method call, then the other two instances become fixed and their definite labels can be used for further inferencing.

A constructor's return type doesn't appear visually in the code, but to the inferencing system, every constructor returns a reference with a container label of `container`. This label becomes a concrete container at the call site by the same inferencing techniques used to resolve generics in the parameter labels.

Returning to the call to `self.insert()`, we see that its first parameter has the container label `:: self` which then becomes the label of the second parameter. Next, the constructor call to `DemoLL(2)` has an expected return type with label: `self` which fully resolves the labels for the construction. Last, the return value of `self.insert()` is also contained within `:: self`.

The insertion of the third list element uses a different insert method, but the method's container labels map to the same containers in the caller's scope, allowing it to be compatible. The final call to `DemoLL.insert()` is a method call on the return value from `Main.insert()`, which in this example is the second node of the list with the container label `:: self` which refers to the instance of `Main`. The parameter `newNode` of `DemoLL.insert()` is typed with container label `:: container` and the mapping of this label into the caller's scope is inferenced using knowledge of the implicit `self` parameter's container label: `self`. As before, the return container label for the call to `DemoLL(3)` can be determined since we now know the label for the parameter to `DemoLL.insert()`.

This is a contrived example and the inferencing process for generics is complex. However, from a programmer's perspective when implementing a method, the parameter labels are abstract and there is no need to be concerned about how they will later resolve. As a

```

entity class DemoLL( id : Int, ref next : DemoLL :: container )

constructor DemoLL( id : Int ) {
  self.id = id;
}

method DemoLL.insert( ref newNode : DemoLL :: container ) ref : DemoLL :: container {
  newNode.next =^ self.next;
  self.next =^ newNode;
  return newNode;
}

method Main.insert(ref node : DemoLL :: `c, ref newNode : DemoLL :: `c) ref : DemoLL :: `c
{
  newNode.next =^ node.next;
  node.next =^ newNode;
  return newNode;
}

class Main ( fixed ref list : DemoLL :: self )
constructor Main() {
  self.list = DemoLL(1);
  self.insert( self.list, DemoLL(2) ).insert( DemoLL(3) );
}

```

Figure 2.7: Generic Container Labels

programmer calling a method, some care needs to be taken to ensure that the parameters have matching container labels, but there is no need to understand every detail of how the mapping between caller labels and callee labels is established.

2.3.2 Container Inferencing Algorithm

A goal of this project was to minimize the syntactic burden of specifying container labels and an inferencing algorithm was developed to allow the container label to be omitted in many cases. A bi-directional typing algorithm was developed to infer labels for variable initialization, assignment statements and method/constructor calls. Note that inferencing only occurs within a single statement. When each expression is typed, an expected container label is provided and this assists with the typing of constructor calls and methods that return generic labels. We'll look at the method call reasoning in detail.

When a method call is typed, we have multiple sources of information; the expected return container label is provided by higher-level code, typing information is in the method declaration. We also have the expressions passed to the method from the program's ab-

stract syntax tree. Specifically, we have an expression for the implicit `self` parameter and each of the explicit parameters.

For each expression passed to the method call, we designate a fresh generic container label symbol. Then, we type each of the parameter expressions using the respective generated generic labels as the expected container label for the expression. E.g. `typeExpr (expr, generated-generic-label)`. Often, expressions will immediately be typed with a concrete container label, but if an expression type comes back with the same generic type we passed in, then we know that the expression is free in its container label. At this point in the process, we have a list of expressions as well as preliminary typings for each of them. Note that after the inferencing process completes, we will re-type these expressions.

Next, we need to deal with the fact that the method is declared with container labels relative to the method's own scope. We need to take the method signature and associate the labels declared with each parameter with labels in the calling scope. To do this, we first use a function called `expressionToContainerLabel ()`, which can take each of the parameter expressions and determine what the container label should be if that expression result is used as a container. Recall the example in [figure 2.6](#), where the second parameter was contained within the first parameter. If a situation like that occurs, then `expressionToContainerLabel ()` tells us what the appropriate label should be in the caller's scope. We now have a mapping from parameter names to caller-centric labels. Using this mapping, we map the method's declared container labels into the scope of the caller.

The inferencing can now be done in the container label space of the caller. [Figure 2.8](#) presents pseudo code for this operation. For simplicity, we've attenuated the inputs to just two lists of labels: `decl-labels` and `passed-labels`. These represent the labels for each passed parameter, based on their preliminary typing, as well as the expected return label. The two lists are aligned to each other, meaning `index` in each list represents the `th` parameter.

The inferencing consists of three steps. First, we create a map with the domain being the set of generic container labels declared in the method signature. These generic container labels should not be confused with the temporary unique-free-generics that we injected when we typed the parameter expressions; these generic labels are from the method signature. Each generic label may appear in the method signature more than once, and for each generic, we record the passed-in container labels corresponding to the generic label. The function `findGenericParameters` implements this operation in [figure 2.8](#). The result is a map from generic labels to a list of the passed container labels.

For each generic mapping, we merge the list of labels into a single label. The function `mergeAllLabels` does the merging. In this process, the injected free-generics are replaced

```

inferGenericParameters( free-generics , decl-labels , passed-labels ) =
  let
    // Map each declared generic label to a list of passed labels
    generic-map = findGenericParameters( zip( decl-labels , passed-labels ) )

    // For all generics, merge the occurrences
    mergedGenericMap= genericMap.mapRange( mergeAllLabels )

    // If any free generics remain after merging then the program
    // has no concern for which label is chosen so we use - local
    mergedGenericMap = mergedGenericMap.substRange( label 2 free-generics -> local]

findGenericParameters [] = {}
findGenericParameters( hdecl-label ; passed-label i tail ) =
  let
    generic-map = findGenericParameters( tail )
  in
    if decl-label is generic then
      if decl-label 2 dom(generic-map) then
        generic-map.set( decl-label ?! passed-label generic-map.get( decl-label ) )
      else
        generic-map.set( decl-label ?! [passed-label ] )

mergeLabels( unknown-cont, label2 ) = error: unknown-cont is incompatible with label2
mergeLabels( label1 , unknown-cont ) = error: unknown-cont is incompatible with label1
mergeLabels( null-cont, label2 ) = label2
mergeLabels( label1 , null-cont ) = label1
mergeLabels( label1 , label2 ) =
  if label1 2 free-generics
  if label2 2 free-generics
    label1 // Both are free, keep label1
  else
    label2 // Use the non-free label
  else if label2 2 free-generics
    label1 // Use the non-free label
  else if label1 == label2 then
    label1
  else
    error Labels label1 and label2 are incompatible

mergeAllLabels( [ label1 ] ) = label1
mergeAllLabels( label1 tail-labels ) =
  mergeLabel( label1 , mergeAllLabels( tail-labels ) )

```

Figure 2.8: Pseudo-code for Container Label Inferencing

with container labels from the caller's scope. Errors are raised if there are inconsistencies. The final result is that we have a concrete container label for every generic container label in the method signature.

With this map in place, the typing of the method call can be completed. As a final step, each parameter expression is re-typed now that we know the exact container label needed for each expression. This is how we can call a method and pass a constructor call as one of the parameters. Once the specific parameter knows its container label, then the constructor call is re-typed, which sets the container for the newly constructed object.

With this inferencing logic in place, most container specification can be eliminated from code. Declarations of methods and fields will need explicit syntax, but there are defaults in place that can help there as well. Any proposed language feature faces a cost/benefit analysis and the inferencing logic was developed to reduce the overhead cost of using container types.

2.4 Self Containment

In this section, we look at four capabilities of self-contained objects. The language allows you to declare data classes and entity classes, which mean self-contained and non-self-contained. This distinction allows for extra capabilities to be added only for self-contained objects.

An entity is entangled with its environment. Imagine a machine that could clone a person. After one person becomes two, then which one will receive a pay check from their employer? Which one gets to live with their spouse? The point of this example is that copying an object that has relationships to its environment is complicated. Logic to copy an entity class needs to be customized for that particular class. Our language takes the opinion that entities should never be copied automatically. Similar to Java, you can write a clone method if you desire, but there is no language support. This can be seen by the rule that entity objects are always managed by reference, never by value.

Self-contained objects, on the other hand, are pure data. It is natural to copy them because they have no entanglements with their environment. One novel aspect of this system is that inside a data object, there can exist entity objects with complex relationships. Entity objects are copied as part of copying the outer data object, but an entity is never copied by itself. The well-defined boundary means we can recreate an identical graph of the inner entities.

Figure 2.9: Deep Copy On Assignment: Assigning by value creates a new disjoint object.

Contrast this to a language like C++. If you built a class using shared smart pointers, you will get a shallow copy from the default copy constructor. If you use direct containment, then then you get a deep copy. There is no in-between. If you have a nested object with multiple aliases, then there is no immediate way to copy the nested object only once and have every alias in the copy point at the new nested object. Of course, you can implement this yourself, but there is no automatic support. Next, we'll explain how we can solve the deep copy vs. shallow copy problem for self-contained objects.

2.4.1 Object Copying

By declaring a class as self-contained and having the type system enforce that containment, you can be certain that every object reachable from the self-contained object belongs to the object and should also be copied. After a deep-copy operation, the resulting object is completely disjoint from the original and modifications to the original will have no effect on the copy. If the original object had multiple aliases that could modify the object unexpectedly, then a local, unaliased copy can be created, which is immune to outside mutation.


```

copySelfContained( heap, src-location , class )
  // Begin the copy process with an empty copy-map
  hcopy-map, dst-location i = copy(heap, ?, src-location , class )
  return dst-location

// Recursively copy a tuple within an accumulated copy-map
// Tuples that have already been copied simply return the heap location
// of the copy.
copy( heap, copy-map src-location , class )
  if src-location 2 copy-map then
    return hcopy-map copy-map(src-location )i

  dst-location = heap.alloc( class )
  copy-map = copy-map[src-location 7! dst-location ]
  src-fields = heap(src-location )

  for( i = 0; i < len( src-fields ); i++ )
    if src-fields [i] is primitive then
      heap(dst-location )[i] = src-fields [i]
    else
      hcopy-map, dst-field i =
        copy(heap, copy-map, src-fields [i], classOf( src-fields [i]))
      heap(dst-location )[i] = dst-field

  return ( copy-map, dst-location )

```

Figure 2.10: Pseudo-code for Deep-Copying a Self-Contained Tuple

Copying objects in the container language is automatic and there is `copy` function or operator. Copying occurs automatically when assigning by value or passing a parameter by value. As an example, once you can treat an object like data, then a feature like an undo button in a word processor could simply be implemented by taking periodic snapshots of the document and simply reverting to an earlier snapshot when needed.

Figure 2.10 outlines the algorithm to deep-copy tuples. The key component is the `copy-map`, which records a mapping of old source tuples to their copied counterparts. The copy process is recursive, and when a tuple is found for the first time, it is copied and added to the map. If there is a second alias to the tuple, then a map lookup is used to find the matching copied tuple and avoid copying the same tuple multiple times. This ensures that cycles of references are handled correctly and the code does not get stuck in an infinite loop.

```

compareTuple( heap; compare-map lhs-location ; rhs-location ; class )
  if lhs-location == 2 compare-map then
    // If we are comparing a tuple we compared before, then the lhs must also be the
    // same tuple
    return hcompare-map compare-map lhs-location == rhs-location i
  else
    compare-map = compare-map lhs-location != lhs-location ]
    lhs-fields = heap( lhs-location )
    rhs-fields = heap( rhs-location )
    field-types = getFieldTypes( class )

    // Process each field
    for( i = 0; i < len( lhs-fields ); i++ )
      if field-types [i] is primitive then
        if lhs-fields [i] != rhs-fields [i] then
          return hcompare-map false i
      else if ( lhs-fields [i] == null and rhs-fields [i] != null ) or
        ( lhs-fields [i] != null and rhs-fields [i] == null ) then
        return hcompare-map false i
      else
        hcompare-map same = compareTuple( heap; compare-map lhs-fields [i],
          rhs-fields [i], field-types [i].class )
        if not same then
          return hcompare-map false i
    return hcompare-map true i

```

Figure 2.11: Pseudo-code for Deep Object Comparison

2.4.2 Equality Comparison

Another benefit of self-containment is a well-defined equality comparison. Comparison by reference equality is common practice in object oriented-languages. Here, we have the ability to compare complex objects by their content using the common `==` operator. The process works very much like the copy algorithm. When two tuples are compared, a map is kept to associate nested tuples with the corresponding tuples in the other object. Figure 2.11 outlines a comparison algorithm. The algorithm is recursive and it short-circuits when a tuple is encountered a second time, just like the copy algorithm.

This algorithm would work on a non-self-contained object as well, however, we disallow it based on the opinion that entities are better compared by reference. As with copying objects, custom comparison code can be written for entities. Custom code has the advantage of knowing the purpose of an entity's outside relationships and can make better decisions than a generic comparison.

2.4.3 Serialization

The last benefit of self-contained objects we'll discuss is serialization. The language supports the statements `load` and `save` that can be used to persist and reload any self-contained object. Serialization can be done using the similar recursive logic with a map tracking previously encountered objects as was used to copy objects. By using containers, we can guarantee that an object serialized to a file and then de-serialized back into a new object will compare as equal to the original using the `==` operator.

Serialization of data covers a wide range of use cases such as persistence and messaging between threads or processes across a network. By requiring self-containment, we don't require additional meta-data in order to marshal data correctly. Because of container enforcement, there is also nothing a programmer could do to make a self-contained class un-serializable.

2.5 Demonstration System

The demonstration language was implemented in Haskell and implemented as an interpreter using a style that loosely resembles small-step semantics. This design was chosen so that this code base could inform the subsequent work to formalize the language. The source code for the project can be found on gitlab at <https://git.uwaterloo.ca/mthode/woven-c>.

There are a large number of test cases in the `test/` folder and although they were written for testing, they would also be useful to a human looking to get a feel for how the language works. On a Linux system, the interpreter can be built using the `build` command in the root directory. The `ghc` Haskell compiler must be installed beforehand. In particular, the test `test/stmt29.woven` would be a good example to start with. It implements a linked list and implements a sort method on the list. Each test also includes a `.chk` file which verifies the output when the `-t` option is used.

```
Usage: woven [args] (<source-file> | <test-name>)
-t      Testing Mode.  Runs a test if specified, otherwise runs all tests
-i      Interactive execution
-vt     Verbose output of scanned tokens
-va     Verbose output of abstract syntax tree
-vta    Verbose output of annotated/typed abstract syntax tree
-ve     Verbose output of the type environment
-vx     Verbose output of each step of execution
-vg     Write graph file graph.dot after completion
```

The interpreter breaks execution down into steps, which can be observed using the command line option. When this option is used, the heap and stack are displayed after each step is executed. Also of interest is the -vg option, which will output a diagram of the heap into a file called graph.dot which can be interpreted by the Graphviz application. This was a useful tool to verify object copying semantics were working correctly.

```
dot -T pdf graph.dot > graph.pdf
```

Chapter 3

Containers Formalized

In this chapter, we will formalize a type system and language for containers. The language is reduced in scope for the system described in chapter 2 to reduce the complexity. Methods have been replaced with functions. At the beginning of development, it was assumed that methods would need special treatment in order to manage their containers. However, through the development of generic containers, it was found that the special containers `self` and `container-of-self` could be supported by the same generic container mechanism used for other parameters. Since no special treatment was necessary, the formal system was simplified to plain functions and constructors were replaced with a tuple initializer command.

Indirect references were also removed for two reasons. Firstly, in the demonstration system, the implementation of `iref` didn't add any significant challenges beyond what was needed for `plairref` variables. Secondly, this work in its present state hasn't properly motivated `iref`'s existence. In chapter 7 there is a brief discussion about adding a concept of a role to the language, and `andref` was meant to be a building block for that work.

Primitive types and operators have also been removed, as these can be modeled as special tuples and functions. With these changes in mind, significant complexity still remains and we are left with many details that are often omitted in formal works. Types are defined as a 5-tuple in order to include container information as well as carefully tracking properties like mutability which are critical to the type rules.

Although methods are replaced with functions in the formalism, tuples remain a key ingredient and the intent is still to consider these tuples as objects even though an abstraction mechanism and polymorphism are not present. We will refer to a tuple's definition as a class, remaining consistent with object-oriented terminology.

In addition to the planned changes for the abstract language, there were many incremental changes made as the formalism was developed. Even though the interpreter was implemented in a style similar to the formalism, modifications were needed in order to facilitate proving the soundness of the semantics. There is no claim that the soundness of the formalism implies soundness of the demonstration system as there are many differences. No doubt, some of the changes that were needed in the abstract language represent bugs in the demonstration system.

The type system defined in this chapter is a dependent type system which introduces significant complexity. The types defined in this chapter include a container label which refers to another in-scope construct (object or scope) to indicate the container of an object or the container constraint of a reference. One of the primary concerns of this formalism is stability and the importance of types depending solely on immutable symbols.

3.1 Types

This section details the structure of a container-type. Figure 3.1 defines the formal representation of a type including a list of the possible container labels. We'll discuss each of the components in turn.

Variables and fields have `kind`, which can be either `aref` or a value, and each variable also has a mobility and mutability component, which are represented with the symbols `mob` and `mut`, respectively. Mobility is similar to Java's `final` keyword. Here, a reference is considered movable if the identity of the object referred to can change. Mutability indicates whether the destination of a reference can be modified. We also need to distinguish between an object being mutable and a reference that is able to mutate the object it refers to; the former being a guarantee that a referenced object will never change, and the latter simply a permission granted to the reference. These distinctions are essential because the container language uses dependent types which require certainty that all dependent properties are preserved.

One option considered was representing mutability as two properties: one describing the reference (this-alias-has-mutation-permission) and the other describing the object (this-object-is-mutable). One issue is that they are not orthogonal, because you can't have permission to modify something that is immutable. This scheme was rejected in favor of a single property with three possible values: mutable, read-only, immutable. The read-only state disallows a reference from modifying its object, but it does not guarantee that the object can't be modified by another alias. Therefore the type rules will be

Identifiers

C Class name
F Function name
f Field name
v Variable name (parameter or local variable)
g Generic container label name

Container Labels

Path ::= tuple | var(v) | step(path; f) path 2 Path
ContainerLabel ::= unknown-label 2 ContainerLabel
 | null-label
 | local
 | container-of-tuple
 | generic(g)
 | path

Types

Kind ::= value | ref kind 2 Kind
Mutability ::= mutable | read-only | immutable 2 Mutability
Mobility ::= movable | unfixed | fixed 2 Mobility
Type ::= hkind; C; ; ; i T 2 Type

Type Deconstructors

With T = hkind; C; ; ; i
 $T_{(kind)} = kind \quad T_{(class)} = C \quad T_{(label)} = \quad T_{(mut)} = \quad T_{(mob)} =$

Figure 3.1: Representation of Container Types

pessimistic, and read-only will disqualify an expression from being used to indicate the container of a symbol.

Having a read-only state simplified the type rules because it allows uncertainty about the true mutability of the object. Precise knowledge of every object's mutability would be a nice property to have, but functions would lose generality if immutable parameters were distinct from read-only-mutable parameters.

References can't be declared as immutable following the principle above about the generality of passed parameters. It's not permitted for values to be declared as read-only; this wouldn't make sense because there is no ambiguity at the original declaration site. Ambiguity of the mutability of an object only arises after aliases are created and passed around.

The story is similar for mobility, where movable means that a symbol can be moved to reference a different object. Only when an expression types with = movable can that expression be used on the left-hand side of an assignment statement. The other use of mobility is to help determine when an expression can be converted into a container label. This is important for typing field access expressions, which will be discussed further in section 3.4.1. Only when = fixed can an expression be converted into a container label. Similar to read-only, unfixed is a middle ground where stability is not guaranteed and assignment is not possible.

The final component of a type is the container label, which indicates the container relative to symbols in the current scope or to the scope itself. This is the key feature of this system which enables the typing of containers.

```

ref r : Foo :: local => null;           // <ref, Foo, local, mutable, movable>
ref r : Foo :: x.y => null;           // <ref, Foo, step(var(x), y), mutable, movable>
readonly ref r : Foo :: local => null; // <ref, Foo, local, read-only, movable>
fixed ref r : Foo :: local => null;    // <ref, Foo, local, mutable, fixed>

val x : Bar = Bar();                 // <value, Bar, local, immutable, fixed>
var x : Bar = Bar();                 // <value, Bar, local, mutable, movable>
fixed var x : Bar = Bar();           // <value, Bar, local, mutable, fixed>

```

Figure 3.2: Formal Type Examples: Type declarations mapped to their formal representation.

When references are typed with container labels, they describe the container of the referred object. Value types behave differently, and as we'll see in the next section, the type rules require that value symbols must be typed with self-contained classes. This means that unlike references, the container for values is always the scope in which they are

declared. For example, a local variable of a self-contained class will always have container label `local` and a field value will always have container label `tuple`. In chapter 2, details like this were implicit, whereas in the formalism everything is made explicit.

The label `unknown-label` can be used whenever the code has no concern for which container a referenced object belongs in. As mentioned in chapter 2, a goal is to minimize the burden of writing code with containers, and `unknown-label` allows for the omission of a definite label. Fields are an exception: `unknown-label` is not permitted as the label for a field. Since objects exist within containers, a field cannot be `unknown-label`, as that would allow for container violations, because all fields in an object must obey the object's container boundaries.

A special label exists solely for typing the expression `null`, namely `null-label`. It's not permitted to declare a symbol with this label. The label `local` represents the scope of a function invocation, and all references constrained by label `local` can only reference local variables. The label `var(v)` indicates that the container is a parameter or variable within the current typing environment.

The labels `local` and `var(v)` are dependent on their scope, and this raises the question of how parameters can be passed across scopes and retain the correct container label. For example, how can a reference to a local variable in the caller scope be passed to the callee? This issue is solved by the use of generic container labels. Functions that use parameters typed with generic labels are generic functions parameterized by the set of generic labels appearing in the parameter types. When calling such a function, the generic substitutions are made based on the container labels of the passed parameters. Parameters, return values and local references with generic container labels are indicated by the label `generic(g)`. Note that function parameters can only be generic in their labels; no mechanism is provided for parameters to be generic in their class.

For reference fields, the labels `container-of-tuple` and `tuple` are used, with `container-of-tuple` meaning the container of the tuple that the field belongs to, and `tuple` meaning the tuple itself is the container.

The last container label to discuss is `path`. A path begins with a root container label that is also an object; this excludes `local` and `container-of-tuple`. Starting with the root, a series of field accesses can be taken to reach the final container. This allows you to declare a reference as contained within a series of nested fields.

In all, there are 5 components to a type, with each playing a role in the correct enforcement of containment. The type rules are defined in section 3.4, but first comes the definition of the abstract language.

ClassContainment ::= self-cont by-container	cc2 ClassContainment
ClassDe nition ::= $h\bar{f} : \bar{T}; c\bar{d}$	class-def2 ClassDe nition
FunctionSignature ::= $h\bar{p} : \bar{T}; T_{ret}i$	S 2 FunctionSignature
FunctionDe nition ::= $hS; si$	
LabelMap ::= $f \xrightarrow{\text{callee}} \xleftarrow{\text{caller}} g$	2 LabelMap
Expression ::= $var(v)$	e2 Expression
$field(e; f)$	
$call(F; ; e)$	
$init(C; ; e)$	
$null$	
Statement ::= $seq(s_1; s_2)$	s 2 Statement
$let(T; v; e; s)$	
$assign-value(T; e_{lhs}; e_{rhs})$	
$assign-ref(e_{lhs}; e_{rhs})$	
$return(e)$	

Figure 3.3: Abstract Language De nition

3.2 Language

The grammar of our abstract container language is in gure 3.3. Class de nitions consist of a list of elds and their types as well as a `ClassContainment` setting, which determines if instances of this class are self-contained, meaning no reference can point to an external object, or contained by the container that the object is placed into. Function signatures contain a list of parameter types and the return type.

A `LabelMap` holds a mapping of container labels mapping labels in a callee scope back to labels in the caller's scope. Chapter 2 described an algorithm which could calculate this mapping, however we omit this complication in this abstract language. The speci cation of the `LabelMap` is part of the program.

There are ve forms of expressions: variable access, eld access, function call, tuple initialization and the null literal. Both `call()` and `init()` take a parameter. For functions, it maps parameter labels, and for tuple initialization, provides a mapping for the container-of-tuple label.

For statements we have `seq()` that simply allows multiple statements to exist in a function body, which by de nition is a single statement. The statement `let` de nes a new local variable and initializes it. The nal parameter to `let` is a statement which runs with the new variable de ned. Once this statement completes, the symbol is out of scope and

$\text{::=} f \mathcal{C} \quad \text{ClassDe nition } j \mathcal{F}! \quad \text{FunctionDe nition } g$
 $\text{::=} hf \mathcal{V}! \quad \mathcal{T}_g; \mathcal{T}_{\text{ret}} i$

$\text{With } = hf::; v_i \mathcal{T}_i; :::g; \mathcal{T}_{\text{ret}} i$
 $(v_i) = \mathcal{T}_i$
 $(\text{ret}) = \mathcal{T}_{\text{ret}}$
 $[v_{\text{new}} \mathcal{T}_{\text{new}}] = hf::; v_i \mathcal{T}_i; :::; v_{\text{new}} \mathcal{T}_{\text{new}}g; \mathcal{T}_{\text{ret}} i$

Figure 3.4: Global and Local Typing Environments

not available to subsequent statements. There are two assignment statements to assign values and references. The distinction between these two statements isn't as important as it was in the demonstration language in chapter 2 because there is no indirect reference kind where you could assign a value through a reference. In this formalism, the left-hand side of assign-ref must always be a reference and the left-hand-side assign-value must always be a value. Finally, the return statement returns a value from a function as one would expect.

3.3 Typing Environments

For typing environments, we use the symbol for global de nitions of classes and functions. For typing of local environments for statements and expressions, is used. Both are de ned in gure 3.4 as well as notation for deconstructing components from and adding new symbols to .

3.4 Type Rules

In this section, we de ne the typing rules for the abstract container language. To determine if a program is valid, all classes in must satisfy ` class-ok (gure 3.13) and all functions must satisfy ` hS; si func-ok (gure 3.18). However, there are many components of classes and functions that need individual validation and we will begin by de ning a set of functions needed to support the type rules.

```

// Find the set of container labels use by the parameters of F
// that need to be in the domain of phi.
distinctParameterLabels( ; F) =
  let
    function(  $\overline{hp} : T; T_{ret} i; \_$ ) = ( F )
  in
    distinctParameterLabels2(  $T_{ret} \overline{T}$  )

distinctParameterLabels2(  $T_1 \overline{T}$  ) =
  let
    distinct = distinctParameterLabels2(  $\overline{T}$  )
    = requiredLabelMapping(  $T_{1(label)}$  )
  in
    if  $\notin$  unknown-label  $\wedge \neq$  distinct
    then distinct else distinct

requiredLabelMapping( )
  case of
    var( v) ! var( v)
    generic( g) ! generic( g)
    step(  $_{base}$ , f) ! requiredLabelMapping(  $_{base}$ )
    otherwise ! unknown-label

// Take a container label relative to a scope and compose a
// container label relative to the callers scope based on phi.
mapLabel( , ) =
  case of
    step(  $_{base}; f$ ) !
    step(mapLabel( ;  $_{base}$ ); f)
    otherwise !
    ( )

// Take a type relative to the current and compose a type relative to
// the calling scope based
exportType( , T) = hT $_{(kind)}$  ; T $_{(class)}$  ; mapLabel( ; T $_{(label)}$  ); T $_{(mut)}$  ; T $_{(mob)}$  i

// Construct a new type environment for the body of a function.
// The function's parameters populate the scope.
funcEnv( ; F)
  let
     $\overline{hp} : T_p; T_{ret} i = ( F )$ 
  in
     $\overline{hp} \overline{T}_p; T_{ret} i$ 

```

Figure 3.5: call () Related Typing Functions

```

expressionToLabel( ; ; e) =
; \ e: T
if T(mob) ∈ fixed then
// The expression e does not evaluate to a fixed container
unknown-label
else
case e of
var( v)! var( v)
field( eobj; f)!
case expressionToLabel( ; ; eobj) of
path! step( path; f )
otherwise ! unknown-label
otherwise ! unknown

```

Figure 3.6: Field Access Related Typing Functions (1)

3.4.1 Type Rule Helper Functions

The first set of functions in figure 3.5 help with typing function calls. In order to determine if the label mapping parameter `toCall` () is correct, we need to know what the domain of `toCall` should be. The `distinctParameterLabels` () function builds the list of labels that must be in the domain of `toCall`. This list includes all of the generic `var` container labels appearing in the function parameters.

The function `mapLabel`() takes care of applying the mapping. Its main behavior is to look up `toCall` (`callee`), but when paths are present, it recursively processes the paths until it reaches the root, then the root is replaced according to `toCall` and the path is reconstructed using the new root. The `exportType` () function is essentially a wrapper around `mapLabel`() to map the entire type tuple. All components other than the container label are directly mapped.

The function `funcEnv`() builds the initial typing environment for the function containing the types of each parameter and \bar{T}_{ret} for the return value.

The functions defined in figures 3.6 and 3.7 all relate to the typing of the field access expression. When accessing a field we immediately know the type of the field with respect to the tuple it is defined in, but type we need must be relative to the local scope and a translation is needed.

Fields are typed with container labels relative to either tuple or container-of-tuple and these must be converted to labels relative to the current scope. For example, using the field access notation of the demonstration language, in expression `a.b.c`, we begin with the tuple-relative type of `c`. If `c` is typed with a container label rooted in container-of-tuple `α`, then the container label of `a.b.c` is the same as the container label of `α.b`, which can be

```

// Take a field's declared container label and compose a new label
// relative to the local scope
mapFieldLabel( ; ; cont ; tuple ; field ) =
  case field of
    container-of-tuple ! cont
    tuple ! tuple
    step( base, f ) ! case mapFieldLabel( ; ; cont ; tuple ; base ) of
      unknown-label ! unknown-label
      base ! step( base, f )

// Determine the type of a field access expression
fieldAccessType( ; ; e_obj ; f )
let
  ; \ e_obj : T_obj
  T_field = ( T_(class) )(f)
  tuple = expressionToLabel( e_obj )
  = mapFieldLabel( ; ; T_obj(cont) ; tuple ; T_field (cont) )

  = if T_field (kind) == value then
    // When a value field is mutable, the expressions mutability
    // is inherited from the base object which might not be mutable
    if T_field (mut) == mutable then
      T_obj (mut)
    else
      immutable
  else
    // Field is a reference
    if T_obj (mut) == mutable ^ T_field (mut) == mutable then
      mutable
    else
      read-only

  = if T_field (kind) == ref ^ == unknown-label then
    // If we can't determine a precise container then no updates
    // to references are allowed
    unfixed
  else if T_field (mob) = movable then
    if T_obj (mut) = mutable then
      movable
    else
      unfixed // Object immutability overrides a movable field
  else
    // Field is fixed
    if T_obj (mob) == fixed then
      fixed
    else
      // When the base expression is unstable then everything
      // following is also unstable
      unfixed
in
hT_field (kind) ; T_field (class) ; ; ; i

```

Figure 3.7: Field Access Related Typing Functions (2)

determined by recursion.

Otherwise, if field `c` had the container label `tuple`, then the container is the object that the expression `a.b` evaluates to. The function `expressionToLabel()` shown in figure 3.6 is used to make the conversion from an expression to a container label. It builds a container label path from a root variable and a series of field-access steps. There are many expressions that can't be converted to a specific container label. In these cases `unknown-label` will be returned. Container labels need to be stable, so all dependencies present in them must be immutable. For an expression to be converted, it must only contain fixed variables and fields.

The final two functions in figure 3.7 determine the type of a field access expression. There is special logic for the container label, mutability and mobility of the type. The function `fieldAccessType()` makes use of `expressionToLabel()` to determine `tuple` which is the container label representing the object `obj` evaluates to. With `tuple` computed, `mapFieldLabel()` does the appropriate substitutions to create a label for the accessed field which is relative to the symbols in the local scope.

The remainder of `fieldAccessType()` computes the mutability and mobility components of the type. For `value` fields can have their mutability overridden by the object. For example, if a field is mutable but the object is immutable, then you should not be able to modify the field. References are treated similarly, but `read-only` is always used instead of `immutable`.

Figure 3.8 provides two examples of container label determination for field access expressions. First all the relevant information is gathered, then `mapFieldLabel()` selects the appropriate container label. In figure 3.9, the if-then-else logic of figure 3.7 is unraveled. The tables show the full set of possibilities for determination of `mutable` and `mobile` for a field access expression. The final column shows the result that `fieldAccessType()` will return, based on the values in the preceding columns.

For the mobility component, the determination of `mobile` also takes into account the previously computed `mutable`. The reason is that `unknown-label` is overloaded in meaning. When a reference parameter or local variable is declared with container label `unknown-label`, it indicates that it doesn't matter which container it references. This means that the type rule for reference assignments permits any right-hand-side container label to be assigned to `unknown-label`. However, in all other areas, `unknown-label` should be taken literally, as there is no knowledge available about what container the referred object is inside. To resolve this ambiguity, expressions typed with `unknown-label` that are not actually don't-care are also typed with `mutable = unfixed` to prevent assignments that would break the containment system. The expression `var(v)` is the only expression that can be typed with

Figure 3.8: Container Labels for Field Access Expressions

Field Kind	Field	Base Expr.	Expression
value	mutable		
value	immutable	any	immutable
ref	mutable	mutable	mutable
ref	read-only	any	read-only
ref	any	not mutable	read-only

Field Kind	Expr	Field	Base Expr.	Base Expr.	Expr.
ref	unknown	any	any	any	unfixed
any	any	movable	any	mutable	movable
any	any	movable	any	not mutable	unfixed
any	any	fixed	fixed	any	fixed
any	any	fixed	not fixed	any	unfixed

Figure 3.9: Mutability and Mobility for Field Access Expressions

both movable and unknown-label .

The remainder of the logic for determination is less subtle. An immutable object prevents assignments to a field which is movable, and fixed fields become unfixed if the object is not also fixed . This ensures that any attempt to convert this expression's type into a container label only succeeds when both the object and field are fixed .

3.4.2 Type Validation

$$\begin{array}{c}
 \text{T-SP-Step} \\
 \frac{;; \text{hC}; \overline{f : \overline{T_i}} \quad \text{path} : C_{\text{base}} \quad C_{\text{base}} \in C \quad T = (C_{\text{base}})(f) \quad T_{(\text{mobility})} = \text{fixed}}{;; \text{hC}; \overline{f : \overline{T_i}} \quad \text{step}(\text{path}; f) : T_{(\text{class})}}
 \end{array}$$

$$\begin{array}{c}
 \text{T-SP-Self-Step} \\
 \frac{;; \text{hC}; \overline{f : \overline{T_i}} \quad \text{path} : C_{\text{base}} \quad C_{\text{base}} = C \quad f = f_i \quad \overline{f : \overline{T}} \quad T_{(\text{mobility})} = \text{fixed}}{;; \text{hC}; \overline{f : \overline{T_i}} \quad \text{step}(\text{path}; f) : T_{(\text{class})}}
 \end{array}$$

$$\begin{array}{c}
 \text{T-SP-Tuple} \\
 \frac{C \ 2}{;; \text{hC}; \overline{f : \overline{T_i}} \quad \text{tuple} : C}
 \end{array}
 \qquad
 \begin{array}{c}
 \text{T-SP-Var} \\
 \frac{T = (v) \quad T_{(\text{mobility})} = \text{fixed}}{;; \text{h?}; ? \ i \quad \text{var}(v) : T_{(\text{class})}}
 \end{array}$$

$$\begin{array}{c}
 \text{T-SP-Parm} \\
 \frac{\overline{v = v_i \ 2 \ v : \overline{T}} \quad T_{(\text{mobility})} = \text{fixed} \quad T_{(\text{kind})} = \text{ref}}{; \ ? \ ; \ \text{h?}; \ \overline{v : \overline{T_i}} \quad \text{var}(v) : T_{(\text{class})}}
 \end{array}$$

Figure 3.10: Path Validation

With our helper functions defined, we can begin defining the type rules, beginning with container label path validation. In figure 3.10, the judgment $;; \text{hC}; \overline{f : \overline{T_i}} \quad \text{step}(\text{path}; f) : T_{(\text{class})}$ determines if a container label path is valid. Paths are distinguished from other container labels in that they refer to containers that must be objects and type as a class instance. This isn't true for container labels `local`, `container-of-tuple` and `generic (g)`.

The class name $\in C$ indicates that the path exists in a class declaration context in the container label of a field. A field's type must only reference preceding fields so that de-

dependencies can be initialized beforehand, and the specification of enables this logic. Otherwise, \emptyset is used in place of \mathcal{C} when not in a class declaration context.

The $\overline{v : T}$ part of the judgment is also for managing the order of dependencies. Both parameter declarations and field declarations populate a list of previous declarations. When local variables are typed, \emptyset is used and variables are found in $\overline{v : T}$ by the rule T-SP-Var . When field and parameter declarations are typed, $\overline{v : T}$ will be empty.

The critical feature that all four path typing rules share is $\text{isT}_{(\text{mobility})} = \text{fixed}$. This ensures that all steps along a path are stable and that a container label will continue to specify the same container throughout its lifetime. Parameters have an additional restriction that any dependency must be a reference parameter. Since value parameters are passed by value, a copy is made when the function is invoked. Therefore, no other passed parameter could reference this fresh object.

$$\begin{array}{c}
\text{V-Field-Label} \\
\frac{\text{kind} = \text{ref}}{\text{; } \overline{\text{hC}; \text{f} : \overline{\text{T}}\text{i}}; \text{kind} \text{ ` container-of-tuple label-ok-f}} \\
\\
\text{V-Field-Label-Path} \\
\frac{\text{; } ?; \overline{\text{hC}; \text{f} : \overline{\text{T}}\text{i}} \text{ ` step (path; f) : C}}{\text{; } \overline{\text{hC}; \text{f} : \overline{\text{T}}\text{i}}; \text{kind} \text{ ` step (path; f) label-ok-f}} \\
\\
\text{V-Parameter-Label-Local} \qquad \text{V-Parameter-Label} \\
\frac{\text{= local} \quad \text{kind} = \text{value}}{\text{; } \overline{\text{v} : \overline{\text{T}}}; \text{kind} \text{ ` label-ok-p}} \qquad \frac{2 \text{ f unknown-label, generic(_) } \text{ g} \quad \text{kind} = \text{ref}}{\text{; } \overline{\text{v} : \overline{\text{T}}}; \text{kind} \text{ ` label-ok-p}} \\
\\
\text{V-Parameter-Label-Path} \\
\frac{\text{; } ; \text{ h? ; } \overline{\text{v} : \overline{\text{T}}\text{i}} \text{ ` step (path; f) : C} \quad \text{kind} = \text{ref}}{\text{; } \overline{\text{v} : \overline{\text{T}}}; \text{kind} \text{ ` step (path; f) label-ok-p}} \\
\\
\text{V-Local-Label-Local} \qquad \text{V-Local-Label} \\
\frac{\text{= local}}{\text{; } ; \text{ kind} \text{ ` label-ok-l}} \qquad \frac{2 \text{ f unknown-label, generic(_) } \text{ g} \quad \text{kind} = \text{ref}}{\text{; } ; \text{ kind} \text{ ` label-ok-l}} \\
\\
\text{V-Local-Label-Path} \\
\frac{\text{; } ; \text{ h? ; } ? \text{ i} \text{ ` step (path; f) : C} \quad \text{kind} = \text{ref}}{\text{; } ; \text{ kind} \text{ ` step (path; f) label-ok-l}}
\end{array}$$

Figure 3.11: Container Label Validation

The remainder of the container label validation rules are in figure 3.11. Here, there are different judgments for fields, parameters and local variables. Different rules are applied for these contexts. For example, a field is not permitted to be labeled `unknown-label`. Only `container-of-tuple` and a path are valid for fields. This is because objects are placed in containers and a reference field can never be more permissive than its object.

For typing parameter labels, there are restrictions where parameters can only depend on parameters declared before them. The label `unknown-label` is permitted, so function parameters can be declared without needing to specify a generic container label when there is no need to know the container. Reference parameters are not allowed to be contained within value parameters. This is because the pass-by-value semantics make a fresh copy of

value parameters, so no passed reference parameter could refer to the fresh object.

The typing of values is consistent in that their container is the scope they are declared in, meaning `tuple` for elds and `local` for parameters and local variables. Other labels are only used for references.

$$\begin{array}{c}
 \text{V-Field-Type} \\
 \frac{\text{` } T \text{ type-ok-common} \quad ; T_{(\text{kind})} \text{ ` } T_{(\text{label})} \text{ label-ok-f}}{\quad ; hC; \bar{f} : \bar{T}i \text{ ` } T \text{ type-ok-f}} \\
 \\
 \text{V-Parameter-Type} \\
 \frac{\text{` } T \text{ type-ok-common} \quad ; ; T_{(\text{kind})} \text{ ` } T_{(\text{label})} \text{ label-ok-p}}{\quad ; v : \bar{T} \text{ ` } T \text{ type-ok-p}} \\
 \\
 \text{V-Local-Type} \\
 \frac{\text{` } T \text{ type-ok-common} \quad ; ; T_{(\text{kind})} \text{ ` } T_{(\text{label})} \text{ label-ok-l}}{\quad ; \text{ ` } T \text{ type-ok-l}} \\
 \\
 \text{V-Common-Type} \\
 \frac{\text{` } T_{(\text{class})} \text{ class-ok} \quad T_{(\text{kind})} = \text{ref } _ (T_{(\text{class})})_{(\text{containment})} = \text{self-cont}}{\text{ ` } T \text{ type-ok-common}}
 \end{array}$$

Figure 3.12: Type Validation

The verification of types is also broken into separate judgments for the three contexts that symbols can appear in. The rules defined in figure 3.12 are straightforward as the complexities are all handled in the `label-ok` judgments. The rule `V-Common-Type` checks that the class is well defined and that values can only be declared for self-contained classes. The reason for the self-contained restriction is that value types have pass-by-value and assign-by-value semantics and non-self-contained objects can't be copied.

$$\begin{array}{c}
\text{V-Fields} \\
\frac{f_1 : T_1 = \text{head}(\overline{f : T}) \quad ; \text{hC}; \overline{f_{\text{dep}} : T_{\text{dep}} i} \quad T_1 \text{ type-ok-f} \\
\quad ; \text{hC}; (f_1 : T_1) \quad \overline{f_{\text{dep}} : T_{\text{dep}} i} \quad \text{tail}(\overline{f : T}) \text{ fields-ok}}{; \text{hC}; \overline{f_{\text{dep}} : T_{\text{dep}} i} \quad \overline{f : T} \text{ fields-ok}} \\
\\
\text{V-Class} \\
\frac{; \text{hC}; ? i \quad \overline{f : T} \text{ fields-ok}}{\quad \quad \quad \backslash \text{ C class-ok}}
\end{array}$$

Figure 3.13: Class Validation

To validate a class, the judgment $\backslash \text{ C class-ok}$ is defined in figure 3.13. Each field must satisfy type-ok-f where the context consists of the fields declared before it.

3.4.3 Language Typing

$$\begin{array}{c}
\text{V-Initializer-Label-Map} \\
\frac{; \text{ caller } \quad \backslash \quad (\text{container-of-tuple} \quad) \text{ label-ok-l}}{; \text{ caller } \quad \backslash \quad \text{htuple} ; i \text{ lmap-ok}} \\
\\
\text{V-Function-Label-Map} \\
\frac{= \text{funcEnv}(; F) \quad \delta \quad 2 \text{ distinctParameterLabels} \quad (; F): \quad 2 \text{ dom}(\quad) \\
\delta \quad 2 \text{ dom}(\quad): ; \quad \backslash \quad \text{label-ok-l} \quad \delta \quad 2 \text{ range}(\quad): ; \quad \text{ caller } \quad \backslash \quad \text{label-ok-l}}{; \text{ caller } \quad \backslash \quad \text{hF}; i \text{ lmap-ok}}
\end{array}$$

Figure 3.14: Label Map Validation

Label maps play an important role in managing containers. Before a function call can be type checked, there needs to be a way to compare labels in the caller's scope with labels in the callee. The mapping takes labels as they were declared in the function and maps them back to the caller's symbol space. For example, if the function contains a reference parameter with labelgeneric (g), then the mapping must contain that label in its domain. The function distinctParameterLabels () establishes the set of labels that must appear in the domain of .

The lmap-ok judgment verifies that the domain of is valid with respect to function F, and then validates each label in the domain in the scope of the function and each label in the range of in the caller's scope.

$$\begin{array}{c}
 \text{Value-Initializable} \\
 \frac{T_{\text{lhs}}(\text{kind}) = \text{value} \quad T_{\text{rhs}}(\text{class}) = T_{\text{lhs}}(\text{class})}{; \quad \backslash \quad T_{\text{lhs}} \text{ e } T_{\text{rhs}}}
 \\
 \\
 \text{Ref-Initializable} \\
 \frac{T_{\text{lhs}} = \text{href}; C; \text{lhs}; \text{lhs}; _i \quad T_{\text{rhs}} = \text{h_}; C; \text{rhs}; \text{rhs}; _i}{\text{lhs} = \text{rhs} _ \text{lhs} = \text{read-only} \quad ; \quad \backslash \quad \text{lhs} \text{ l-match } \text{rhs}} \\
 ; \quad \backslash \quad T_{\text{lhs}} \text{ e } T_{\text{rhs}}
 \\
 \\
 \text{Value-Assignable} \\
 \frac{T_{\text{lhs}} = \text{hvalue}; C; _ _ ; \text{movable}i \quad T_{\text{rhs}}(\text{class}) = C}{; \quad \backslash \quad T_{\text{lhs}} \quad T_{\text{rhs}}}
 \\
 \\
 \text{Ref-Assignable} \\
 \frac{T_{\text{lhs}} = \text{href}; C; \text{lhs}; \text{lhs}; \text{movable}i \quad T_{\text{rhs}} = \text{h_}; C; \text{rhs}; \text{rhs}; _i}{\text{lhs} = \text{rhs} _ \text{lhs} = \text{read-only} \quad ; \quad \backslash \quad \text{lhs} \text{ l-match } \text{rhs}} \\
 ; \quad \backslash \quad T_{\text{lhs}} \text{ b } T_{\text{rhs}}
 \\
 \\
 \text{Container-Label-Match} \\
 \frac{\text{lhs} = \text{rhs} _ \text{lhs} = \text{unknown-label} _ \text{rhs} = \text{null-label}}{; \quad \backslash \quad \text{lhs} \text{ l-match } \text{rhs}}
 \end{array}$$

Figure 3.15: Assignment and Initialization

Next, in figure 3.15, we define a set of rules to determine which types are assignable to which destination types. The operator e is used for initialization where the mobility of the type is ignored. For regular assignments, is used for value assignment, and b is used for reference assignment. Value assignment has fewer conditions than reference assignment, because a copy of the right-hand side is made and the mutability of the copy is independent of the original. The left-hand side type must be value and the classes must match in order to be value-assignable. No subtyping is implemented in this system.

For reference assignment, the mutability of the left-hand side must match the right-hand side, unless the left-hand side is typed read-only, which is universally compatible.

The container labels must also match using the `match` judgment.

The normal case for labels to be compatible by `match` is if the left-hand side has the same label as the right-hand side. There are two special cases. First, if the left-hand side has `unknown-label`, then any right-hand side label is permitted. Second, if the right-hand side is `null-label` (can only occur with the `null` expression) then the assignment is permitted.

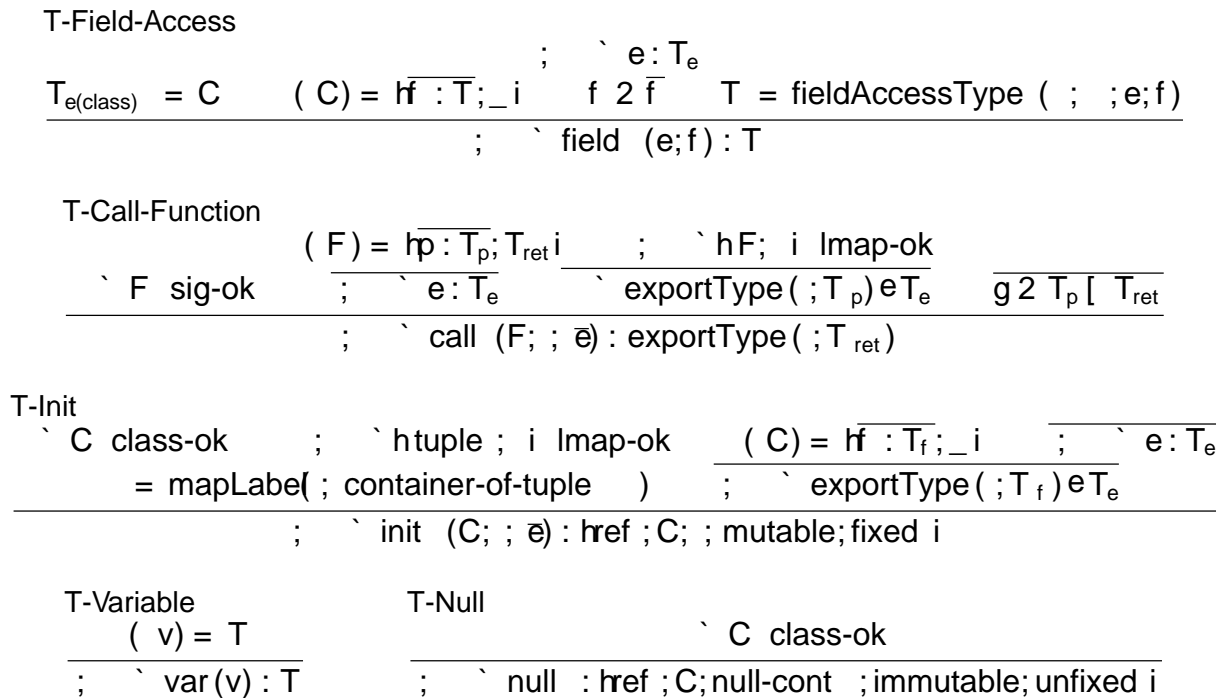


Figure 3.16: Expression Typing

With the building block rules defined, we can now type the expressions and statements of the language. Figure 3.16 contains the type rules for expressions. For typing field accesses, most of the work is done in the function `fieldAccessType()` which was described in section 3.4.1. Otherwise, the base expression must type correctly and `field` must be defined in the class of `e`.

In the rule `T-Call-Function`, we see how label maps are used in practice. The `exportType()` function (figure 3.5) applies the `map` to the parameter types and return type of the call `()` to map them into the container label space of the caller. Once they are

mapped, the type of each passed parameter is checked using the relation. Finally, the return type is also mapped and becomes the type of the `call ()` expression.

The rule `T-Init` behaves in a very similar way to a function call. The semantics of `init ()` is simply to allocate a new tuple in the heap and initialize each of the fields using the passed parameters. Even though there is no function code to run in a new scope, the new tuple is a scope with container labels relative to the tuple. Each parameter's type is checked against the respective field's type after it has been mapped into the space of the caller using `.`. Here, `.` is much simpler than `call ()` and only `container-of-tuple` is in its domain.

The rules for `var (v)` and `null` are straightforward. The `C` in `T-Null` is a free variable.

$$\begin{array}{c}
 \text{T-Let} \\
 \frac{\text{e} : T_{\text{init}} \quad T \text{ ok-I} \quad [\text{vars}]! \quad (\text{vars}) [h v ; T_i] \quad \text{s} : \text{Void}}{\text{let } (T ; v ; e ; s) : \text{Void}} \\
 \\
 \text{T-Statement-Sequence} \\
 \frac{\text{s}_1 : \text{Void} \quad \text{s}_2 : T}{\text{seq}(\text{s}_1 ; \text{s}_2) : T} \\
 \\
 \text{T-Assign-Value} \\
 \frac{\text{e}_{\text{rhs}} : T_{\text{rhs}} \quad T_{\text{lhs}} = \text{hvalue} ; _ ; _ ; _ ; \text{movable} \quad \text{e}_{\text{lhs}} : T_{\text{lhs}} \quad T_{\text{lhs}} \quad T_{\text{rhs}}}{\text{assign-value } (T_{\text{lhs}} ; \text{e}_{\text{lhs}} ; \text{e}_{\text{rhs}}) : \text{Void}} \\
 \\
 \text{T-Assign-Reference} \\
 \frac{\text{e}_{\text{rhs}} : T_{\text{rhs}} \quad T_{\text{lhs}} = \text{href} ; _ ; _ ; _ ; \text{movable} \quad \text{e}_{\text{lhs}} : T_{\text{lhs}} \quad T_{\text{lhs}} \quad T_{\text{rhs}}}{\text{assign-ref } (\text{e}_{\text{lhs}} ; \text{e}_{\text{rhs}}) : \text{Void}} \\
 \\
 \text{T-Return} \\
 \frac{\text{(ret)} = T_{\text{ret}} \quad \text{e} : T_{\text{e}} \quad T_{\text{ret}} \text{ e } T_{\text{e}}}{\text{return } (\text{e}) : T_{\text{e}}}
 \end{array}$$

Figure 3.17: Statement Typing

The rules for statements are defined in figure 3.17. The first rule `T-Let` types statement

s with variable $v : T$ added to its environment. For the assignment statements, the left-hand-side expression must be movable and the details of what types can be assigned are specified in figure 3.15. The last statement is `return` where the return expression $e : T$ must be initialization-assignable to the return type.

V-Parameters

$$\frac{\overline{p_{\text{dep}} : T_{\text{dep}}} \quad \overline{T_1} \text{ type-ok-p} \quad \overline{v_1 : T_1 = \text{head}(\overline{p : T})} \quad \overline{(p_1 : T_1) \quad \overline{p_{\text{dep}} : T_{\text{dep}}} \quad \overline{\text{tail}(\overline{p : T})} \text{ parms-ok}}}{\overline{p_{\text{dep}} : T_{\text{dep}}} \quad \overline{p : T} \text{ parms-ok}}$$

V-Signature

$$\frac{\overline{? \quad \overline{p : T} \text{ parms-ok}} \quad \overline{S = \overline{hp : T_p; T_{\text{ret}}}} \quad \overline{T_{\text{ret}} \text{ type-ok-p}} \quad \overline{T_{\text{ret}(\text{mob})} = \text{unfixed}}}{\overline{S \text{ sig-ok}}}$$

V-Function

$$\frac{\overline{S = \overline{hp : T_p; T_{\text{ret}}}} \quad \overline{S \text{ sig-ok}} \quad \overline{= \text{funcEnv}(\overline{; F})} \quad \overline{; \quad \overline{s : T_{\text{ret}}}}}{\overline{hS; si \text{ func-ok}}}$$

Figure 3.18: Function Validation

The final set of rules complete the typing of the container language and are defined in figure 3.18. These rules type function definitions, which are a pair of a function signature and a function body statement: $hS; si$. All parameters and the return value must satisfy `type-ok-p`. The parameter labels can only be dependent on parameters that come before and the return value must have the additional constraint that its mobility is `unfixed`.

Chapter 4

Operational Semantics

In this chapter, we define a small-step operational semantics for the language defined in chapter 3. We follow the methodology of structural operational semantics developed by Gordon Plotkin [17]. More directly, the style of our semantics were inspired by the MJ system by G.M. Bierman et al. [4].

4.1 Overview

The semantics of the language operate within a configuration which contains the program heap, stack and all other constructs needed to evaluate a program. This is illustrated in figure 4.1. The semantic rules relate a configuration to its subsequent configuration $C \rightarrow C'$.

The heap consists of references, boxes and tuples stored at abstract locations and the program is expressed by frames. A frame encodes an operation to be evaluated by the abstract machine. Frames include all of the expressions and statements of the language, but in addition, there are extra frames to support complex operations, which need to be decomposed into multiple scalar steps.

The configuration holds a current frame, which represents the next operation to be evaluated. In addition, the configuration contains a stack of stacks. The outer stack holds local configurations, which represent the state for a function invocation. Within a local configuration, there is a stack of frames, which we call the frame-stack. During the evaluation of a function, frames will be pushed and popped from the frame-stack as needed.

Figure 4.1: Configuration Overview

There are many definitions needed before we can present the semantic rules. In the next section, we will formally define the structure of the heap.

4.2 Heap

The heap structure shown in figure 4.2 is a mapping from an abstract location to a heap construct. There are three different constructs in the heap: tuples, boxes and references. We distinguish heap locations by which construct they refer to. A heap location is either a `TupleLocation` or a `SymbolLocation`, with `SymbolLocation` further subdivided into a `BoxLocation` or a `RefLocation`. The mappings for tuple locations are immutable once a tuple's fields are fully initialized. Mappings for symbol locations are mutable.

A tuple construct itself is an immutable mapping from field names to `SymbolLocations`. Later, we will see that local variables are also an immutable mapping from symbols to `SymbolLocations`. All mutation of the fields of a tuple occurs by modifying the individual `Box` and `Ref` constructs.

The distinction between `Box` and `Ref` is that a `Box` maps to a value tuple (self-contained) which is assigned by value, meaning the source tuple must be copied and the

```

Location ::= TupleLocation | SymbolLocation           ` 2 Location
SymbolLocation ::= BoxLocation | RefLocation
Tuple ::= FieldName SymbolLocation
Box ::= TupleLocation
Ref ::= TupleLocation | null-ref
Heap ::= f TupleLocationh Tuple                       H 2 Heap
        | BoxLocationh Box
        | RefLocationh Ref g

```

Figure 4.2: Heap Structure. TupleLocation, BoxLocation and RefLocation are abstract addresses in the heap. The tuple construct is immutable once fully initialized, but boxes and references can be updated.

box's mapping is set to a fresh tuple. The Ref construct, on the other hand, may reference objects that are not self-contained and is also permitted to be null-ref. No copy is made when a reference is assigned.

4.3 Frames

Our language consists of statements and expressions, but in order to implement a small-step semantics for this language, we need to introduce a number of additional frames to the language. Frames are the unit of computation, and each step of the machine applies a rule to the current frame to produce the next state of the configuration.

In addition to the statement and expression frames, figure 4.3 introduces a number of new frames. Expressions evaluate to a result frame, which is a pair of a heap location and the type of that heap location. The framevoid is simply the result of executing a statement. The statementlet introduces a symbol that is only in scope for the body statement and pop-local (v) is responsible for removing the variable when it goes out of scope.

The remainder of the frames relate to initialization of symbols and copying tuples. The frame init-symbols is used to marshal parameters for a function call, with init-symbol handling individual parameters. Similarly, init-fields and init-field are responsible for initializing a new tuple's fields.

There are a number of rules associated with copying an object. Although this process is similar to initializing a new tuple, the addition of a copy environment is required to track

```

Result ::= hLocation; Typei
CopyResult ::= CopyEnv | hResult; CopyEnv
AnyResult ::= Result | CopyResult
CopyEnv ::= f Locationl Locationg
Frame ::= ClosedFrame | OpenFrame

ClosedFrame ::= Statement
    | Expression
    | AnyResult
    | void
    | pop-local (VarName)
    | init-symbols ( hName; Resulti )
    | init-symbol ( Name; Result )
    | init-fields ( Result; hName; Resulti )
    | init-field ( Result; Name; Result )
    | assign-copied ( Result; Result )
    | copy-tuple ( Type, Result )
    | copy-tuple2 ( Type, Result, CopyEnv )
    | copy-fields ( Result; hName; Resulti, CopyEnv )
    | copy-init ( Result, Name, hResult; CopyEnv )
    | discard-copy-env ( hResult; CopyEnv )

```

```

R 2 Result          R(loc) = ` where R = h; T i
R 2 AnyResult
A 2 CopyEnv
F 2 Frame
CF 2 ClosedFrame

```

Figure 4.3: Closed Frames

```

OpenFrame ::= let (T, l, s)
  | assign-value (T, , e)
  | assign-value (T, l, )
  | assign-ref ( , e)
  | assign-ref (l, )
  | return ( )
  | field ( , f)
  | call (F, , R1 :: Ri-1; ; ei+1 :: en)
  | init (C, , R1 :: Ri-1; ; ei+1 :: en)
  | init-symbol (v; )
  | init-field (Rtuple; f; )
  | copy-fields (Rtuple; f; Ri, )
  | copy-init (Rtuple; f; )
  | discard-copy-env ( )

```

OF 2 OpenFrame

Figure 4.4: Open Frames

which tuples have already been copied. CopyEnv maps tuples contained within the source tuple to new tuples constructed within the destination tuple. When recursively processing the source tuple, a nested tuple could be aliased multiple times and cycles could also exist. The mapping ensures that each tuple in the source is only copied once, and when an alias is encountered, the destination reference is set to src-tuple .

The frame `copy-tuple` is the entry point into the copy environment. An empty CopyEnv is created and a `discard-copy-env` frame is pushed to the frame-stack to strip the CopyEnv from the resulting copied tuple. The frame `copy-tuple2` is then set as the current frame. Any recursive copies bypass `copy-tuple` and directly use the frame `copy-tuple2` so they can continue to use the same environment. As copies are made, the environment accumulates the tuple mappings. The frames `copy-fields` and `copy-init` behave much like `init-fields` and `init-field` except that they contain the copy environment as an extra parameter. The last frame associated with copying tuples is `assign-copied`, which completes an `assign-value` operation after the copy has been done.

Complex statements and expression must be evaluated in multiple steps. To facilitate this, there are decomposition rules which will extract the next sub-expression from a complex frame and replace it with a hole which is written as `.`. A frame is closed if it has no

Variables ::= $\overline{\text{Name}} \quad \overline{\text{SymbolLocation}}$	V 2 Variables
FrameStack ::= $\overline{\text{Frame}}$	F 2 FrameStack
ScopeID ::= Z^+	ID 2 ScopeID
Context ::= initial tuple FunctionName	ctx 2 Context
LocalConfig ::= $\text{hV}; \overline{\text{F}}; \text{hctx}; i; \text{ID}i$	L 2 LocalConfig

With $L = \text{hV}; \overline{\text{F}}; \text{hctx}; i; \text{ID}i$

$L_{(\text{vars})} = V$	$L[\text{vars } \overline{v}! V^0] = \text{hV}^0; \overline{\text{F}}; \text{hctx}; i; \text{ID}i$
$L_{(\text{fstack})} = \overline{\text{F}}$	$L[\text{decl } \overline{v}! \overline{v}] = \text{h}(\overline{v}! \overline{v}) \quad V; \overline{\text{F}}; \text{hctx}; i; \text{ID}i$
$L_{(\text{ctx, lmap})} = \text{hctx}; i$	$L[\text{fstack } \overline{f}! \overline{\text{F}}^0] = \text{hV}; \overline{\text{F}}^0; \text{hctx}; i; \text{ID}i$
$L_{(\text{lmap})} =$	$L[\text{push } \overline{\text{F}}] = \text{hV}; \overline{\text{F}} \quad \overline{\text{F}}; \text{hctx}; i; \text{ID}i$
$L_{(\text{id})} = \text{ID}$	$L[\text{push } \overline{\text{F}}_{\text{next}}] = \text{hV}; \overline{\text{F}}_{\text{next}} \quad \overline{\text{F}}; \text{hctx}; i; \text{ID}i$
$\text{Config} ::= @ \overline{\text{LocalConfig}}^A$	C 2 Config
ClosedFrame	$C = @ \overline{\text{L}} \quad \overline{\text{A}}$
	CF

Figure 4.5: Configuration

holes and open if there is a hole. After a sub-expression has completed, the result will be used to fill the hole and execution can continue. Figure 4.4 lists each of the open frames. Note that some frames will contain lists of expressions that will progress through a series of nested evaluations and hole-filling before the complete set of results can be operated upon.

4.4 Configuration

The state of the abstract machine as well as its component parts are defined in figure 4.5. The heap, stack and current frame are the three components of the configuration, which is written as a vertical vector. We've previously defined the heap and closed frames, but the stack is new and we'll introduce it now. The configuration contains a stack of local configurations, with each local configuration representing a scope. When a function call is made, a new local configuration is created and pushed onto the stack. When tuples are initialized, a new local configuration is also created.

Within a local configuration, variables and a frame-stack are stored. Variables map local symbols and function parameter symbols to locations in the heap. The frame-stack contains

frames that represent the remainder of the code that needs to run within a function. There are two common patterns of usage with the frame-stack. The first is when the current frame evaluates to a result with an open frame at the top of the stack. The open frame will have a hole () in it and the result will fill in the hole, making a closed frame which becomes the new current frame. The other main pattern is when the current frame is void, indicating that a statement has completed. In this case, a frame is popped off the frame-stack to become the new current frame.

With the stack of local configurations and a nested frame-stack within each local configuration, we have a stack of stacks. This decision was originally made to support an early return in the middle of a function, as is commonly supported in mainstream imperative languages. When a function returned early, the remaining inner frame-stack in the local configuration could be immediately discarded and control would pop directly back to the caller using the outer stack. However, other features required for a short-circuit return did not make it into this formalism so we are left with the complexity without the payoff. Despite this limitation, the stack of stacks complication was manageable and could be used in future research. Perhaps exceptions and break statements could be supported as well using this technique.

There are two more bookkeeping components of a local configuration. First is a LabelMap linking the container labels of the current scope with the caller's scope. Along with the mapping is the context for the mapping so that the stack typing rules can validate the map. The context in the normal case is the name of the current function. Otherwise, it is the context marked as tuple, which indicates a tuple initialization context, or initial, which is a special context which calls main() to bootstrap the machine.

The last component of the local configuration is a unique identifier for the scope. This identifier is used for configuration typing which will be explained in the next section.

Also in figure 4.5 are notations for extracting sub-components of a local configuration and also notations for updating L. These are used heavily in the rules for the operational semantics.

4.4.1 Configuration Typing Environment

In order to show soundness for the operational semantics, we need to type the configuration including all of our new frames. But before we address the typing, we need to introduce the configuration typing environment and basic definitions of physical types in contrast with the logical types that we have discussed so far. The new definitions are in figure 4.6.


```

Container ::= TupleLocation | ScopeID | unknown-cont | null-cont
PhysicalType ::= hClass Container; Mutability; Mobilityi
ContainerMap ::= f ContainerLabel Container g
Uninitialized ::= f TupleLocation g
Con gTyping ::= hf Locationl PhysicalType; f ScopeID? ContainerMapgi

```

```

    2 Container          T 2 PhysicalType
M 2 ContainerMap      2 Con gTyping

```

```

With = hphysical-typescontainer-map$
( ` ) = physical-type$
(cmap)(ID) = container-map$(ID)

```

Figure 4.6: Configuration Typing Environment

Container labels are relative to their scope, and the same container can be represented with different container labels in different scopes. When a reference parameter is passed to or returned from a function, we need to know that the container indicated in both scopes is indeed the same physical construct in the configuration.

We used lowercase as a container label and the physical configuration construct will be represented with uppercase. A physical container can be a tuple location in the heap or the unique identifier of a local configuration. Where container labels can only be interpreted relative to their scope, physical containers are global and can be compared independently of their origin.

Along with a physical container, we need a physical type \bar{T} , which is similar to our logical types T except the container label is replaced by a physical container. With these physical constructs defined, any typing we do will need to be able to convert logical types to physical types. We'll soon define the `thetype()` function in [figure 4.8](#), which performs these conversions, but for now we introduce the mapping $\bar{\cdot}$. We saw in [chapter 3](#) the mapping which maps labels to labels in the calling scope. Here, $\bar{\cdot}$ maps container labels to physical labels for a specific scope. Notation for accessing $\bar{\cdot}$ is also presented in [figure 4.6](#). An example is provided in [figure 4.7](#), with two scopes shown, a caller and callee. Concrete mappings are shown to demonstrate how the container labels and physical containers are connected.

With these definitions in place, $\bar{\cdot}$ is the typing environment for the configuration. Every location in the heap is mapped to a physical type. The environment also contains a container-map for every local configuration, which allows container labels to be mapped

Figure 4.7: Container Mappings Example

```

// Build a PhysicalType for a type with respect to a configuration
ptype( , H, L, T) =
  hT(class) , labelToCont( , H, L, T(label) ), T(mut) , T(mob)

// Determine the physical container for a container label within a
// local environment
labelToCont( , H, L, ) =
  case of
    unknown-label -> unknown-cont
    null-label -> null-cont
    variable( v ) -> locToCont( H, L(vars) ( v) )
    step( base, f ) ->
      let `tuple = labelToCont( , H, L, base) in
        locToCont( H, H(`tuple) ( f ) )
      -> (cmap)(L(id) )( )

// De-reference references and boxes
locToCont( H, l ) =
  case l of
    TupleLocation -> l
    BoxLocation -> H(l )
    RefLocation -> locToCont( H, H(l))

// De-reference references and boxes
toTuple( H, l ) =
  case l of
    TupleLocation -> l
    BoxLocation -> H(l )
    RefLocation -> H(l )

```

Figure 4.8: Operational Rules Support Functions (1)

to their physical container. Finally, we also introduce U which is a set of uninitialized tuples. Combined, Σ ; U comprise the entire typing environment for a configuration. The operational semantics do not depend on U in any way; they exist solely to enable the type safety proofs.

4.5 Supporting Functions

Similar to chapter 3, there are a number of supporting functions needed to type the configuration. Two important functions are reused from chapter 3, namely `exportType()` and `mapLabel()`, defined in figure 3.5.

```

// Create an initial local configuration for tuple initialization
tupleLocalConfig( h, ID ) = hfg; ? ; htuple ; i ; ID i

// Create an initial local configuration for function F
// with the function body s placed on the frame-stack
funcLocalConfig( F; ; s ) = hfg; s ; hF; i ; uniqueID() i

// Create a mapping for container labels to physical containers for
// a function body scope
labelMappings( h, H, L_caller, ) =
  f ; ! labelToCont ( ; H; L_caller ; ( i )) j ; i 2 dom( ) g [ f local ; ! L(id) g

// Create a mapping for container labels to physical containers for a tuple scope
tupleMappings( h, H, L_caller, , tuple ) =
  let cont = (container-of-tuple) in
    f container-of-tuple ; ! labelToCont ( ; H; L_caller ; cont ); tuple ; ! tuple g

// Take a type relative to the caller's scope current and compose a type
// relative to the current scope.
importType( T_caller ; T_local ) = hT_caller (kind) ; T_caller (class) ; T_local (label) ; T_caller (mut) ; T_caller (mob) i

```

Figure 4.9: Operational Rules Support Functions (2)

Figure 4.8 defines four functions. The `fst, ptype()`, is used extensively; it converts a logical type T to a physical type T . The components for class, mutability and mobility are directly copied and the container label is converted to a physical container using `labelToCont()`. The conversion is handled by cases. When the container is a variable, the physical container is simply the heap location of that variable. To obtain that address, the function `locToCont()` automatically follows references.

When converting a label `ofstep (base; f)`, the `base` is recursively processed, yielding a tuple location. From that tuple, the `field f` is accessed, resulting in the desired container. For all other container labels, for example `generic (g)`, the physical container is derived from the mapping M , which is associated with every local configuration.

The fourth function of figure 4.8, `toTuple()`, is used in the operational semantics to automatically follow references if present and return the location of the desired tuple.

Next, there are 5 additional functions in figure 4.5, which are all related to establishing new local configurations and label-to-container mappings. The function `tupleLocalConfig()` creates a new scope for initialization of a new tuple. The `h` parameter in this case must be the heap location of the tuple. The `uniqueID()` function is abstract and generates a globally unique identifier. We rely on global uniqueness to permit a tuple heap location to also serve a double duty as the key for $(cmap)$.

When initializing a new scope for a function call `funcLocalConfig ()` builds the initial local configuration to execute the body of a function. The statement `body` is placed on the frame-stack, the mapping is recorded and a unique identifier is assigned.

The `labelMappings ()` function takes care of building the mapping M for a new function body scope. We have the label-to-label mapping, which maps container labels in the callee's scope to the caller's scope. We also have the mapping M_{caller} , so we can take a callee-relative label, convert it into a caller-relative label and finally use the caller's M_{caller} to obtain the appropriate physical container. These steps are taken for all container labels in `in` to generate M .

4.6 Frame Typing

Each of our frames requires typing. All statements and expressions of the language are frames, and these frames retain their typing as in figures 3.16 and 3.17. However, for the new frames, we need new rules, which in turn need an extended typing environment. Statements and expressions are typed in context Γ and the extended frames are typed within $\Gamma; \Delta; H; L$. However, for the purposes of frame typing, we consider these to be the same judgment with the original statement and expression typing simply ignoring the extra configuration context.

Figures 4.10 and 4.11 define the typing rules for our new frames. Most rules are straightforward in that they validate their component parts. A few rules make use of judgments that have not been defined yet. For reference, the judgment `match` is defined in figure 4.24 and `copy-map-ok` is in figure 4.25. Although a hole `()` is not a frame, it needs to be typed, as it represents a future result. To do so, we add its type to Δ and use the rule `T-Hole` to type it.

The rule `T-Location` in particular does more than just type a location. It also establishes consistency between the location's logical and physical types. This type of consistency requirement is discussed in full in section 4.9, but for now it suffices to understand that frames are typed in the configuration typing context Δ and must be consistent with their corresponding physical types.

4.6.1 Program Type Annotations

All frames are typed prior to applying any semantic rules. As part of the typing process, we create an annotated version of the program where each frame is tagged with the typing

$$\begin{array}{c}
\text{T-Hole} \\
\frac{}{;;; H;L \text{ ` } : T} \\
\\
\text{T-Void} \\
\frac{}{;;; H;L \text{ ` } \text{void} : \text{Void}} \\
\\
\text{T-Location} \\
\frac{(\text{ ` }) = T = \text{ptype} (; H;L;T) \quad ; H;L \text{ ` } T \text{ t-match } T}{;;; H;L \text{ ` } h \text{ ` }; T_i : T} \\
\\
\text{T-CopyEnv} \quad \text{T-Location-CopyEnv} \\
\frac{\text{ ` } A \text{ copy-map-ok}}{;;; H;L \text{ ` } A : \text{CopyEnv}} \quad \frac{\text{ ` } A \text{ copy-map-ok} \quad ;;;; H;L \text{ ` } h \text{ ` }; T_i : T}{;;; H;L \text{ ` } hh \text{ ` }; T_i; A_i : hT; \text{CopyEnv}} \\
\\
\text{T-Pop-Local} \\
\frac{v \ 2}{;;; H;L \text{ ` } \text{pop-local} (v) : \text{Void}} \\
\\
\text{T-Copy-Discard-Env} \\
\frac{\text{ ` } A \text{ copy-map-ok} \quad ;;;; H;L \text{ ` } h \text{ ` }; T_i : T}{;;; H;L \text{ ` } \text{copy-discard-env} (hh \text{ ` }; T_i; A_i) : T} \\
\\
\text{T-Init-Symbols} \\
\frac{\overline{T_i = (v_i)} \quad ;;;; H;L \text{ ` } \text{ ` }_{init \ i} : T_{init \ i} \quad (\overline{T_i(\text{kind}) = \text{ref} \wedge T_i \in T_{init \ i}}) \text{ _ } (T_i = T_{init \ i})}{;;; H;L \text{ ` } \text{init-symbols} (h \text{ ` }; h_{init \ i}; T_{init \ i \ i}) : \text{Void}} \\
\\
\text{T-Init-Symbol} \\
\frac{T = (v) \quad ;;;; H;L \text{ ` } \text{ ` }_{init} : T_{init} \quad (\overline{T(\text{kind}) = \text{ref} \wedge T \in T_{init}}) \text{ _ } (T = T_{init}) \quad T(\text{kind}) = \text{ref} \text{ _ } \text{ ` }_{init} \notin \text{null}}{;;; H;L \text{ ` } \text{init-symbol} (v; h_{init \ i}; T_{init \ i}) : \text{Void}} \\
\\
\text{T-Init-Fields} \\
\frac{C = T_{\text{tuple}(\text{class})} \quad \overline{T_{\text{field} \ i} = (C)(f_i)}}{;;; \text{ ` }_{init \ i} : T_{init \ i} \quad (\overline{T_{\text{field} \ i}(\text{kind}) = \text{ref} \wedge T_{\text{field} \ i} \in T_{init \ i}}) \text{ _ } (T_{\text{field} \ i} = T_{init \ i})}{;;; H;L \text{ ` } \text{init-fields} (h_{\text{tuple} \ i}; T_{\text{tuple} \ i}; hf; h_{init \ i}; T_{init \ i}) : T_{\text{tuple}}} \\
\\
\text{T-Init-Field} \\
\frac{C = T_{\text{tuple}(\text{class})} \quad T_{\text{field}} = (C)(f) \quad ;;;; H;L \text{ ` } \text{ ` }_{init} : T_{init} \quad (\overline{T_{\text{field}}(\text{kind}) = \text{ref} \wedge T_{\text{field}} \in T_{init}}) \text{ _ } (T_{\text{field}} = T_{init}) \quad T_{\text{field}}(\text{kind}) = \text{ref} \text{ _ } \text{ ` }_{init} \notin \text{null}}{;;; H;L \text{ ` } \text{init-field} (h_{\text{tuple} \ i}; T_{\text{tuple} \ i}; f; h_{init \ i}; T_{init \ i}) : \text{Void}}
\end{array}$$

Figure 4.10: Frame Typing (1)

$$\begin{array}{c}
\text{T-Copy-Tuple} \\
\frac{\begin{array}{c} ; ; \quad \text{\`}`\`src : T_{src} \quad T_{dst} \quad T_{src} \\ ; ; ; \quad H; L \text{\`}` copy-tuple (T_{dst}; h_{src}; T_{src} i) : T_{dst} \end{array}}{} \\
\\
\text{T-Copy-Tuple-2} \\
\frac{\begin{array}{c} ; ; \quad \text{\`}`\`src : T_{src} \\ \text{\`}` T_{dst} \text{ type-ok-common} \quad T_{dst} \quad T_{src} \quad \text{\`}` A \text{ copy-map-ok} \\ ; ; ; \quad H; L \text{\`}` copy-tuple2 (T_{dst}; h_{src}; T_{src} i; A) : h_{T_{dst}}; \text{CopyEnv} \end{array}}{} \\
\\
\text{T-Copy-Fields} \\
\frac{\begin{array}{c} C = T_{\text{tuple}(\text{class})} \\ ; ; \quad \text{\`}`\`tuple : T_{tuple} \quad \frac{T_{\text{field } i} = (C)(f_i)}{\text{\`}`\`init } i : T_{\text{init } i} \\ (T_{(\text{kind})} = \text{ref} \wedge T_{\text{field } i} e_{T_{\text{init } i}}) _ (T_{\text{field } i} = T_{\text{init } i}) \quad \text{\`}` A \text{ copy-map-ok} \\ ; ; ; \quad H; L \text{\`}` copy-fields (h_{tuple}; T_{tuple} i; hf; h_{init}; T_{init } ii; A) : \text{CopyMap} \end{array}}{} \\
\\
\text{T-copy-init} \\
\frac{\begin{array}{c} C = T_{\text{tuple}(\text{class})} \quad T_{\text{field}} = (C)(f) \quad ; ; ; \quad H; L \text{\`}`\`tuple : T_{tuple} \\ ; ; ; \quad H; L \text{\`}`\`init : T_{init} \quad (T_{(\text{kind})} = \text{ref} \wedge T_{\text{field}} e_{T_{\text{init}}}) _ (T_{\text{field}} = T_{\text{init}}) \\ T_{(\text{kind})} = \text{ref} _ \text{\`}`\`init \text{\`}` null \quad \text{\`}` A \text{ copy-map-ok} \\ ; ; ; \quad H; L \text{\`}` copy-init (h_{tuple}; T_{tuple} i; f; h_{init}; T_{init } i; Ai) : \text{CopyEnv} \end{array}}{} \\
\\
\text{T-Assign-Copied} \\
\frac{\begin{array}{c} ; ; ; \quad H; L \text{\`}`\`lhs; T_{lhs} i : T_{lhs} \\ ; ; ; \quad H; L \text{\`}`\`rhs; T_{rhs} i : T_{rhs} \quad T_{lhs(\text{mobility})} = \text{movable} \quad T_{rhs} = T_{lhs} \\ ; ; ; \quad H; L \text{\`}` assign-copied (h_{lhs}; T_{lhs} i; h_{rhs}; T_{rhs} i) : \text{Void} \end{array}}{}
\end{array}$$

Figure 4.11: Frame Typing (2)

environment that it is typed in. For example, if we have the judgment; $\vdash e : T$, then in the annotated program, we will write e , making both e and T available to the conditions of the rules. By binding T to each frame, we enable further configuration typing judgments to be made.

4.7 Special Configurations

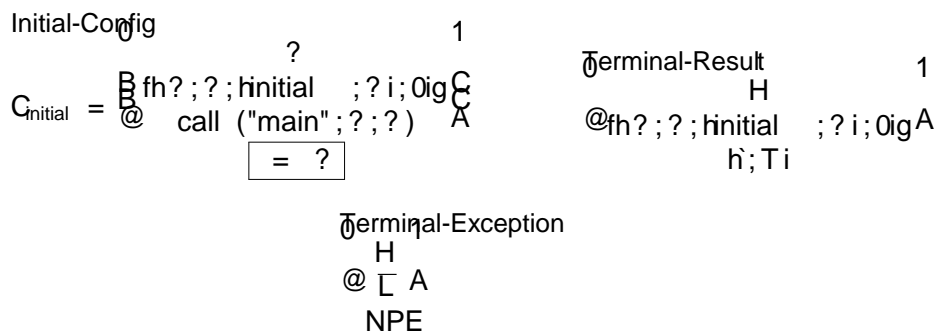


Figure 4.12: Initial and Terminal Configurations

Figure 4.12 defines $C_{initial}$ which is the initial configuration. All programs are expected to have a `main()` function defined. There are no variables in this special initial scope and the heap is empty. Also shown in a boxed region is the initial configuration typing environment. The boxing is intended to separate distinct concerns. Primarily, we are defining the semantics (unboxed), which do not require \vdash . However, as the rules of the operational semantics are defined, we also show how \vdash should be updated to reflect any new typing needed after a rule is applied.

Also in 4.12 are two terminal state patterns. **Terminal-Result** represents normal termination of a program with \vdash as the heap address containing the return value of `main()`. Similarly, **Terminal-Exception** represents termination of a program after an attempt to dereference `null`. No further rules can be applied if the configuration matches one of these patterns.

4.8 Transition Rules

In this section, we define the rules of the operational semantics. The rules are written with the current configuration in the lower left, related to the subsequent configuration in the lower right ($C \rightarrow C'$). When relevant, the typing environment Γ is shown boxed as a fourth element of the configuration's vertical vector and information about uninitialized tuples U appears boxed to the right of the heap component H . Again, this is typing information distinct from the operational semantics. Most rules operate with $U = ?$, and in these cases, we omit any mention of U in the rules. However, for initialization-related frames that do operate with uninitialized tuples, there will be a box describing how the rule impacts U .

Above the line are the conditions required for the rule. Some conditions are simply definitions of symbols used to compose the next configuration. In general, most rules have their conditions formatted into two columns with the more significant ones in the left column and simple definitions on the right. This convention is broken in a few cases for better layout with lengthy conditions.

4.8.1 Machinery

First we look at three simple foundational rules in figure 4.13. First, rule `Next-Frame` runs after a statement rule completes and leaves `void` result as the current frame. A closed-frame is popped from the frame-stack and set as the new current frame.

For processing expression results, the rule `Next-Frame` takes a result, pops an open frame from the frame-stack and replaces the `in` in the open frame with the result. The now-closed frame is set as the new current-frame.

The last of our utility rules is `Pop-Local` which removes a local variable from the local configuration after a `let` statement has completed its execution.

4.8.2 Statements

The next set of rules are for statements and are defined in figure 4.14. The rule `let` statements pushes σ and a `pop-local` frame which removes the variable v after the body of the `let` has completed. The current frame is set to `init-symbol` which defines and initializes the new variable.

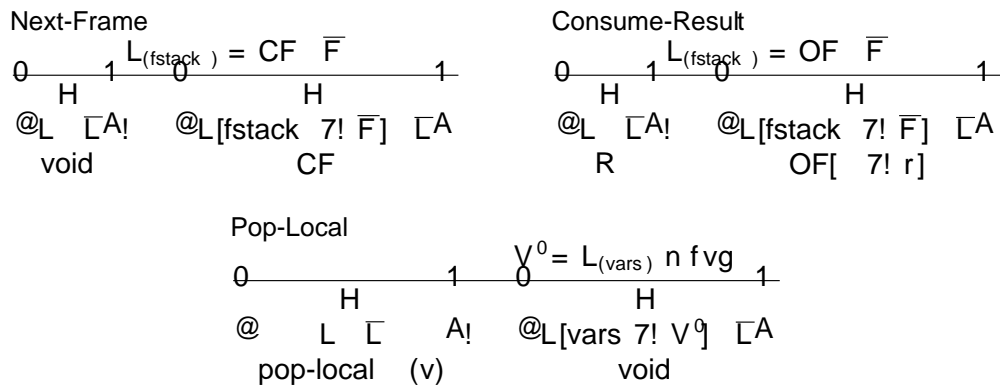


Figure 4.13: Transition Rules Machinery

The rule for `assign-value` uses `copy-tuple` to implement assign-by-value semantics. After the copy is complete, `assign-copied` consumes the result and completes the assignment by updating the heap with the left-hand side. Box updated to reference the new tuple. For reference assignment, `assign-ref` doesn't have assign-by-value semantics so it directly updates the heap to reference the right-hand side tuple.

For `return` statements there are two rules, one for values and one for references. When returning a value, `copy-tuple` is invoked, but it runs in the scope of the calling function and the result of the copy is directly consumed by the calling function. Note the use of the `exportType()` function, which converts the return type's container labels to be relative to the function caller's scope. Again, returning references is similar, but skips the `copy-tuple` steps.

4.8.3 Expressions

Rules for expressions appear in figure 4.15. For `eld access expressions`, `direct-access` first normalizes `tuple` using `toTuple()`, which automatically de-references `tuple` if it is a reference. The `eld` is then found in the heap by looking up `symbol` in the tuple structure. The logic for `var(v)` is also simple. The local configuration holds the local variables in `L(vars)`.

In the rule `Call-Function`, there are a number of critical steps. Recall that one of the most important aspects of this system is the management of container labels when passing parameters across scopes. The first mapping is provided by the program, and it maps labels in the new called function body back to the caller's scope. This `map` is stored in the newly created local configuration `Lfn`. The second `map` `M` is built using the

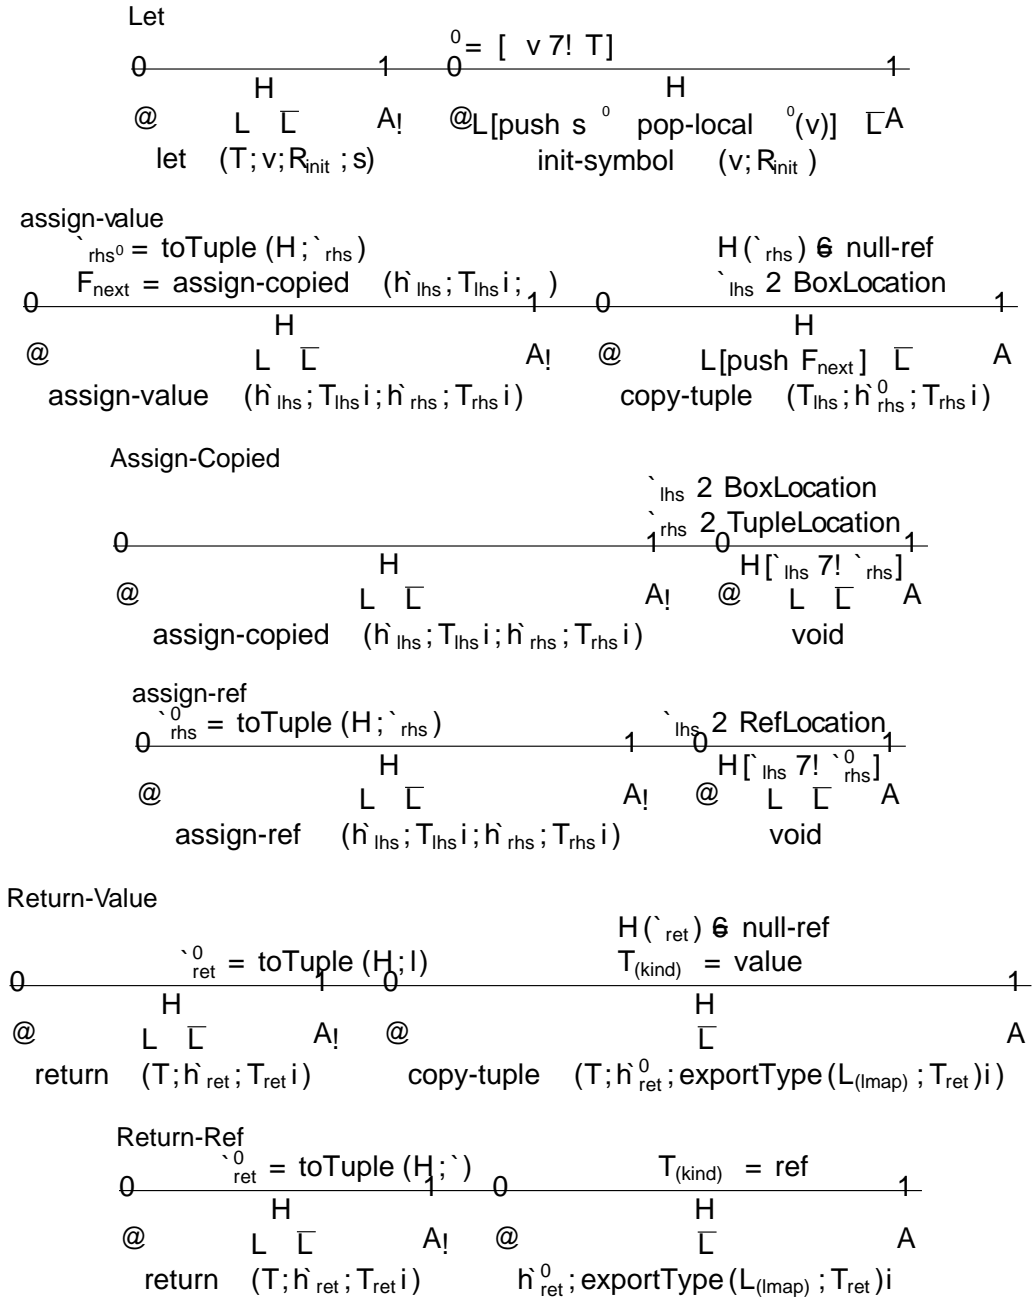


Figure 4.14: Transition Rules Statements

$$\begin{array}{c}
\text{field-access} \\
\frac{\begin{array}{c} \text{H}(\text{tuple}) \notin \text{null-ref} \\ \text{field}(\text{:::}): T_{\text{field}} \end{array}}{\text{field}(\text{tuple}; f)} \quad \frac{\text{H}(\text{tuple}) \notin \text{null-ref}}{\text{field}(\text{:::}): T_{\text{field}}} \\
\frac{\text{H}(\text{tuple}) \notin \text{null-ref}}{\text{field}(\text{tuple}; f)} \quad \frac{\text{H}(\text{tuple}) \notin \text{null-ref}}{\text{field}(\text{:::}): T_{\text{field}}} \\
\text{field}(\text{tuple}; f) \quad \text{field}(\text{:::}): T_{\text{field}} \\
\text{var} \\
\frac{\text{H}(\text{vars})(v)}{\text{var}(v)} \quad \frac{\text{H}(\text{vars})(v)}{\text{var}(v)} \\
\text{var}(v) \quad \text{h}; T_i
\end{array}$$

call-function

$$\begin{array}{c}
L_{fn} = \text{funcLocalConfig}(F; ; s^{fn}) \quad \frac{f F g = \text{function}(_; S; s)}{S = \text{hp}; T; _ i} \\
\frac{\text{H}(\text{init}) \notin \text{null-ref}}{\text{init-symbols}(fn(\text{hp}; \text{h}_{init}; T_{init}^0; T))} \quad \frac{\text{H}(\text{init}) \notin \text{null-ref}}{\text{init-symbols}(fn(\text{hp}; \text{h}_{init}; T_{init}^0; T))} \\
\text{init-symbols}(fn(\text{hp}; \text{h}_{init}; T_{init}^0; T)) \\
[\text{L}_{fn}(\text{id}) \text{map} \text{labelMappings}(_; H; L; _)]
\end{array}$$

Initialize

$$\begin{array}{c}
\text{H}(\text{new}) \notin \text{null-ref} \quad \frac{\text{H}(\text{new}) \notin \text{null-ref}}{\text{new} = \text{uniqueID}()} \\
\text{new} = \text{uniqueID}() \\
= \text{container-of-tuple} \\
T_{\text{new}} = \text{href}; C; ; \text{mutable}; \text{movable} \\
\frac{\text{H}(\text{new}) \notin \text{null-ref}}{\text{init-fields}(h_{\text{new}}; T_{\text{new}}; hf; h_{init}; T_{init}^0; T)} \quad \frac{\text{H}(\text{new}) \notin \text{null-ref}}{\text{init-fields}(h_{\text{new}}; T_{\text{new}}; hf; h_{init}; T_{init}^0; T)} \\
\text{init-fields}(h_{\text{new}}; T_{\text{new}}; hf; h_{init}; T_{init}^0; T) \\
[\text{new} \text{map} \text{tupleMappings}(_; H; L; ; \text{new})] \\
\text{tupleLocalConfig}(_; \text{new}) \\
\text{H}^0 = \text{H}[\text{new} \text{map} \text{tupleMappings}(_; H; L; ; \text{new})] \\
\text{U} = f \text{new} g \\
[\text{new} \text{map} \text{tupleMappings}(_; H; L; ; \text{new})] \\
[\text{new} \text{map} \text{tupleMappings}(_; H; L; ; \text{new})]
\end{array}$$

Figure 4.15: Transition Rules Expressions

labelMappings () function, and it translates container labels into physical containers for the new scope.

With both of these maps established, the function call operation begins with the frame init-symbols which takes each passed parameter from $\overline{h_{init}; T_{init}}$ and defines the respective symbols within the new scope. The function funcLocalConfig () places the function body into the new frame-stack.

The rule for init () initializes a new tuple using the passed parameters as the initial values. It works like a function call even though there is no function body to run. Here, the frame init-fields handles the parameters. The function tupleLocalConfig performs the same role as funcLocalConfig () and tupleMappings performs the same role as labelMappings ().

The unique portion of init () is the allocation of a new heap location $\overline{n_{new}}$ and its insertion into U recognising it as an uninitialized tuple. When the init-fields operation completes, it will remove $\overline{n_{new}}$, and U will be empty again.

4.8.4 Parameter Passing and Symbol Initialization


Figure 4.16 details rules dealing with symbol initialization and passing parameters to functions. The frame init-symbols takes a list of variables to be initialized and processes the first variable. Either Init-Symbols-Value or Init-Symbols-Ref can apply depending on the kind of the first variable. When initializing a value, a copy must be made using copy-tuple, and in this case, init-symbol is pushed onto the stack to consume the copy and do the variable assignment. After init-symbol is complete, the remainder of the variables are processed by another init-symbols frame, which is also pushed to the stack. When the variable list becomes empty, rule Init-Symbols-Value returns void which completes the operation.

Rule Init-Symbols-Ref is simpler, as it bypasses the copy-tuple step. Rules Init-Symbol-Value and Init-Symbol-Ref add the symbol v to the local configuration and map the variable to a newly allocated and initialized Box or Ref. The mapping from v to Box=Ref is fixed and access to the actual value of a variable requires a heap lookup.


4.8.5 Field Initialization

Field initialization behaves similarly to parameter initialization. Instead of defining symbols in the local configuration, the tuple construct in the heap has fields added to it


Init-Symbols-Value




Init-Symbols-Ref



Init-Symbols-Complete



Init-Symbol-Value



Init-Symbol-Ref

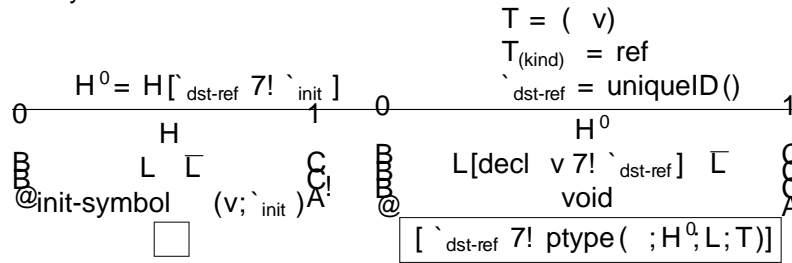


Figure 4.16: Transition Rules Symbol Initialization

incrementally. When all elds have been initialized, the tuple is removed from the set of uninitialized tuples U . Rules for eld initialization are defined in two figures: 4.17 and 4.18. Each of the five rules is similar to its counterpart for parameter initialization. Also similar is the immutability of a tuple as each eld name maps to a location in the heap for that eld.

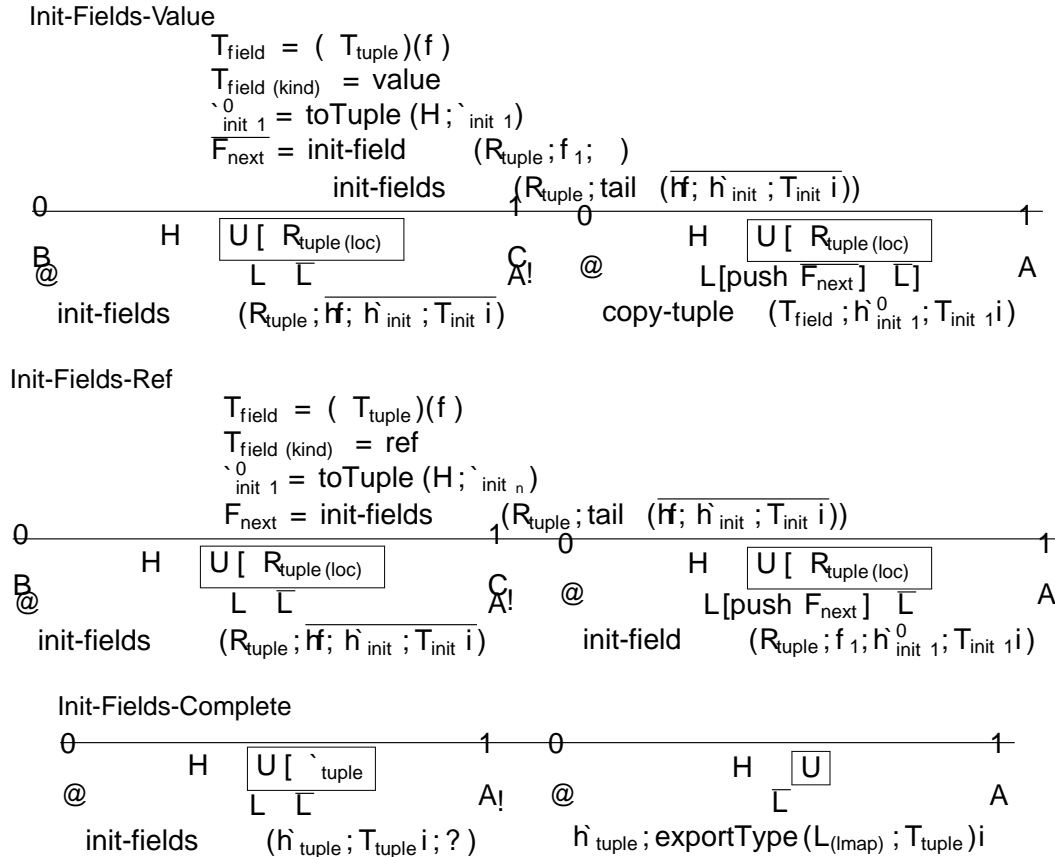



Figure 4.17: Transition Rules Field Initialization (1)

Init-Field-Value



Init-Field-Ref





Figure 4.18: Transition Rules Field Initialization (2)

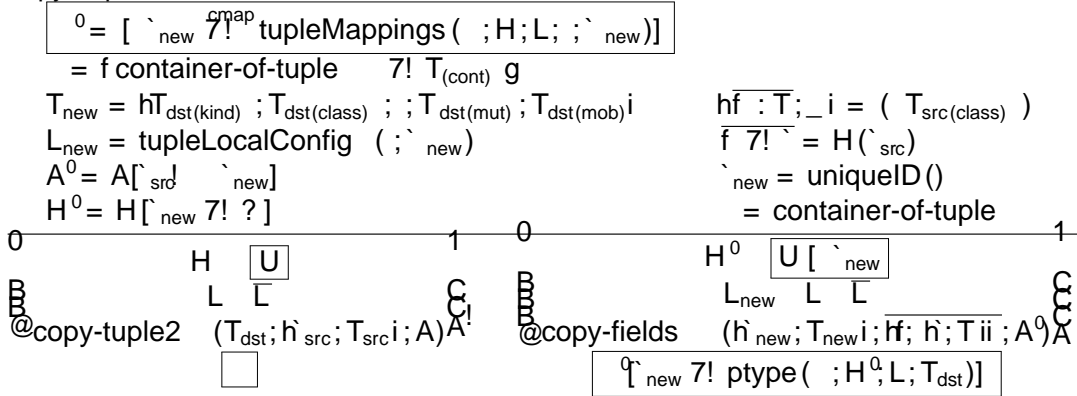
4.8.6 Object Copy Rules

The rules to copy objects are a third instance of the initialization strategy we've discussed with parameter and eld initialization. We are initializing elds, so the greatest similarity is to tuple initialization. However, there is also another layer of complexity with copying, because we carry an accumulated record of all tuples that have been copied as we recursively copy a self-contained object.

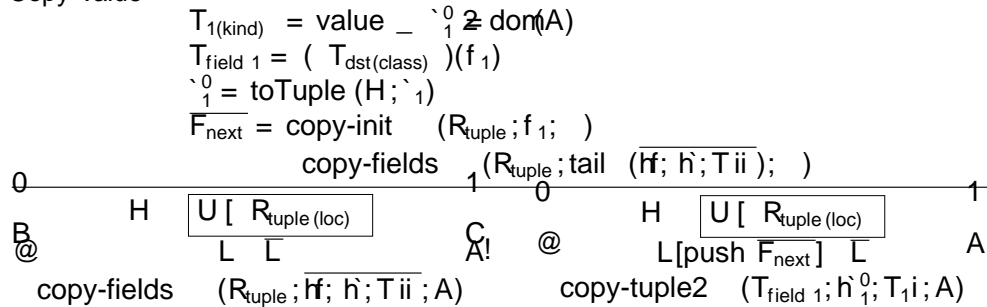
The copy-map_A has tuple heap locations as its domain and range. The domain represents the source tuple being copied and the range are the new tuples. At the completion of the copy process, the domain of α_A will contain the original tuple to copy plus all tuples reachable from that tuple. The copy-map is used to make sure that source tuples are only copied once. When a nested tuple is encountered a second time through a second alias, the copy-map is used to determine the appropriate tuple location that should be used for the respective alias in the copy. As an example, if a self-contained tuple contains a cycle of references within it, the use of the map will prevent an infinite loop.



Copy-Tuple-2



Copy-Value



Copy-Ref




Figure 4.19: Transition Rules Copy Tuple (1)

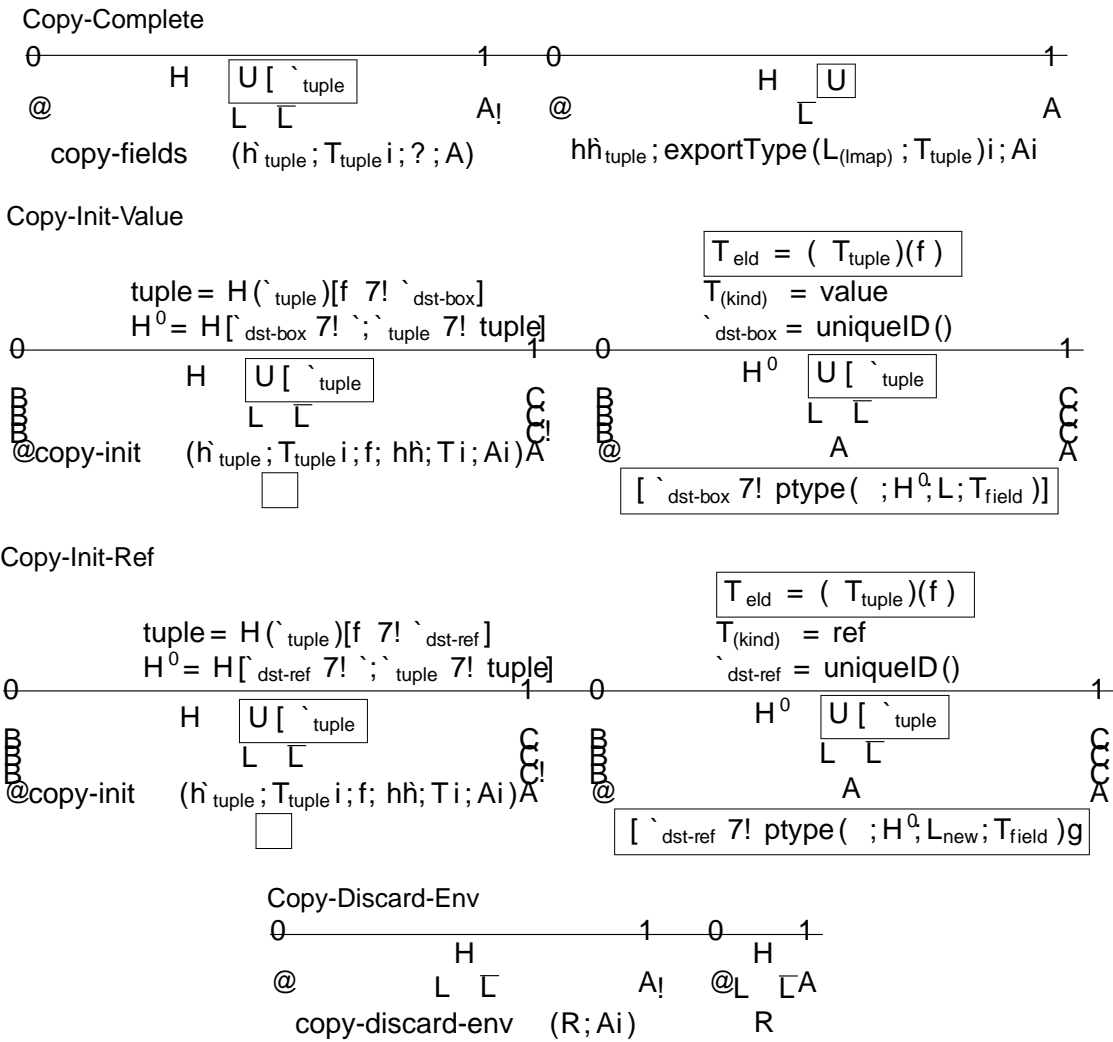


Figure 4.20: Transition Rules Copy Tuple (2)

4.8.7 Null De-Reference Guards

There are a number of places where the transition rules require access to a tuple. In these cases, if `null-ref` were used as the tuple, the machine would break. To prevent this, there are a number of rules defined in figure 4.21 to trap null violations and safely halt the machine. There is one rule for each frame that cannot tolerate `null-ref`.

Assign-Value-NPE

$$\begin{array}{c}
 0 \xrightarrow{H(\hat{r}_{\text{rhs}}) = \text{null-ref}} 1 \quad 0 \xrightarrow{H} 1 \\
 @ \quad L \sqsubset \quad A! \quad @L \sqsubset A \\
 \text{assign-value} \quad (\hat{h}_{\text{lhs}}; T_{\text{lhs}} i; \hat{h}_{\text{rhs}}; T_{\text{rhs}} i) \quad \text{NPE}
 \end{array}$$

Return-Value-NPE

$$\begin{array}{c}
 0 \xrightarrow{H(\hat{r}_{\text{ret}}) = \text{null-ref}} 1 \quad 0 \xrightarrow{H} 1 \\
 @ \quad L \sqsubset \quad A! \quad @L \sqsubset A \\
 \text{return} \quad (T; \hat{h}_{\text{ret}}; T_{\text{ret}} i) \quad \text{NPE}
 \end{array}$$

Field-NPE

$$\begin{array}{c}
 0 \xrightarrow{H(\hat{r}_{\text{tuple}}) = \text{null-ref}} 1 \quad 0 \xrightarrow{H} 1 \\
 @ \quad L \sqsubset \quad A! \quad @L \sqsubset A \\
 \text{field} \quad (\hat{h}_{\text{tuple}}; T_{\text{tuple}} i; f) \quad \text{NPE}
 \end{array}$$

Init-Symbols-NPE

$$\begin{array}{c}
 T = (v_1) \quad T_{(\text{kind})} = \text{value} \quad H(\hat{r}_{\text{init } 1}) = \text{null-ref} \\
 0 \xrightarrow{H} 1 \quad 0 \xrightarrow{H} 1 \\
 @ \quad L \sqsubset \quad A! \quad @L \sqsubset A \\
 \text{init-symbols} \quad (\overline{h}; \hat{h}_{\text{init}}; T_{\text{init}} ii) \quad \text{NPE}
 \end{array}$$

Init-Fields-NPE

$$\begin{array}{c}
 T = (v_1) \quad T_{(\text{kind})} = \text{value} \quad H(\hat{r}_{\text{init } 1}) = \text{null-ref} \\
 0 \xrightarrow{H} 1 \quad 1 \quad 0 \xrightarrow{H} 1 \\
 @ \quad L \sqsubset \quad A! \quad @L \sqsubset A \\
 \text{init-fields} \quad (\hat{h}_{\text{tuple}}; T_{\text{tuple}} i; \overline{h}; \hat{h}_{\text{init}}; T_{\text{init}} i) \quad \text{NPE}
 \end{array}$$

Copy-Fields-NPE

$$\begin{array}{c}
 T = (v_1) \quad T_{(\text{kind})} = \text{value} \quad H(\hat{r}_{\text{init } 1}) = \text{null-ref} \\
 0 \xrightarrow{H} 1 \quad 1 \quad 0 \xrightarrow{H} 1 \\
 @ \quad L \sqsubset \quad A! \quad @L \sqsubset A \\
 \text{copy-fields} \quad (\hat{h}_{\text{tuple}}; T_{\text{tuple}} i; \overline{h}; \hat{h}; T_{ii}; A) \quad \text{NPE}
 \end{array}$$

Figure 4.21: Transition Rules Null Pointer Violations

4.8.8 Decomposition Rules

Our final set of transition rules in figure 4.22 and 4.23 take care of decomposing complex frames into discrete operations. Sub-expressions are removed from their parent frame, leaving a hole () in place of the expression. The expression is placed on the frame-stack and becomes a result after it is executed. Some frames contain a list of expressions (e.g., call ()) and each one is extracted in turn until all sub-expressions have been replaced with results. At this stage, the completed frame is ready to be applied to one of previous rules we've discussed.

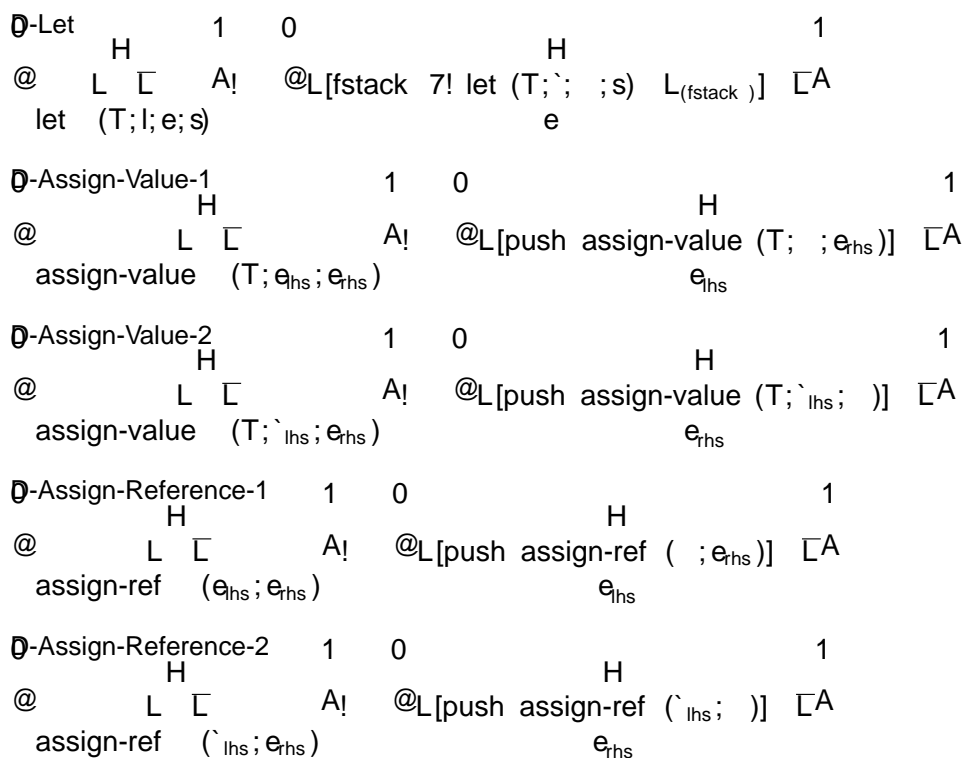


Figure 4.22: Transition Rules Decompositions (1)

$$\begin{array}{c}
\text{D-Return} \quad \begin{array}{c} 1 \quad 0 \quad 1 \\ \text{H} \end{array} \\
@ \quad \begin{array}{c} L \quad \sqsubset \quad A! \\ \text{return } (e) \end{array} \quad \begin{array}{c} \text{H} \\ @L[\text{push return } (e)] \quad \sqsubset A \end{array} \\
\text{D-field} \quad \begin{array}{c} 1 \quad 0 \quad 1 \\ \text{H} \end{array} \\
@ \quad \begin{array}{c} L \quad \sqsubset \quad A! \\ \text{field } (e;f) \end{array} \quad \begin{array}{c} \text{H} \\ @L[\text{push field } (e;f)] \quad \sqsubset A \end{array} \\
\text{D-Call-Function} \quad \begin{array}{c} 1 \quad 0 \quad 1 \\ \text{H} \end{array} \\
@ \quad \begin{array}{c} L \quad \sqsubset \quad A! \\ \text{call } (F; \text{e}_1 \dots \text{e}_n) \end{array} \quad \begin{array}{c} \text{H} \\ @L[\text{push call } (F; \text{e}_1 \dots \text{e}_n)] \quad \sqsubset A \end{array} \\
\text{D-Initialize-Tuple} \quad \begin{array}{c} 1 \quad 0 \quad 1 \\ \text{H} \end{array} \\
@ \quad \begin{array}{c} L \quad \sqsubset \quad A! \\ \text{init } (C; \text{e}_1 \dots \text{e}_n) \end{array} \quad \begin{array}{c} \text{H} \\ @L[\text{push init } (C; \text{e}_1 \dots \text{e}_n)] \quad \sqsubset A \end{array}
\end{array}$$

Figure 4.23: Transition Rules Decompositions (2)

4.9 Configuration Typing

With the transition rules defined, we turn our attention back to typing. We typed our frames in section 4.6 and now we address the remainder of the configuration. Figure 4.24 contains the `config-ok` judgment, which validates the entire configuration. This includes fully typing the heap, all frame-stacks and all local variables. Establishing `config-ok` relies on many other rules to type each of the various components. First, we'll look at the `heap-ok` judgment.

The heap is comprised of a set of abstract locations, and one of the conditions of `heap-ok` is that every location have a corresponding typing in Σ . For each of these locations we apply the `location-ok` judgment to ensure that the heap and its typing are in agreement.

There are four rules, also in figure 4.24, that establish `location-ok`. The rule `Object-OK` handles all tuple locations that are fully initialized (not in U). It looks up the set of fields from Σ and confirms that each field is present and the heap location associated with the field has a compatible type using the `field-match` judgment which we'll detail soon.

$$\frac{\text{Configuration-OK} \quad ; ; U \setminus H \text{ heap-ok} \quad ; ; H_0 \setminus L[\text{pysh CF}] \quad \Gamma : (\text{Void} \quad T; U \quad ?)}{; ; U \setminus @L \quad \Gamma A \text{ config-ok} \quad \text{CF}}$$

$$\frac{\text{Heap-OK} \quad \forall i \in \text{dom}(H) : ; ; H \setminus i \text{ location-ok} \quad \text{dom}(H) = \text{dom}(H)}{; ; U \setminus H \text{ heap-ok}}$$

$$\frac{\text{Object-OK} \quad \forall \text{ TupleLocation } \setminus \in U \quad T_{\text{obj}} = (\setminus) \quad \forall f_i \in (T_{\text{obj(class)}}) : ; ; H \setminus \setminus (T_{\text{obj(class)}})(f_i) \text{ f-match } (H(\setminus)(f_i))}{; ; U; H \setminus \setminus \text{ location-ok}}$$

$$\frac{\text{Object-OK-Uninit} \quad \forall \text{ TupleLocation } \setminus \in U \quad T_{\text{obj}} = (\setminus) \quad \forall f_i \in \text{dom}(H(\setminus)) : ; ; H \setminus \setminus (T_{\text{obj(class)}})(f_i) \text{ f-match } (H(\setminus)(f_i))}{; ; U; H \setminus \setminus \text{ location-ok}}$$

$$\frac{\text{Null-Ref-OK} \quad \forall \text{ RefLocation } \setminus \quad H(\setminus) = \text{null-ref}}{; ; U; H \setminus \setminus \text{ location-ok}}$$

$$\frac{\text{Ref/Box-OK} \quad H(\setminus) \notin \text{null-ref} \quad \forall \text{ SymbolLocation } \setminus \quad T_{\text{ref}} = (\setminus) \quad T_{\text{tuple}} = (H(\setminus)) \quad T_{\text{ref}} \text{ p-match } T_{\text{tuple}}}{; ; U; H \setminus \setminus \text{ location-ok}}$$

Figure 4.24: Configuration Typing Rules (1)

For uninitialized objects, the rule `Object-OK-Uninit` is used and the validation is similar, but the set of fields that are type checked are only the fields that are currently present in the heap. Next we have a special rule, for reference heap locations that are set to `null-ref`. This case requires no further validation, as all reference types are permitted to be `null-ref`.

For the normal reference case as well as boxes, rule `Box-OK` is used. This validates that the physical type of the box or reference matches the physical type of the tuple that is referenced by the box or reference.

In figure 4.25, we have the next set of rules. The `f-match` judgment is defined in the rule `Field-Match`, and contains the condition `within`, which is perhaps the most important condition for establishing proper containment properties. The `within` judgment is defined in rule `Containment` and requires that for a container c_1 to be considered `within` another container c_2 , it either must be the same container ($c_1 = c_2$) or the container of c_1 must be `within` c_2 . Nested containers unwrap as many layers of containers as needed to find the container of interest. Note that we are considering physical containers, which are identified by a heap location scope identifier within the configuration. Establishing this properly means that containers behave the way our common sense dictates they should. This means that no reference that is inside the container may refer to an object that is outside container.

Returning to the judgment `f-match`, it relies on three additional judgments: `t-match`, to compare logical to physical types; `p-match`, which checks that two physical types match, and finally `c-match`, which checks that containers match. We'll now look at this sequence of judgments in more detail.

The second condition `off-match` is `t-match` which verifies that the logical type T is compatible with its physical type T . The `t-match` judgment uses the `type()` function to convert it into a physical type. Now, with two physical types, we validate them using the `p-match` judgment.

`p-match` validates the sub-components of the physical type. For a class there must be an exact match. For mutability, there must either be a match or the left-hand side container must be read-only. This mutability logic mirrors the type rules of the language. The third condition validates the containers and delegates this to our last matching rule `c-match`.

The rule `Container-Match` takes two physical containers ref and $value$, and determines if they are compatible. Either $ref = value$ or ref must be set to `unknown-cont`, which indicates that it is compatible with any container.

Field-Match

$$\frac{L = \text{tupleLocalConfig } (?;`) \quad \text{` } T_{(\text{cont})} \text{ within } (\text{` tuple})_{(\text{cont})} ; H; L \text{ ` } T \text{ t-match } T}{; H; \text{` tuple ` } T \text{ f-match } T}$$

Type-Match

$$\frac{\text{ptype}(; H; L; T) \text{ p-match } T}{; H; L \text{ ` } T \text{ t-match } T}$$

Physical-Match

$$\frac{T_{\text{ref}}(\text{cont}) \text{ c-match } T_{\text{tuple}(\text{cont})} \quad T_{\text{ref}}(\text{mut}) = T_{\text{tuple}(\text{mut})} \text{ -- } T_{\text{ref}}(\text{mut}) = \text{read-only}}{T_{\text{ref}}(\text{class}) = T_{\text{tuple}(\text{class})} \quad T_{\text{ref}} \text{ p-match } T_{\text{tuple}}}$$

Container-Match

$$\frac{\text{ref} = \text{value} \text{ -- } \text{ref} = \text{unknown-cont}}{\text{ref} \text{ c-match } \text{value}}$$

Vars-OK

$$\frac{8v_i \text{ ` } 2 \text{ don}(V): ; \text{ ` } \text{var}(v_i): T_i ; H; L \text{ ` } T_i \text{ t-match } (V(v_i))}{; ; ; L \text{ ` } V \text{ vars-ok}}$$

Containment

$$\frac{1 = 2 \text{ -- } 2 = \text{null-cont} \text{ -- } (1)_{(\text{cont})} \text{ within } 2}{\text{` } 1 \text{ within } 2}$$

Copy-Map-OK

$$\frac{8^i \text{ ` } 2 \text{ don}(A): \text{ ` } i \text{ ` } 2 \text{ ^ } \text{ ` } i \text{ ` } 2 \text{ range}(A) \quad 8^j \text{ ` } 2 \text{ range}(A): \text{ ` } j \text{ ` } 2}{\text{ ` } A \text{ copy-map-ok}}$$

CMap-OK

$$\frac{F_{\text{caller}} = \text{head}(L_{\text{caller}}(\text{fstack})) ; \text{caller} \text{ ` } L_{(\text{ctx}, \text{lmap})} \text{ lmap-ok} \quad = L_{(\text{lmap})} \quad M = (\text{cmap})(L_{(\text{id})}) \quad \text{don}() = \text{don}(M)}{8^i \text{ ` } 2 \text{ don}(M): M(i) = \text{labelToCont}(; H; L; i) = \text{labelToCont}(; H; L_{\text{caller}}; \text{mapLabel}(; i)) ; H \text{ ` } L_{\text{caller}} \text{ cmap-ok } L}$$

Figure 4.25: Configuration Typing Rules (2)

We've completed the `vars-ok` judgment in the rule `Vars-OK`. Like `vars-ok`, each variable must satisfy the `vars-ok` judgments. Unlike `vars-ok`, there is no check needed for `vars-ok`, because local symbols aren't `vars-ok` of contained objects.

The `copy-map-ok` is a simple rule that ensures that when copying an object, none of the tuples reachable from the source tuple appear in the range of `vars-ok`. The converse case must be true as well; none of the new tuples in the range of `vars-ok` can appear in the domain.

The last rule of [figure 4.25](#) is `CMap-OK`. Recall that a container map `M` converts container labels in a scope `L` to physical containers. Also, container maps are built using information from mapping `M` as well as physical container labels from the parent scope `Lcaller`. The judgment, written $\vdash; H \vdash L_{\text{caller}} \text{ cmap-ok } L$, is judging `M` relative to a caller and callee. All labels in `M` must be mapped to physical containers which are in agreement with the physical containers of the corresponding container labels in the caller's scope.

This completes all the judgments needed to validate the heap in our configuration. What remains is the stack of local configurations and the frame-stacks within them.

Local Configuration Typing

The second component of our master `config-ok` judgment is the typing of the local configuration stack, which is typed by the rules in [figure 4.26](#). This judgment works by induction, but first we'll look at how it validates the current local configuration `L`. The local variables in `L` are checked using `vars-ok`. `ScopeL` is created by `scope-head(L)` and these two scopes are used to validate `L`'s container map `M` using `cmap-ok`. The frame-stack of `L` is validated using the frame-stack typing rules in the next section.

The last condition of the local configuration typing is the inductive step to verify the tail of the local configuration stack. The final typing of the stack is of the form $(T \vdash T^{00}; U \vdash ?)$. The first relation $T \vdash T^{00}$ means that this local configuration stack

$$\begin{array}{c}
 \text{T-Local-Config-Stack} \\
 \hline
 \vdash; H \vdash L \text{ vars-ok} \quad \vdash; H \vdash \text{head}(L) \text{ cmap-ok } L \\
 \vdash; H \vdash L \text{ L}_{(\text{fstack})} : (T \vdash T^{00}; U \vdash ?) \quad \vdash; H \vdash L : (T \vdash T^{00}; U \vdash ?) \\
 \hline
 \vdash; H \vdash L : (T \vdash T^{00}; U \vdash ?) \\
 \\
 \text{T-Local-Config-Stack-Empty} \\
 \hline
 \vdash; H \vdash ? : (T \vdash T; U \vdash U)
 \end{array}$$

Figure 4.26: Local Configuration Typing

is a computation taking T as input and producing T^{00} as output. The first T matches the input type of $L_{(fstack)}$ and the T^{00} matches with the inductive typing of the stack tail.

The second component of the typing is $U!$, which indicates that this stack begins with a configuration where the heap locations in U are uninitialized and the computation ends with no uninitialized objects.

Returning to the definition of `config-ok`, notice how it melds the current frame into the current local configuration's frame-stack before typing it. By placing the current frame CF into $L_{(fstack)}$ before typing the local configuration stack, we don't need to special case the typing of CF as it happens automatically as part of the local configuration typing.

Frame-Stack Typing

To type the frame-stacks, we first separate the frames into groups. Open frames are typed differently than closed frames, and frames that work with uninitialized tuples need special treatment. Except for stacks with the `return` on the top, all frame-stacks can be typed generically without specific rules for each frame. Return statements have special treatment to support the case of returning early from a function, and as mentioned earlier, this capability is not taken advantage of in the current system.

$UF_{same} = \{ R, \text{void}, \text{copy-tuple-2}, \text{init-field}, \text{copy-init} \}$
 $UF_{done} = \{ \text{init-fields}, \text{copy-fields}, \text{copy-complete} \}$
 $UF = UF_{done} \cup UF_{same}$

We define sets of frames to split them according to their usage of uninitialized tuples. Frames that type with $U!$ are in UF_{same} and frames that type as $U [\text{!}] U$ are in UF_{done} . No frame-stack type ever needs to add elements to U , because frames that do expand U always push an additional frame to the stack to remove the new uninitialized tuple. This preserves the $U!$ typing.

The frame-stack typing rules appear in figure 4.27. Similar to local configuration typing, frame-stacks are typed as $T!$ ($T^0; U!$ or U^0). The different rules vary based on which frame is on the top of the stack. In general, the top frame is typed directly and the remainder of the stack is typed inductively, with the rule `T-Empty-Stack` terminating the induction. The output of the top frame must be consistent with the input of the inductively typed frames. Frames are logically typed as \mathbb{F} , and the frame-stack uses physical type \mathbb{F} . The function `ptype()` is found in each rule to do the conversion.

Stacks with an open frame OF on top only differ from stacks with closed frames CF in that the input type is always `void` in the rule `T-Closed-Stack`, and it's the type of the

T-Empty-Stack

$$\frac{}{;; H;L \text{ ` } [] : (T! \quad T;U! \quad U)}$$

T-Open-Stack

$$\frac{\begin{array}{l} \text{CF} \text{ \# } \text{return } (_) \quad ; \quad [f : Tg; \quad \text{` } \text{OF} : T^0 \quad ; ; \quad H;L \text{ ` } \text{F} : (T! \quad T^{00};?! \quad ?) \\ T = \text{ptype}(\quad ;H;L;T) \quad T^0 = \text{ptype}(\quad ;H;L;T^0) \end{array}}{;; H;L \text{ ` } \text{OF} \quad \text{F} : (T! \quad T^{00};?! \quad ?)}$$

T-Closed-Stack

$$\frac{\begin{array}{l} \text{CF} \text{ \# } \text{UF} \quad \text{CF} \text{ \# } \text{return } (_) \quad ; ; ; \quad H;L \text{ ` } \text{CF} : T \\ T = \text{ptype}(\quad ;H;L;T) \quad ; ; \quad H;L \text{ ` } \text{F} : (T! \quad T^0;?! \quad ?) \end{array}}{;; H;L \text{ ` } \text{CF} \quad \text{F} : (\text{Void} \quad T^0;?! \quad ?)}$$

T-Immediate-Return

$$\frac{\begin{array}{l} ; ; ; \quad H;L \text{ ` } \text{return } (\text{CF}) : T_{\text{callee}} \\ T = \text{exportType}(L_{(\text{imap})}; T_{\text{callee}}) \quad T = \text{ptype}(\quad ;H;L;T) \end{array}}{;; H;L \text{ ` } \text{return } (\text{CF}) \quad _ : (\text{Void} \quad T;?! \quad ?)}$$

T-Unit-Closed-Stack

$$\frac{\begin{array}{l} \text{CF} \text{ \# } \text{UF}_{\text{same}} \quad ; ; ; \quad H;L \text{ ` } \text{CF} : T \\ T = \text{ptype}(\quad ;H;L;T) \quad ; ; \quad H;L \text{ ` } \text{F} : (T! \quad T^0;U! \quad U^0) \end{array}}{;; H;L \text{ ` } \text{CF} \quad \text{F} : (\text{Void} \quad T^0;U! \quad U^0)}$$

T-Unit-Open-Stack

$$\frac{\begin{array}{l} \text{OF} \text{ \# } \text{UF}_{\text{same}} \quad ; \quad [f : Tg; \quad \text{` } \text{OF} : T^0 \quad ; ; \quad H;L \text{ ` } \text{F} : (T! \quad T^{00};U! \quad U^0) \\ T = \text{ptype}(\quad ;H;L;T) \quad T^0 = \text{ptype}(\quad ;H;L;T^0) \end{array}}{;; H;L \text{ ` } \text{OF} \quad \text{F} : (T! \quad T^{00};U! \quad U^0)}$$

T-Unit-Closed-Stack-Done

$$\frac{\begin{array}{l} \text{CF} \text{ \# } \text{UF}_{\text{done}} \quad ; ; ; \quad H;L \text{ ` } \text{CF} : T \quad ; ; \quad H;L \text{ ` } \text{F} : (T! \quad T^0;U! \quad U^0) \\ T = \text{ptype}(\quad ;H;L;T) \quad T^0 = \text{ptype}(\quad ;H;L;T^0) \end{array}}{;; H;L \text{ ` } \text{CF} \quad \text{F} : (\text{Void} \quad T^0;U[f \text{ ` } g \quad U^0)}$$

T-Unit-Open-Stack-Done

$$\frac{\begin{array}{l} ; \quad [f : Tg; \quad ; \quad H;L \text{ ` } \text{CF} : T^0 \quad ; ; \quad H;L \text{ ` } \text{F} : (T! \quad T^{00};U! \quad U^0) \\ T = \text{ptype}(\quad ;H;L;T) \quad T^0 = \text{ptype}(\quad ;H;L;T^0) \end{array}}{;; H;L \text{ ` } \text{OF} \quad \text{F} : (T! \quad T^{00};U[f \text{ ` } g \quad U^0)}$$

Figure 4.27: Frame-Stack Typing

open frame in the rule $T\text{-Open-Stack}$.

The rules $T\text{-Uninit-Closed-Stack}$ and $T\text{-Uninit-Open-Stack}$ mirror the rules for working with frames that always run with $U = ?$. The only difference is that they type with a free variable U rather than explicitly specifying $?$.

The last two rules handle frames that remove an element from U . These are rules that complete an initialization operation. Here, there is an extra free variable which is removed from U . Otherwise, these rules operate like the others.

This completes the frame-stack typing, which is also the final set of rules needed to establish config-ok . With our configuration fully typed, the next chapter will prove the soundness of the operational semantics.

Chapter 5

Type Safety

Here, we establish type safety for our container operational semantics and show that container restrictions expressed by the type rules are indeed obeyed during execution. Given a properly typed program and a configuration initialized to `Initial-Config` as defined in section 4.8.1, the machine will either reach a terminal state or run forever.

Recall that the specification of the operational semantics in chapter 4 simultaneously defined the semantics as well as configuration typing rules which were segregated into boxed regions. We will write $\text{hC}; \text{U} \mid \text{h} \text{ C}^0; \text{U}^0$ to refer to the combined relation defined by the rules of chapter 4. In general, arguments will begin with a valid configuration, a configuration typing, as well as the set of uninitialized heap locations, and will show that the subsequent state is also valid. This relation is also a function, however, we do not require determinism for this proof and we will ignore this fact.

Following the approach developed by Wright and Felleisen [20], progress and preservation theorems are presented. Note that the key novel property of this system that we are proving is the enforcement of container restrictions. This property is encoded into the configuration typing; the `within` judgment defined in the rule `Containment` gives us the result we want. Showing that the `config-ok` property holds establishes the containment result as well. In other words, any program that satisfies the container type rules is guaranteed to execute with correct containment.

5.1 Progress and Preservation

Theorem 5.1.1 Progress

If $\exists C$ config-ok then the current frame of C will either be in a terminal state or there exists a rule to advance the configuration such that $\exists C' ; UI h \ C'; \sigma, U \sigma$.

Proof: First we make the observation that for each frame defined in figure 4.3 there is at least one rule defined in section 4.8. However, many of the rules have conditions that must be satisfied, and it must be shown that for each frame, the conditions will indeed be satisfied.

First we'll address sub-expressions. For every frame that contains sub-expressions there exists a decomposition rule in section 4.8.8. Further, none of these rules impose any conditions. For the remainder of the frames we will discuss, it is assumed that all sub-expressions have been evaluated and will appear as results in their respective frames.

There are 6 frames that are handled by rules with no conditions: `void`, `null`, `result`, `statement-seq`, `copy-tuple` and `copy-discard-environment`. The remainder of the frames will be addressed individually.

`pop-local` : By rule `T-Pop-Local` variable v and by `Vars-OK` v is present in the local configuration.

`assign-value` : The rule `Assign-Value-NPE` catches any null de-reference attempt. The type rule `T-Assign-Value` asserts that the left-hand-side is `movable` and a value which ensures that lhs is a `BoxLocation`. `T-Location` ensures lhs and rhs are valid heap addresses.

`assign-ref` : The type rule `T-Assign-Ref` asserts that the left-hand-side is `movable` and a `ref` which ensures that lhs is a `RefLocation`. `T-Location` ensures lhs and rhs are valid heap addresses.

`return` : Either the rule `Return-Value` or `Return-Ref` is used depending on the kind of return value and `Return-Value-NPE` will protect against null de-reference. The type rule `T-Location` ensures ret is a valid heap address.

`field` : The rule `Field-NPE` catches any null de-reference attempt. The rule `T-Field` requires that the `elfd` exists in the class of tuple . The rules `T-Location` and `heap-ok` ensure that tuple is a valid heap address referring to an object with a `elfd`.

`variable` : The rule `T-Variable` requires that the variable v exists in Σ and `Vars-OK` assures that it is also present in the local configuration.

`call-function` : The rule `T-Function` requires that the function f exists in Σ .

`init` : The rule `T-Init` requires that the class C exists in Σ .

init-symbols : When $T_{(kind)} = \text{value}$, Init-Symbols-NPE will protect against null de-reference. Heap locations \hat{s}_{init} are valid by T-Location .

init-symbol : The rule T-Init-Symbol ensures $v \neq \text{null}$ and that there is no null de-reference.

init-fields : When $T_{(kind)} = \text{value}$, Init-Fields-NPE will protect against null de-reference. Heap locations in \hat{n}_{init} are valid by T-Location .

init-field : The rule T-Init-Field ensures that f is a field of τ_{tuple} and that there is no null de-reference.

copy-tuple2 : The rule T-Copy-Tuple-2 ensures that T_{dst} is a valid type and T-Locations ensures that \hat{s}_{src} is valid.

copy-fields : When $T_{(kind)} = \text{value}$, Copy-Fields-NPE will protect against null de-reference. Heap locations in \hat{n}_{init} are valid by T-Location .

copy-init : The rule T-copy-init ensures that f is a field of τ_{tuple} and that there is no null de-reference.

Therefore, all well typed frames have corresponding rules to process them and progress is assured.

Theorem 5.1.2 Preservation

If $\vdash ; U \vdash C \text{ config-ok}, hC; ; U \vdash h \ C^0, \ \hat{q}; U \hat{q}$
then $\vdash ; \ \hat{q}; U^0 \vdash C^0_{(heap)} \text{ config-ok}$

Proof: By lemma 5.2.1, the heap-ok condition is satisfied in C^0 , and the local configuration stack is satisfied by lemma 5.3.1. Therefore, config-ok is preserved in C^0 .

5.2 Heap Validation Lemmas

Lemma 5.2.1 Heap OK

If $\vdash ; U \vdash C \text{ config-ok}, hC; ; U \vdash h \ C^0, \ \hat{q}; U \hat{q}$
then $\vdash ; \ \hat{q}; U^0 \vdash C^0 \text{ heap-ok}$

Proof: By lemma 5.2.3, we know that all non-modified heap locations are properly typed. Lemma 5.2.4 establishes the typing for newly allocated heap locations. And finally, lemma 5.2.6 shows that modified heap locations preserve their typing. Therefore, we know all heap locations in C^0 are location-ok .

To establish that $\text{dom}(H) = \text{dom}(H^0)$, we observe that for the eight rules that expand the heap (Init-Symbol-Value, Init-Symbol-Ref, Init-Field-Value, Init-Field-Ref, Copy-Tuple-2, copy-init-Value, copy-init-Ref and Initialize), each makes equivalent additions to both H and H^0 . No operation removes elements from either set.

Lemma 5.2.2 Immutable Heap Locations

If $\vdash C \text{ config-ok}, hC; \vdash U \text{ h } C^0, U^0$
then $\exists H$ where $\not\subseteq U$ and $(\cdot)_{(\text{mobility})} = \text{fixed} : H^0(\cdot) = H(\cdot)$

Proof: We consider the two rules in the operational semantics that modify existing heap locations: Assign-Copied and Assign-Ref. Both of these rules are typed with $\text{hs} : T_{(\text{mobility})} = \text{movable}$, which implies that the physical type $\tau = (\cdot)$ will match with $T_{(\text{mobility})} = \text{movable}$ since the $\text{type}()$ function preserves the mobility property. Therefore, these assignment rules cannot change the values stored in fixed locations.

Lemma 5.2.3 Non-Interference

If $\vdash C \text{ config-ok}, hC; \vdash U \text{ h } C^0, U^0, H = C_{(\text{heap})}, H^0 = C^0_{(\text{heap})}$
then $\exists H$ where $H^0(\cdot) = H(\cdot) : \vdash U; H^0 \text{ location-ok}$

Proof: config-ok implies that \cdot is location-ok in C . Of all the conditions in location-ok across its four rules, the only one that depends on heap contents at locations other than \cdot is f-match . For initialized tuples, the f-match judgment was satisfied in C and continues to hold in C^0 by lemma 5.5.3 for all fields of \cdot .

Otherwise, if $\cdot \in U$, we know we are currently initializing the fields of tuple \cdot in an order that respects their dependencies. This means any previous field cannot depend on the contents of subsequent fields and no other heap locations are modified during initialization of a tuple. Thus, f-match must continue to hold for the initialized fields by lemma 5.5.3.

Therefore, for all unmodified heap locations, if the location-ok judgment held in C , it will continue to hold in C^0 .

Lemma 5.2.4 New Locations OK

If $\vdash C \text{ config-ok}, hC; \vdash U \text{ h } C^0, U^0, H = C_{(\text{heap})}, H^0 = C^0_{(\text{heap})}, \cdot \in H^0$ and $\cdot \not\subseteq H$
then $\vdash U; H^0 \text{ location-ok}$

Proof: There are eight rules in the operational semantics that introduce new heap locations. For each, we show that location-ok holds.

Case Init-Symbol-Value

This rule allocates a new box $\dot{x}_{\text{dst-box}} : T$ and initializes it. By the frame typing, we know that $h_{\text{init}} ; T_{\text{init}}$ satisfies T-Location and $T_{\text{dst-box}} = T_{\text{init}}$. Therefore, $T_{\text{init}} = (\text{init}) = \text{ptype}(\ ; H; L; T_{\text{init}})$ and we define $\dot{x}_{\text{dst-box}} \uparrow! \text{ptype}(\ ; H; L; T_{\text{dst-box}})$ in . After updating the heap with $\dot{x}_{\text{dst-box}} \uparrow! \text{init}$, we have satisfied thep-match condition of Ref/Box-OK , because the $T_{\text{dst-box}} = T_{\text{init}}$ equality makes the two invocations of $\text{ptype}()$ identical.

Case Init-Symbol-Ref

This rule allocates a new reference $\dot{x}_{\text{dst-box}} : T$ and initializes it. By the frame typing, we know that $h_{\text{init}} ; T_{\text{init}}$ satisfies T-Location and $T_{\text{init}} \in T_{\text{init}}$. Therefore, $T_{\text{init}} = (\text{init}) = \text{ptype}(\ ; H; L; T_{\text{init}})$ and we define $\dot{x}_{\text{dst-box}} \uparrow! \text{ptype}(\ ; H; L; T_{\text{dst-box}})$ in . After updating the heap with $\dot{x}_{\text{dst-box}} \uparrow! \text{init}$, we satisfy thep-match condition of Ref/Box-OK by lemma 5.2.5.

Case Init-Field-Value

This rule allocates a new box $\dot{x}_{\text{dst-box}} : T_{\text{field}}$ and initializes it. By the frame typing, we know that $h_{\text{init}} ; T_{\text{init}}$ satisfies T-Location and $T_{\text{field}} = T_{\text{init}}$. Therefore, $T_{\text{init}} = (\text{init}) = \text{ptype}(\ ; H; L; T_{\text{init}})$ and we define $\dot{x}_{\text{dst-box}} \uparrow! \text{ptype}(\ ; H; L; T_{\text{field}})$ in . After updating the heap with $\dot{x}_{\text{dst-box}} \uparrow! \text{init}$, we have satisfied thep-match condition of Ref/Box-OK because $T_{\text{eld}} = T_{\text{init}}$ makes the two invocations of $\text{ptype}()$ identical.

Case Init-Field-Ref

This rule allocates a new reference $\dot{x}_{\text{dst-ref}}$ and initializes it. By the frame typing, we know that $h_{\text{init}} ; T_{\text{init}}$ satisfies T-Location and $T_{\text{field}} \in T_{\text{init}}$. Therefore, $T_{\text{init}} = (\text{init}) = \text{ptype}(\ ; H; L; T_{\text{init}})$ and we define $\dot{x}_{\text{dst-box}} \uparrow! \text{ptype}(\ ; H; L; T_{\text{field}})$ in . After updating the heap with $\dot{x}_{\text{dst-box}} \uparrow! \text{init}$, we satisfy thep-match condition of Ref/Box-OK by lemma 5.2.5.

Case Copy-Tuple-2

The Copy-Tuple-2 rule creates a new uninitialized tuple at location \dot{n}_{new} , which trivially satisfies Object-OK-Uninit because currently no elds of this tuple have been initialized yet.

Case Copy-Init-Value

This rule allocates a new box $\dot{x}_{\text{dst-box}} : T_{\text{field}}$ and initializes it. By the frame typing, we know that $h_{\text{init}} ; T_{\text{init}}$ satisfies T-Location and $T_{\text{field}} = T_{\text{init}}$. Therefore $T = (\text{init}) = \text{ptype}(\ ; H; L; T_{\text{init}})$ and we define $\dot{x}_{\text{dst-box}} \uparrow! \text{ptype}(\ ; H; L; T_{\text{field}})$ in . After updating

the heap with dst-box we have satisfied the p-match condition of Ref/Box-OK because $T_{\text{eld}} = T_{\text{init}}$ makes the 2 invocations of $\text{p-type}()$ identical.

Case Copy-Init-Ref

This rule allocates a new reference $\text{eld}_{\text{dst-ref}} : T_{\text{field}}$ and initializes it. By the frame typing, we know that $\text{init}; T_{\text{init}}$ satisfies $T\text{-Location}$ and that $T_{\text{field}} \in T_{\text{init}}$. Therefore, $T = (\text{ }) = \text{p-type}(\text{ }; H; L; T_{\text{init}})$ and we define dst-box $\text{p-type}(\text{ }; H; L; T_{\text{field}})$ in . After updating the heap with dst-ref we have satisfied the p-match condition of Ref/Box-OK by lemma 5.2.5.

Case Initialize

This rule adds a new uninitialized tuple to the heap. Because the tuple has all elds uninitialized, the Object-OK-Uninit rule will be trivially satisfied.

Lemma 5.2.5 Ref-Initializable Implies p-match

If $\text{ }; ; U \text{ } C \text{ config-ok}$ and two types: T_{src} and T_{dst} , defined within scope $L \in C$ such that $T_{\text{dst}} \in T_{\text{src}}$ and $T_{\text{dst}(\text{kind})} = \text{ref}$ then the associated physical types will match; $T_{\text{dst}} \text{ p-match } T_{\text{src}}$.

Proof: With $T_{\text{dst}} = \text{p-type}(\text{ }; H; L; T_{\text{dst}})$ and $T_{\text{src}} = \text{p-type}(\text{ }; H; L; T_{\text{src}})$, we know that the $\text{p-type}()$ function preserves the class and mutability sub-components which have the same requirements in Ref-Initializable and Physical-Match . Finally, the container-labels must either match resulting in matching $\text{p-type}()$ output or $T_{\text{dst}(\text{label})} = \text{unknown}$ which maps to unknown-cont and satisfies p-match . Therefore, $T_{\text{dst}} \in T_{\text{src}}$ implies $T_{\text{dst}} \text{ p-match } T_{\text{src}}$.

Lemma 5.2.6 Modified Locations OK

If $\text{ }; ; U \text{ } C \text{ config-ok}$, $hC; ; U!$ $h \text{ } C^0; \text{ } U^0$, $H = C_{(\text{heap})}$, $H^0 = C_{(\text{heap})}^0$, $\text{ } \in H$; $H^0(\text{ }) \in H(\text{ })$
then $\text{ }; \text{ } U^0; H^0 \text{ } \text{ location-ok}$

Proof: By cases across the eight rules that make heap modifications; for each case, we show that location-ok holds for .

Cases Init-Field-Value, Init-Field-Ref, copy-init-Value and copy-init-Ref

These four cases are all proved by the same argument; the proof for Init-Field-Value is presented.

This rule updates the uninitialized heap location tuple to add a new eld f located at dst-box , which we know is location-ok by lemma 5.2.4. We have $\text{ } = (\text{ }_{\text{dst-box}}) =$

$\text{ptype}(\ ; H^0; L; T_{\text{field}})$ as the right-hand side of the f -match condition in rule `Object-OK-Uninit` for this new field f . Now, the sub-condition f -match requires that $\text{ptype}(\ ; H^0; L; T_{\text{field}})$ p -match T , which is satisfied because the LHS is identical to the RHS.

All other fields in tuple are unchanged, satisfied f -match in H and will continue to satisfy f -match in H^0 by lemma 5.5.3. Therefore, all conditions of `Object-OK-Uninit` are satisfied and `location-ok` holds.

Case Assign-Ref

For `Assign-Ref`, we know that if the RHS is null then `location-ok` is trivially satisfied by rule `Null-Ref-OK`. If non-null then `Ref/Box-OK` is the relevant rule to establish `location-ok` for our modified heap location lhs . We must show T_{lhs} p -match T_{rhs} . From the type rule `T-Assign-Ref`, we know that $T_{\text{lhs}} \supseteq T_{\text{rhs}}$. Further, from the typing of the `Assign-Ref` frame, we know that $\text{lhs}; T_{\text{lhs}}$ and $\text{rhs}; T_{\text{rhs}}$ satisfy the `T-Location` rule. Therefore, $T_{\text{lhs}} = (\text{lhs}) = \text{ptype}(\ ; H; L; T_{\text{lhs}})$ and $T_{\text{rhs}} = (\text{rhs}) = \text{ptype}(\ ; H; L; T_{\text{rhs}})$. After updating the heap with $\text{lhs} \mapsto \text{rhs}$, we can apply lemma 5.4.1, and we have satisfied the p -match condition of `Ref/Box-OK`.

Case Assign-Copied

In this case, we must satisfy `Ref/Box-OK`. From the type rule `T-Assign-Copied`, we know that $T_{\text{lhs}} = T_{\text{rhs}}$ and $\text{lhs}; T_{\text{lhs}}$ and $\text{rhs}; T_{\text{rhs}}$ satisfy the `T-Location` rule. Therefore, $T_{\text{lhs}} = (\text{lhs}) = \text{ptype}(\ ; H; L; T_{\text{lhs}})$, and for the left-hand side $T_{\text{rhs}} = (\text{rhs}) = \text{ptype}(\ ; H; L; T_{\text{rhs}})$. After updating the heap with $\text{lhs} \mapsto \text{rhs}$, we have satisfied the p -match condition of `Ref/Box-OK`, because the $T_{\text{lhs}} = T_{\text{rhs}}$ equality makes the two invocations of `ptype()` identical.

Cases `Init-Fields-Complete` and `Copy-Complete`

These two cases simply remove tuple from U . With all fields populated, the check for `Object-OK-Uninit` covers every field and is identical to `Object-OK`.

5.3 Frame-Stack Lemmas

Lemma 5.3.1 Local Configuration Stack OK

If $\ ; \ ; \ U \ \text{C config-ok}$, $\text{hC}; \ ; \ U \ \text{h} \ C^0, \ ^0; \ U^0, \ C = @_{\text{CF}}^L \ [A, \ C^0 = @_{\text{CF}^0}^L \ [^0A$
then $\ ; \ ; \ H \ \text{L}^0[\text{fstack} \ \text{?} \ \text{CF}^0 \ \text{L}_{(\text{fstack})}] \ \text{L}^0: (\text{T} \ \text{T}^0; \ U \ \text{?})$

Proof: The conditions of T-Local-Config-Stack are established in two lemmas. The cmap-ok condition is satisfied by lemma 5.3.2. For the frame-stack typing and vars-ok, the combined consideration of the frame-stack typing and as well as the tail of the local configuration stack are satisfied by lemma 5.3.3.

Lemma 5.3.2 Container Map OK

If $\Gamma; U \vdash C \text{ config-ok}, hC; \Gamma; UI \vdash C^0, \Gamma; U^0, C = @L \Gamma A, C^0 = @L^0 \Gamma^0 A$

then $\Gamma; H^0 \vdash \text{head}(\overline{L^0}) \text{ cmap-ok } L^0$

Proof: By lemma 5.5.5, the cmap-ok condition will continue to be satisfied for existing pairs of local configurations in the stack. To cover the remaining cases, we now consider the three operations that add new scopes to the stack.

Call-Function

With $L_{\text{callee}} = \text{funcLocalConfig}()$, we have $L^0 \Gamma^0 = L_{\text{callee}} L \Gamma$

We can see that $\Gamma; H \vdash L \text{ cmap-ok } L_{\text{callee}}$ by first observing that lmap-ok is satisfied by the typing of call() and that $\text{dom}(L) = \text{dom}(M)$ by construction.

To establish the container equality condition, we can see that the mapping $M = (\text{cmap})(L_{\text{fn}}(\text{id}))$ is explicitly set by the labelMappings() function to satisfy this condition of cmap-ok. By definition, $\exists M : \text{dom}(M) : M(\text{label}) = \text{labelToCont}(\Gamma; H; L; \text{mapLabel}(\text{label}))$ which is the right-hand side of the container equality condition in rule Map-OK. On the left-hand side, we have $\text{labelToCont}(\Gamma; H; L_{\text{callee}}; \text{label})$. Examining the implementation of labelToCont(), we can see that it returns $M(\text{label})$ for mapped labels, which as we've seen was defined as $\text{labelToCont}(\Gamma; H; L; \text{mapLabel}(\text{label}))$, thus showing that the two sides of the container equality condition are equivalent. Therefore, Call-Function satisfies all of the conditions of cmap-ok.

Initialize and Copy-Tuple-2

The reasoning for these two cases is identical; initialize is presented.

With $L_{\text{tuple}} = \text{tupleLocalConfig}()$, we have $L^0 \Gamma^0 = L_{\text{tuple}} L \Gamma$

We can see that $\Gamma; H \vdash L \text{ cmap-ok } L_{\text{tuple}}$ by observing that lmap-ok is satisfied by the typing of init() and that $\text{dom}(L) = \text{dom}(M)$ by construction.

The logic for satisfying the container equality condition is the same as it was for the Call-Function case. Briefly, the mapping $M = (\text{cmap})(L_{\text{tuple}}(\text{id}))$ is explicitly set by the tupleMappings() function such that label-matching condition of cmap-ok is satisfied. Therefore, rules Initialize and Copy-Tuple-2 satisfy all the conditions of cmap-ok.

Lemma 5.3.3 Local Configuration Frame-Stack OK

If $\vdash; U \setminus C \text{ config-ok}, hC; \vdash; U \vdash h C^0, \alpha; U \alpha, C = @L \begin{matrix} H \\ CF \end{matrix} \Gamma A, C^0 = @L^0 \begin{matrix} H^0 \\ CF^0 \end{matrix} \Gamma^0 A$

then $\vdash; \alpha; H^0 \setminus L \begin{matrix} \text{fstack} \\ \text{CF}^0 \end{matrix} L^0_{(\text{fstack})} \Gamma : (\text{Void} \vdash; U \vdash U^0)$ and
 $\vdash; \alpha; H^0 \setminus L^0 \Gamma : (\vdash; T \vdash; U \alpha; U \alpha ?)$

Proof: We prove this lemma by parts and we'll group rules of the operational semantics by how they modify the frame-stack.

First, lemma 5.3.4 proves our claim for operations that add a new local configuration to the stack. This includes function calls, which is the most complex case and directly demonstrates the management of container labels and containers across scopes.

Similarly, lemma 5.3.5 proves our claim for operations that remove local configurations from the stack. It shows that containers are managed properly for returned values.

We address a simpler group of rules in lemma 5.3.6, where operations make changes to the current frame-stack, but do not create or remove scopes from the local configuration stack.

Finally, the simplest cases are proven by lemma 5.3.7. Here, no changes are made to the local configuration stack.

The combination of the four lemmas above proves the entire claim.

Lemma 5.3.4 Local Configuration Frame-Stack OK (open scope)

If $\vdash; U \setminus C \text{ config-ok}, hC; \vdash; U \vdash h C^0, \alpha; U \alpha, C = @L \begin{matrix} H \\ CF \end{matrix} \Gamma A, C^0 = @L^0 \begin{matrix} H^0 \\ CF^0 \end{matrix} \Gamma^0 A,$

$\vdash; H \setminus L \Gamma : (\vdash; T \vdash; U \alpha; U \alpha ?)$
 then $\vdash; \alpha; H^0 \setminus L \begin{matrix} \text{fstack} \\ \text{CF}^0 \end{matrix} L^0_{(\text{fstack})} \Gamma : (\text{Void} \vdash; U \vdash U^0)$ and
 $\vdash; \alpha; H^0 \setminus L \Gamma : (\vdash; T \vdash; U \alpha; U \alpha ?)$

Proof: The type of $L \Gamma$ is preserved in C^0 by lemma 5.5.7.

To pass values across scopes, our reasoning must also take into account that the determination of physical types from logical types depends on scope. Each local configuration has a mapping to map container labels to equivalent labels in a parent scope as well as an associated mapping M which is used by `ptype()` to determine the physical type. When passing a parameter of type $\bar{\alpha}_1$ in L_{caller} to a child scope as type $\bar{\alpha}_2$ in L_{callee} , we must show that $\text{ptype}(\vdash; H; L_{\text{caller}}; \bar{\alpha}_1) = \text{ptype}(\vdash; H; L_{\text{callee}}; \bar{\alpha}_2)$.

New scopes do not contain any variables, so they trivially satisfy `vars-ok`.

We prove the type of `L0[fstack 7! CF0 L0(fstack)]` by cases.

Call-Function

We must determine the type of `L0[fstack 7! CF0 L0(fstack)]` where

`CF0 L0(fstack) = f init-symbols (); sg`

For the new scope, we require `cmap-ok L0`, and lemma 5.3.2 establishes this judgment. By the typing of `T-Call-Function` and `V-Function`, we know that `frameinit-symbols()` types as `Void`, since the input parameters will match by lemma 5.4.4. Further, from the call typing, we know the function body frame in `Lcallee` will evaluate to `Tret`. We have

```
;; ; H; L ` call-function (::): exportType ( ; Tret)
= L0(lmap)
T = ptype ( ; H; L; exportType ( ; Tret))
L0 = funcLocalConfig(...)
Tret = ptype ( ; H; L0; Tret)
```

To establish `T = Tret`, we apply lemma 5.4.2. Therefore, the new frame-stack has the required type of `; 0; H0 ` L0[fstack 7! CF0 L0(fstack)]: (Void T; ?! ?)` with `U` and `U0` empty.

Initialize and Copy-Tuple-2

The proofs of these two cases follow the same reasoning; `initialize` is presented.

We must determine the type of `L0[fstack 7! CF0 L0(fstack)]` where

`CF0 L0(fstack) = f init-symbols ()g`

For the new scope, we require `cmap-ok L0`, and lemma 5.3.2 establishes this judgment. By the typing of `T-init`, we know that `frameinit-fields()` types as `Void`, since the input parameters will match by lemma 5.4.4. We have

```
;; ; H; L ` init (::):
Tnew = href ; container-of-tuple ; ; mutable; movable
= L0(lmap)
T = ptype ( ; H; L; exportType ( ; Tnew))
L0 = tupleLocalConfig(...)
Tret = ptype ( ; H; L0; Tnew)
```

To establish `T = Tnew`, we apply lemma 5.4.2. Therefore, the new frame-stack has the required type of `; 0; H0 ` L0[fstack 7! CF0 L0(fstack)]: (Void T; U! U [f ` newg)`.

Lemma 5.3.5 Local Configuration Frame-Stack OK (close scope)

If $\vdash \vdash U \text{ ` } C \text{ config-ok } , hC; \vdash U \text{ h } C^0, \text{ ; } U^0, C = @_{L_{pop}} L \text{ [} A, C^0 = @_{L} \text{ [} A$
 $CF \quad CF^0$

then $\vdash \text{ ; } \text{ ; } H^0 \text{ ` } L[\text{fstack } ?! CF^0 L_{(fstack)}] : (\text{Void } T; U \text{ } U^0)$ and
 $\vdash \text{ ; } \text{ ; } H^0 \text{ ` } L : (\text{! } T^0; U^0 \text{ } ?)$

Proof: The type of L is preserved in C^0 by lemma 5.5.7.

To return a value to the calling scope our reasoning must also take into account that the determination of physical types from logical types depends on scope. Each local configuration has a mapping from container labels to equivalent labels in a parent scope as well as an associated mapping M which is used by `ptype()` to determine the physical type. When returning a construct of type T_1 in L_{callee} to the parent scope as type T_2 in L_{caller} , we must show that $ptype(\text{ ; } H; L_{callee}; T_1) = ptype(\text{ ; } H; L_{caller}; T_2)$.

Return-Value and Return-Reference

We must determine the type of $L[\text{fstack } ?! CF^0 L_{(fstack)}]$ where
 $CF^0 L_{(fstack)} = h_{ret}; \text{exportType}(\text{ ; } T_{ret})i \text{ } L$ and $= L_{pop}(lmap)$.

From T-Local-Config-Stack we know that $L_{pop} \text{ cmap-ok } L$ and the returned value will match the expected input tL , preserving the type of the new frame-stack. We have

$\vdash \text{ ; } \text{ ; } H; L_{pop} \text{ ` } \text{return} (:::): T_{ret}$
 $T_{ret} = ptype(\text{ ; } H; L_{pop}; T_{ret})$
 $T = ptype(\text{ ; } H; L; \text{exportType}(\text{ ; } T_{ret}))$

To establish $T = T_{ret}$, we apply lemma 5.4.2. Therefore, the new frame-stack has the required type of $\vdash \text{ ; } \text{ ; } H^0 \text{ ` } L[\text{fstack } ?! CF^0 L_{(fstack)}] : (\text{Void } T; ?! \text{ } ?)$.

Copy-Complete and Init-Fields-Complete

We must determine the type of $L[\text{fstack } ?! CF^0 L_{(fstack)}]$ where
 $CF^0 L_{(fstack)} = h_{tuple}; \text{exportType}(\text{ ; } T_{tuple})i \text{ } L$ and $= L_{pop}(lmap)$.

From T-Local-Config-Stack, we know that $L_{pop} \text{ cmap-ok } L$ and the returned value will match the expected input tL , preserving the type of the new frame-stack. We have

$\vdash \text{ ; } \text{ ; } H; L_{pop} \text{ ` } \text{copy-fields} (:::): T_{tuple}$
 $T_{tuple} = ptype(\text{ ; } H; L_{pop}; T_{tuple})$
 $T = ptype(\text{ ; } H; L; \text{exportType}(\text{ ; } T_{tuple}))$

To establish $T = T_{tuple}$, we apply lemma 5.4.2. Therefore, the new frame-stack has the required type of $\vdash \text{ ; } \text{ ; } H^0 \text{ ` } L[\text{fstack } ?! CF^0 L_{(fstack)}] : (\text{Void } T; U [T_{tuple}! \text{ } U)$.

Lemma 5.3.6 Local Configuration Frame-Stack OK (same scope)

If $\vdash; U \vdash C \text{ config-ok}, hC; \vdash; U \vdash h C^0, \vdash; U \vdash, C = @L \overline{L} A, C^0 = @L^0 \overline{L} A;$

$L_{(id)} = L_{(id)}^0$
 then $\vdash; \vdash; H^0 \vdash L[fstack \ \ ?; CF^0 \ L_{(fstack)}^0] : (\text{Void} \ \ T; U \ \ U^0)$ and
 $\vdash; \vdash; H^0 \vdash L : (\text{Tail} \ \ T^0; U \ \ U^0 \ \ ?)$

Proof: The set of rules we are considering here modify the current local frame-stack, but do not create new scopes or modify parent scopes. By lemma 5.5.7, we know that the type of \overline{L} is preserved, as well as the type of each individual unmodified frame by lemma 5.5.9. We must show the type of $CF \ L_{(fstack)}$ is the same as the type of $CF^0 \ L_{(fstack)}^0$ so that T-Local-Config-Stack retains the same typing.

When reasoning about the type of the frame-stack, we break the stack into a modified portion and a stable tail portion as follows $CF \ L = \overline{F}_{old} \ \ \overline{F}_{tail}$ and $CF^0 \ L^0 = \overline{F}_{new} \ \ \overline{F}_{tail}$. The argument in each of the cases below is that the type of \overline{F}_{old} is the same as \overline{F}_{new} . Then, by lemma 5.5.8, \overline{F}_{tail} preserves its type. With these two results combined, we have completed the typing of $CF^0 \ L_{(fstack)}^0$ in C^0 and show it to be the same as in C .

Where the typing of the new frames is non-obvious, there will be further justification based on the conditions of the relevant type rule.

Next-Frame

$\overline{F}_{old} = f \text{ void } g : (\text{Void} \ \ \text{Void}; U \ \ U)$
 $\overline{F}_{new} = fg : (\text{Void} \ \ \text{Void}; U \ \ U)$

Consume-Result

$\overline{F}_{old} = f \text{ result}; OFg : (\text{Void} \ \ T; U \ \ U)$
 $\overline{F}_{new} = f OF[\ ?; result]g : (\text{Void} \ \ T; U \ \ U)$

The resulting closed frame will be well typed because its components inherited from the open frame were well typed and the substituted result has the same typing as

Pop-Local

$\overline{F}_{old} = f \text{ pop-local } [\ v(v)g : (\text{Void} \ \ \text{Void}; ?! \ \ ?)$
 $\overline{F}_{new} = f \text{ void } (v)g : (\text{Void} \ \ \text{Void}; ?! \ \ ?)$

pop-local $()$ will remove v from $L_{(vars)}$ and we know the corresponding change is also made to \overline{L} from the previous frame-stack typing. The remainder of the conditions $vars-ok$ are as they were in C , thus $vars-ok$ is preserved in C^0 .

Let

$\overline{F}_{old} = f \text{ let } (v)g : (\text{Void} \ \ \text{Void}; ?! \ \ ?)$

$\overline{F}_{new} = f \text{ init-symbol } (v); s \uparrow^v; \text{ pop-local } \uparrow^v(v)g : (\text{Void} \quad \text{Void}; ?! \quad ?)$

The new frame's components are all extracted from well typed components from the original let frame. Frame `init-symbol ()` will add v to $L_{(vars)}$ and `pop-local (v)` will remove it, leaving the `vars-ok` conditions as they were in C . Thus `vars-ok` is preserved in C^0 .

Assign-Value

$\overline{F}_{old} = f \text{ assign-value } ()g : (\text{Void} \quad \text{Void}; U \quad U)$

$\overline{F}_{new} = f \text{ copy-tuple } (); \text{ assign-copied } ()g : (\text{Void} \quad \text{Void}; U \quad U)$

The new frame's components are all extracted from well typed components from the original assign-value frame. `assign-copied ()` consumes the result from `copy-tuple ()` resulting in a frame of type `Void`.

Init-Symbols-Value

$\overline{F}_{old} = f \text{ init-symbols } ()g : (\text{Void} \quad \text{Void}; ?! \quad ?)$

$\overline{F}_{new} = f \text{ copy-tuple } (); \text{ init-symbol } (); \text{ init-symbols } ()g : (\text{Void} \quad \text{Void}; ?! \quad ?)$

The new frame's components are all extracted from well typed components from the original init-symbols frame. `init-symbol ()` consumes the result from `copy-tuple ()`, resulting in a frame of type `Void`. The next `init-symbols ()` is also typed as `Void`, so the new stack retains its type.

Init-Symbols-Ref

$\overline{F}_{old} = f \text{ init-symbols } ()g : (\text{Void} \quad \text{Void}; ?! \quad ?)$

$\overline{F}_{new} = f \text{ init-symbol } (); \text{ init-symbols } ()g : (\text{Void} \quad \text{Void}; ?! \quad ?)$

The new frame's components are all extracted from well typed components from the original init-symbols frame.

Init-Symbol-Value and Init-Symbol-Ref

$\overline{F}_{old} = f \text{ init-symbol } (v)g : (\text{Void} \quad \text{Void}; ?! \quad ?)$

$\overline{F}_{new} = f \text{ void } \uparrow^v g : (\text{Void} \quad \text{Void}; ?! \quad ?)$

`init-symbol ()` will add v to $L_{(vars)}$ in agreement with the frame typing of `init-symbol`.

The remainder of the `vars-ok` conditions are as they were in C , thus `vars-ok` is preserved in C^0 .

Init-Fields-Value

$\overline{F}_{old} = f \text{ init-fields } ()g : (\text{Void} \quad \text{Void}; U \quad U)$

$\overline{F}_{new} = f \text{ copy-tuple } (); \text{ init-field } (); \text{ init-fields } ()g : (\text{Void} \quad \text{Void}; U \quad U)$

The new frame's components are all extracted from well typed components from the original init-fields frame. `init-field ()` consumes the result from `copy-tuple ()` resulting in a frame of type `Void`, and the next `init-fields ()` is also typed as `Void`, so the new stack retains its type.

Init-Fields-Ref

$\overline{F}_{old} = f \text{ init-fields } ()g : (\text{Void } \text{Void}; U! \ U)$

$\overline{F}_{new} = f \text{ init-field } (); \text{init-fields } ()g : (\text{Void } \text{Void}; U! \ U)$

The new frame's components are all extracted from well typed components from the original `init-fields` frame.

Copy-Tuple

$\overline{F}_{old} = f \text{ copy-tuple } ()g : (\text{Void } T; ?! \ ?)$

$\overline{F}_{new} = f \text{ copy-tuple2 } (); \text{discard-copy-env } ()g : (\text{Void } T; ?! \ ?)$

The frame typing conditions of `copy-tuple2` are satisfied by the typing of `copy-tuple`. After simplifying the result of `copy-tuple2` using `discard-copy-env`, the typing matches the original and the new stack retains its type.

Copy-Value

$\overline{F}_{old} = f \text{ copy-fields } ()g : (\text{Void } \text{Void}; U! \ U)$

$\overline{F}_{new} = f \text{ copy-tuple2 } (); \text{copy-init } (); \text{copy-fields } ()g : (\text{Void } \text{Void}; U! \ U)$

The new frame's components are all extracted from well typed components from the original `copy-fields` frame. `copy-init` consumes the result from `copy-tuple`, resulting in a frame of type `Void`, and the next `copy-fields` is also typed as `Void`. `copy-tuple2` adds location `new` to `U` and `copy-fields` removes it, therefore the new stack retains its type.

Copy-Ref

$\overline{F}_{old} = f \text{ copy-fields } ()g : (\text{Void } \text{Void}; U! \ U)$

$\overline{F}_{new} = f \text{ copy-init } (); \text{copy-fields } ()g : (\text{Void } \text{Void}; U! \ U)$

The new frame's components are all extracted from well typed components from the original `init-fields` frame.

Lemma 5.3.7 Local Configuration Frame-Stack OK (simple replacement) $\frac{H \quad H^0 \quad 1}{H^0 \quad 1}$

If $\vdash; U \setminus C \text{ config-ok}, hC; \vdash U! h \quad C^0, {}^0U^0, C = @L \quad \square A, C^0 = @L \quad \square A$
 $CF \quad CF^0$

then $\vdash; {}^0H^0 \setminus L[fstack \ 7! \ CF^0 \ L_{(fstack)}^0] : (\text{Void } T; U! \ U^0)$ and
 $\vdash; {}^0H^0 \setminus L^0 : (T! \ T^0; U! \ ?)$

Proof: Here we consider rules that only modify the current frame, leaving the the frame-stack unmodified. By lemma 5.5.7, we know the typing of the `CF` is unaltered in `C0`. We will show that the new frame has the same type as the old frame `CF`: `T` and `CF0`: `T`. This equality ensures that `T-Closed-Stack` will return the same stack typing after by applying lemma 5.5.8. Therefore; $\vdash; {}^0H^0 \setminus L[fstack \ 7! \ CF^0 \ L_{(fstack)}^0] : (\text{Void } T; U! \ U^0)$ will have the same type as before.

Rule	CF Type	CF ⁰ Type
Copy-Discard-Env	T	$h; T_i : T$ by lemma 5.2.3
Assign-Copied	Void	void : Void
Assign-Ref	Void	void : Void
Field-Access	T_{field}	$h_{\text{field}}^0; T_{\text{field}} i : T_{\text{field}}$ by lemma 5.2.3
Variable-Access	T	$h; T_i : T$ by lemma 5.2.3
Init-Field-Value	Void	void : Void
Init-Field-Ref	Void	void : Void
copy-init-Value	CopyEnv	A : CopyEnv by old frame typing
copy-init-Ref	CopyEnv	A : CopyEnv by old frame typing

5.4 Logical to Physical Consistency

Lemma 5.4.1 Logical to Physical Type Match

If $;; U \text{ ` } @L \sqsubset A$ config-ok
 CF

and two frames $;;; H; L \text{ ` } F_{\text{lhs}} : T_{\text{lhs}}$ and $;;; H; L \text{ ` } F_{\text{rhs}} : T_{\text{rhs}}$ such that $T_{\text{lhs}} \text{ b } T_{\text{rhs}}$,

then there will also be a match between physical types with
 $\text{ptype}(; H; L; T_{\text{lhs}}) \text{ p-match } \text{ptype}(; H; L; T_{\text{rhs}})$

Proof: The Ref-Initializable rule which defines $T_{\text{lhs}} \text{ b } T_{\text{rhs}}$ provides enough information to meet the three conditions of Physical-Match. First we consider both the class and mutability sub-components of the type, both of which ptype directly maps into the physical type. The requirements of Ref-Initializable are the same as Physical-Match, so these two conditions will be satisfied. For the container condition, we look at two sub-cases. If the left-hand-side container label is unknown, then $\text{ptype}()$ will map it to unknown-cont, which will satisfy Physical-Match. Otherwise, the container labels must be equal and map to the same physical container to satisfy the condition.

Lemma 5.4.2 Returned Values Match

If $;; U \text{ ` } C$ config-ok, $hC; ; U \text{ h } C^0; ; U^0, C = @L \text{ L}_{\text{caller}} \sqsubset A$
 CF

and a type T in scope L then

$\text{ptype}(; H; L; T) \text{ p-match } \text{ptype}(; H; L_{\text{caller}}; \text{exportType}(L_{\text{imap}}; T))$

Proof: `ptype()` and `exportType()` both directly map the class, mutability and mobility components directly into the physical type, guaranteeing that they will match. For the container label, `exportType()` maps the label using `mapLabel(; T(label))`, which will map correctly by lemma 5.4.3, and all components of the physical types will match.

Lemma 5.4.3 Returned Containers Match

If $\vdash ; U \text{ ` } C \text{ config-ok } , hC ; ; UI h \ C^0 ; U^0 , C = @L \ L_{\text{caller}} \ \square A ,$
 CF

$= L_{(lmap)}$ and a heap location h ; $\vdash ; H ; L \ \text{ ` } : T$ with container label $= T_{(label)}$ then `labelToCont(; H ; Lcaller ; mapLabel(;))` c-match `labelToCont(; H ; L ;)`

Proof: If $\neq \text{dom}(h)$, then we can immediately see that the `Map-OK` condition of `T-Local-Config-Stack` gives us the result we need. If, on the other hand, h is a path, then the root container h_{root} of h must exist in U by type rule `lmap-ok`. `CMap-OK` then guarantees that

`labelToCont(; H ; L ; hroot)` c-match `labelToCont(; H ; Lcaller ; mapLabel(; hroot)`

With the base of the path correctly mapped, following the path is a deterministic sequence of immutable field accesses which will lead to the same final container and satisfy c-match.

Lemma 5.4.4 Passed Parameters Match

If $\vdash ; U \text{ ` } C \text{ config-ok } , hC ; ; UI h \ C^0 ; U^0 , C = @L \ L_{\text{caller}} \ \square A ,$
 CF

$= L_{(lmap)}$, and type T_{passed} in L_{caller} and T_{parm} in L such that `exportType(; Tparm) e Tpassed` then `ptype(; H ; L ; importType(Tpassed ; Tparm))` p-match `ptype(; H ; Lcaller ; Tpassed)`

Proof: First, focusing on the container component of the physical type, note that the `importType()` function simply swaps the container label of T_{passed} and sets it to the label from T_{parm} . The `ptype()` function maps labels to containers using the `labelToCont()` function, and for p-match to hold, c-match must be satisfied. With $h_{\text{passed}} = T_{\text{passed}(label)}$ and $h_{\text{parm}} = T_{\text{parm}(label)}$, we must show that `labelToCont(; H ; L ; hparm)` c-match `labelToCont(; H ; Lcaller ; hpassed)`.

From the definition of `e` we know that `mapLabel(; hparm)` l-match h_{passed} , then with l-match established, lemma 5.4.5 provides the c-match judgment we need.

For the remaining components of the physical types, the `ptype` function simply maps them unchanged, and by the `e` relation, we know those components will be satisfied in p-match.

Lemma 5.4.5 Passed Containers Match

$$\text{If } \vdash ; U \text{ ` } C \text{ config-ok, } hC; \vdash UI h \text{ } C^0, \text{ } U^0, C = @L \text{ } L_{\text{caller}} \text{ } \Gamma A, \text{ } CF$$

$$= L_{(\text{lmap})}, \text{ and container label}_{\text{passed}} \text{ in } L_{\text{caller}} \text{ and } \text{parm} \text{ in } L \text{ such that}$$

$$\text{mapLabel}(\text{ ; parm}) \text{ l-match}_{\text{passed}}$$

$$\text{then labelToCont}(\text{ ; H; L; parm}) \text{ c-match labelToCont}(\text{ ; H; L}_{\text{caller}}; \text{passed})$$

Proof: From the definition of l-match, we know that either $\text{passed} = \text{mapLabel}(\text{ ; parm})$ or that $\text{parm} = \text{unknown}$. We'll first consider the equality case. By lemma 5.4.3, we know that the inverse mapping holds

$$\text{labelToCont}(\text{ ; H; L}_{\text{caller}}; \text{mapLabel}(\text{ ; parm})) \text{ c-match labelToCont}(\text{ ; H; L; parm})$$

By our equality, we can substitute in passed , yielding

$$\text{labelToCont}(\text{ ; H; L}_{\text{caller}}; \text{passed}) \text{ c-match labelToCont}(\text{ ; H; L; parm})$$
 which is backwards from what we need. However, we have already assumed that $\text{parm} \neq \text{unknown}$, which means that c-match degrades to a simple equality check, which is symmetric. Therefore the terms can be swapped, giving the desired result.

In the case where $\text{parm} = \text{unknown}$, $\text{labelToCont}(\text{ ; H; L; parm})$ will evaluate to unknown-cont which satisfies c-match.

5.5 Preservation Lemmas

This section contains a series of lemmas showing that the rules of the dependent type system are such that the typing only depends on immutable objects. This leads to the useful property that a typing that is valid in configuration C will continue to be valid in a subsequent configuration C^0 .

Lemma 5.5.1 Stabletype() Evaluation

$$\text{If } \vdash ; U \text{ ` } C \text{ config-ok, } hC; \vdash UI h \text{ } C^0, \text{ } U^0, C = @L \text{ } \Gamma A, \text{ } CF$$

$$C^0 = @L^0 \text{ } \Gamma^0 A, \text{ } CF^0$$

$$\vdash ; H; L \text{ ` } F : T, T = \text{ptype}(\text{ ; H; L; T}), L_{(\text{id})} = L_{(\text{id})}^0$$

$$\text{then } T = \text{ptype}(\text{ }^0; H^0; L^0; T).$$

Proof: We examine the implementation of $\text{ptype}()$ and note that the class, mutability and mobility components pass directly from T to T unaffected by the heap or configuration

typing. The final component of the physical type is the container, which `ptype()` computes using the function `labelToCont()`. By lemma 5.5.2, we know that $\text{labelToCont}(\text{;H;L;T}_{(\text{label})}) = \text{labelToCont}(\text{ }^0\text{;H}^0\text{;L}^0\text{;T}_{(\text{label})})$. Therefore, all components of the physical types are identical and `ptype()` is stable.

Lemma 5.5.2 Stable `labelToCont()` Evaluation

If $\text{; ; U} \setminus \text{C config-ok}, \text{hC}; \text{ ; UI h C}^0, \text{ }^0\text{;U}^0, \text{C} = \text{@}_{\text{CF}} \text{L} \text{ }^1\text{A}, \text{C}^0 = \text{@}_{\text{CF}^0} \text{L}^0 \text{ }^1\text{A},$
 $\text{; ; ; H;L} \setminus \text{F : T}, = \text{labelToCont}(\text{ ;H;L;T}_{(\text{label})}), \text{L}_{(\text{id})} = \text{L}_{(\text{id})}^0$
then $= \text{labelToCont}(\text{ }^0\text{;H}^0\text{;L}^0\text{;T}_{(\text{label})})$.

Proof: For $\text{ } = \text{T}_{(\text{label})}$, we consider the cases in the implementation of the `labelToCont()` function, first addressing the simple cases. The cases `unknown-label` and `null-label` are simple constants with no dependencies on any environment. The default case returns $(\text{cmap})(\text{L}_{(\text{id})})(\text{ })$, which will be consistent because the container map is immutable.

Finally, when is a path, the heap will be accessed. Since `type` is a valid type, the path of has satisfied `path-ok`. By rules `V-SP-Tuple` or `V-SP-Var`, the base of the path must be fixed. Further, by the rule `V-SP-Step`, any fields accessed along the path will also be fixed, and by lemma 5.2.2, we know that all of these locations will have the same value in H and H^0 . Therefore, all heap information that the path depends on is the same in H and H^0 , which forces the results to match.

Therefore, for all possible values for , `labelToCont()` is stable.

Lemma 5.5.3 f-match Preserved

If $\text{; ; U} \setminus \text{C config-ok}, \text{hC}; \text{ ; UI! h C}^0, \text{ }^0\text{;U}^0, \text{H} = \text{C}_{(\text{heap})}, \text{H}^0 = \text{C}_{(\text{heap})}^0; \setminus \text{ }^2 \text{H};$
 $\text{ ; H}; \setminus \setminus \text{T f-match T}$
then $\text{ }^0\text{;H}^0; \setminus \setminus \text{T f-match T}$

Proof: We know from `t-match` in C that $\text{T} = \text{ptype}(\text{ ;H;L;T})$ and by lemma 5.5.1 we know that $\text{T} = \text{ptype}(\text{ }^0\text{;H}^0\text{;L}^0\text{;T})$. Therefore the `t-match` condition of `f-match` continues to hold.

The second condition of interest is the `within` judgment, which is preserved by lemma 5.5.4. If all dependencies of `offmatch` are preserved, we conclude that `fmatch` is also preserved.

Lemma 5.5.4 Containment within Preserved

If $\text{; ; U} \setminus \text{C config-ok}; \text{hC}; \text{ ; UI! h C}^0, \text{ }^0\text{;U}^0, \text{H} = \text{C}_{(\text{heap})}, \text{H}^0 = \text{C}_{(\text{heap})}^0, \setminus$

Γ_1 within Γ_2
then $\Gamma_0 \vdash \Gamma_1$ within Γ_2

Proof: The within judgment uses Γ to recursively look up parent containers. The type environment only grows and is never modified. Therefore $\Gamma_0(\Gamma_1) = \Gamma(\Gamma_1)$, making the entire judgment preserved in \mathbb{C}^0 .

Lemma 5.5.5 CMap-OK Preserved

If $\Gamma; \Gamma; U \vdash C$ config-ok ; $hC; \Gamma; UI h \mathbb{C}^0, \Gamma; U^0; \Gamma; H \vdash L_{\text{caller}} \text{ cmap-ok } L_{\text{callee}},$
 $L_{\text{caller}}^0 \vdash \mathbb{C}^0, L_{\text{caller}}^0(\text{id}) = L_{\text{caller}}(\text{id}), L_{\text{callee}}^0 \vdash \mathbb{C}^0, L_{\text{callee}}^0(\text{id}) = L_{\text{callee}}(\text{id})$
then $\Gamma; H^0 \vdash L_{\text{caller}}^0 \text{ cmap-ok } L_{\text{callee}}^0$

Proof: Looking at each of the conditions in rule CMap-OK, we'll start with lmap-ok, which is a pure typing rule and independent from the configuration. The lookup of L from the caller scope is equivalent in \mathbb{C}^0 , because the frame-stack of the caller is fixed while the callee executes. Therefore, the lmap-ok condition continues to hold in \mathbb{C}^0 .

Both Γ and M are immutable, which leaves the final container equality to consider. Here, we appeal to lemma 5.5.2 to establish that the left and right sides of this equality will evaluate to the same values in \mathbb{C}^0 and \mathbb{C} for each $i \in \text{dom}(M)$. Since the equality held in \mathbb{C} , it will continue to hold in \mathbb{C}^0 .

Lemma 5.5.6 vars-ok Preserved

If $\Gamma; \Gamma; U \vdash C$ config-ok ; $hC; \Gamma; UI h \mathbb{C}^0, \Gamma; U^0; L \vdash C; L^0 \vdash \mathbb{C}^0, L_{(\text{vars})} = L^0_{(\text{vars})}$
then $\Gamma; \Gamma; \Gamma^0 \vdash L^0_{(\text{vars})} \text{ vars-ok}$

Proof: We know that we have heap-ok in \mathbb{C} , which means that every local symbol is mapped to a box or reference which is location-ok. $T_{\text{variable}} = T_{\text{box/ref}}$ implies that $(\Gamma)_{\text{box/ref}} = \text{ptype}(\Gamma; H; L; T_{\text{box/ref}}) = \text{ptype}(\Gamma; H; L; T_{\text{variable}})$. Therefore, we conclude that the t-match judgment is satisfied for all local variables.

Lemma 5.5.7 Local Configuration Stack Typing Preserved

If $\Gamma; \Gamma; U \vdash C$ config-ok ; $hC; \Gamma; UI h \mathbb{C}^0, \Gamma; U^0; H = C_{(\text{heap})}, H^0 = C^0_{(\text{heap})}, \Gamma; \Gamma; H \vdash$
 $L \vdash (\Gamma; T^0; U^0; ?)$
then $\Gamma; \Gamma; H^0 \vdash L \vdash (\Gamma; T^0; U^0; ?)$

Proof: The cmap-ok judgment continues to hold in \mathbb{C}^0 by lemma 5.5.5. Since the set of variables on the stack has not changed, the vars-ok judgment remains true, by lemma 5.5.6. Next, by lemma 5.5.8, we have $\Gamma; H^0 \vdash L_{(\text{fstack})} : (\Gamma; T^0; U^0; U^0)$

Finally, by induction on this lemma, we have; $\overset{0}{H} \overset{0}{L} \vdash \Gamma : (T^? \ T^{00}; U^? \ ?)$. Rule T-Local-Config-Stack-Empty provides the base case and guarantees that the induction terminates. Therefore, all conditions of the rule T-Local-Config-Stack are met, and we can conclude that an unmodified stack of local configurations retains its typing id .

Lemma 5.5.8 Frame-Stack Typing Preserved

If $\overset{0}{H}; \overset{0}{L} \vdash C \text{ config-ok}; \overset{0}{h}C; \overset{0}{U} \vdash h \ C^0; \overset{0}{U}^0, H = C_{(\text{heap})}, H^0 = C_{(\text{heap})}^0,$
 $\overset{0}{H}; \overset{0}{L} \vdash \bar{F} : (T^! \ T^{00}; U^! \ U^0)$
 then $\overset{0}{H}; \overset{0}{L} \vdash \bar{F} : (T^! \ T^{00}; U^! \ U^0)$

Proof: We consider the conditions of the frame-stack typing rules as defined in section 4.9 in their entirety. There is significant similarity among the rules, and we'll reason by cases over the kinds of conditions that need to be satisfied rather than over each rule individually.

First, all conditions of the form $T = \text{ptype}()$ will continue to hold by the stability of $\text{ptype}()$ established in lemma 5.5.1. All induction on the tail of the frame-stack is satisfied by induction on this lemma. The base case is typed by rule Empty-Stack, which guarantees that the induction terminates. Finally, typing of individual frames is preserved by lemma 5.5.9. This completes the set of conditions common to most of the frame-stack typing rules.

The rule T-Immediate-Return is special, and we consider its unique condition $\bar{\pi} = \text{exportType}(L_{(\text{lmap})}; T_{\text{callee}})$. We know T_{callee} is preserved by lemma 5.5.1 and lmap is immutable. Therefore, this condition is also preserved, and we conclude that all conditions appearing in the frame-stack typing rules continue to hold id^0 .

Lemma 5.5.9 Frame Typing Preserved

If $\overset{0}{H}; \overset{0}{L} \vdash C \text{ config-ok}; \overset{0}{h}C; \overset{0}{U} \vdash h \ C^0; \overset{0}{U}^0, H = C_{(\text{heap})}, H^0 = C_{(\text{heap})}^0,$
 $\overset{0}{H}; \overset{0}{L} \vdash F : T$
 then $\overset{0}{H}; \overset{0}{L} \vdash F : T$

Proof: Here, we distinguish frames that are the original statements and expressions of the language from the extended frames that are typed in section 4.6. For the proper language components, we don't need to concern ourselves with the extended typing environment as the rules only depend on env and env^0 . Therefore we conclude that all statements and expressions have the same typing id^0 .

For the extended typing rules of section 4.6, we use the same strategy of lemma 5.5.8 and show preservation for the various kinds of conditions in these rules rather than address each rule directly. All type lookups in env or env^0 will be the same because env and env^0 are the

same and existing entries in Δ are immutable. For frames that hold other frames and inductively type those sub-frames, we inductively apply this lemma to show the sub-frames retain their same typing.

Simple logical predicates on types will remain the same, because all the type information consumed is preserved. For example, the condition $\text{is_ref}(T_i) = \text{ref} \wedge T_i \in T_{\text{init } i}$ from rule T-Init-Symbols is composed of typings that are preserved.

For rules that test copy-map-ok, we know that this condition will hold, because Δ is the same as before. Note that in each instance Δ in the rules, it is part of the frame itself, and this lemma is only concerned with typing identical frames.

Finally, the rule T-Location has a `ptype()` condition, which is preserved by lemma 5.5.1, and `at-match` condition which does not depend on the heap.

Therefore, all of the conditions present in the rules of section 4.6 will be preserved in C^0 .

Chapter 6

Related Work

6.1 Ownership Types

Aliasing causes many problems for imperative programs. Reasoning about a system with aliases is problematic. Generally, when considering a method's behavior, there will be limited or no knowledge of what external aliases can exist. This interferes with proving code correctness and forces an optimizing compiler to make pessimistic choices.

Although aliasing is not the focus of our container language, there can be reductions in aliasing with this system. We encourage pass-by-value semantics by providing well-defined copying. Once you've made a fresh copy of an object, external aliasing is eliminated. Although our semantics define a potentially expensive copy algorithm, an alternative implementation could use persistent data structures like a functional language. The cost of a copy would become zero at the expense of increased update time. If such a system were implemented, pass by value wouldn't incur a performance penalty and aliasing could be avoided in many cases.

Ownership types were first proposed by David G. Clarke et al. in [9] as a way to enforce encapsulation, which results in all aliasing concerns being local and easier to reason about. Significant work in this area has been continued by many researchers and two surveys can be found in [8] and [14].

Although our containment system doesn't involve owners, the fact that a container has no outgoing references means it is self-contained and implicitly everything in the container is owned by the container. We make no attempt to enforce encapsulation, but an extension of container-typing could add restrictions on the use of containers. If a container is declared

as encapsulated, then only privileged code would be able to hold references into that container.

In this thesis, we are focused on containment, which is converse of encapsulation in that it disallows outgoing references rather than incoming. However, despite having converse goals, the fundamental type-system machinery needed to manage containment has similarities to ownership types. With ownership types, an object can contain references to owned sub-objects. Types are augmented to indicate that certain fields represent owned objects. When an object is owned the type system must prevent the creation of external aliases. Both systems must control the assignment of references.

The paper *Sheep Cloning with Ownership Type* by Paley et al. [12] recognizes that ownership types can be leveraged for the purpose of object copying. This thesis fundamentally builds on their insight. They recognize the difficulty in manually implementing cloning code, and they present a hybrid between deep cloning and shallow cloning. Their work uses the declared owner of nested objects to determine which objects should be copied. Our approach mirrors theirs, and although we use containers instead of owners, there is fundamental similarity. This thesis, differs in that we remove the goals of ownership types so that we can exclusively focus on automated functionality related to self-contained objects. Their hybrid approach contrasts with our system as we have chosen to explicitly prevent a hybrid copy. Their work is further refined in [16]. Here, they refactor their work to facilitate proving its soundness.

Although still primarily encapsulation-focused, the work of Bettini et al. [3] introduces the concept of boxes, which is similar to our notion of a container. They recognize that making every object encapsulated is too restrictive. Their duality between boxed and un-boxed classes is similar to the data and entity classes of the container language. In contrast, the ownership type system of Boyapati et al. [5] also tries to make a less restrictive system, but takes a different approach. It uses module boundaries as encapsulation barriers. Inner classes would have full access to objects owned by the outer class. It's appealing that they could avoid introducing a new concept like a box or container, but using a module as a container would be too broad, and we would lose pass-by-value semantics.

Ownership types are a dependent type system, as is our container type system. In *Ownership Type Systems and Dependent Class* Dietl et al. [10] implement ownership types on top of a dependent type system. They conclude that specialized syntax for ownership types is still desirable, even though ownership types can be fully expressed in a more general way. However, by establishing an implementation of ownership types in a general dependent-type system, they have shown a new way to reason about ownership types. A similar treatment of our container work could yield additional insights.

The work of Cameron et al. [7] also evaluates ownership types from a more theoretical viewpoint and compares them with more fundamental dependent type systems such as Dependent ML. Again, applying this type of analysis to our container type system would be an worthwhile exercise.

Huang et al. [11] develop type inferencing of ownership types. They argue that the overhead of ownership specification hinders the adoption of ownership types. We make a similar argument with respect to containers. Any adoption is unlikely unless systems like these are simple and easy to use.

6.2 Serialization

Serialization in object-oriented systems is another broad field of research and is also widely used in practice. In our system, we claim that objects that are candidates for serialization should always be self-contained, and our container system enforces that property. This greatly simplifies the problem of serialization. Here, we'll compare our container system to a number of other systems.

In Instant Pickles by Miller et al. [13], they develop a pickle combinator with the flexibility to customize to many data formats. To enhance our system so the output format could be customized would require a reflection system. If this was in place, containment could continue to provide the nice property that no per-class code would be needed for serialization. A generic piece of code could reflect on the meta-data and using an algorithm similar to the one presented in this thesis, produce a custom serialized form of a self-contained object. This paper also considers inheritance and versioning which is absent from our work.

In the Fibonacci system [1], Albano et al. have developed a programming language for object databases. They build a system with similar ideas to an entity-relationship model. They state that any value, irrespective of its type, has the same rights to persistence. Here, we take a different stance, where non-self-contained objects can only be persisted as part of a larger self-contained object. Although this language is object oriented, they have added concepts like associations, which effectively create a data model like the database schema in a relational database.

Nestmann et al. [15] develop a system not to copy objects, but to migrate them. They go to great lengths to achieve transparency in their distributed system with surrogate objects working locally to forward calls to remote objects. Effectively, they extend their

application's heap across physical machines. This work is pursuing a goal opposite to ours in that we want to make message passing easier to reduce the need for a shared heap.

6.3 Operational Semantics

The operational semantics of chapter 4 were initially modeled using the MJ technical report [4] as inspiration, although the design drifted apart as our semantics were developed. In the frame stack typing of the MJ system, they duplicated logic from their type rules in order to re-build an equivalent typing environment needed for the stack typing. In our system, with many more frames to consider, this would have become unmanageable. Like the MJ system, we need the typing environment for our frame-stack typing as well. Our typing environment is only built by the type rules, and we bind it to our frames to avoid rebuilding it a second time. This blunt-force strategy worked well and saved significant duplication of work which would have cluttered the frame stack typing.

Chapter 7

Future Work

This chapter looks at three areas where the ideas of this thesis can be extended and presents ideas to address them.

7.1 Sub-Containers

One significant issue with the container type system is that it requires precise specification of a container. For example, if a reference has container `then` that reference is not allowed to point to an object contained with `c.nested`. This could be fixed by introducing a sub-container constraint for references. For example, figure 7.1 contains a proposed specifier which constrains a reference to be within a container and not necessarily have that precise container. This allows for nesting containers without exposing internal structure.

With an imprecise container constraint, it would then also be natural to allow narrowing a reference to a specific container. This could be done by a matching predicate which refines the type to a precise container label. E.g. `within nested then ...`

One concern with this narrowing behavior is that up until now, the soundness of the container type checking implied that there is no extra run-time overhead associated with container typing. In order to support the narrowing behavior, there would need to be an extra field in each object to indicate which container it was in. Perhaps this could be minimized with global optimization as only objects that participate in narrowing would require this extra overhead.

```

class Inner( value : Int )
constructor Inner( value : Int ) {
    self.value = value;
}

class Outer( fixed nested : Inner )
constructor Outer( value : Int ) {
    self.nested = Inner( value );
}

class Main ()
constructor Main() {
    fixed var multiLevel = Outer(1);

    iref r1 : Int :: multiVar =^ null;
    r1 =^ multiLevel.inner.value; // COMPILER ERROR! container mismatch

    iref r2 : Int :: multiVar.nested =^ null;
    r2 =^ multiLevel.inner.value; // OK, containers match

    iref r3 : Int ::> multiVar =^ null;
    r3 =^ multiLevel.inner.value; // OK, with new ::> sub-container!
}

```

Figure 7.1: Sub-Container Notation

7.2 Roles

In chapter 2 we presented an algorithm to do a deep comparison between two self-contained tuples (figure 2.11). Conceptually, this algorithm must assign an identity to each of the nested objects. A mapping is created indicating a correspondence between nested tuples, saying tuple x on the left hand side matches with tuple y on the right hand side. A natural extension of this is to say that an entity has an identity relative to its container. This is in contrast to the typical instance-identity notion, where the event of an object's creation defines its identity. With a container-relative identity, it becomes possible to have two container objects and relate nested entities that fulfill the same role with respect to their container.

The identities from the deep comparison algorithm are simply the first path to reach a nested object by a depth-first search. If a concept of a role were added to the language, then this mechanically-derived identity could be replaced with something more meaningful. The idea would be that all entities would be assigned a role. The indirect reference `iref` shown in the demonstration language was added with the idea that it would be a reference to a role. In the most direct case, an entity could be described as a having the role granted

by a special *role* field of a container object. This idea can be extended by longer paths to reach nested entities. Every entity's identity could be described in the form *(role-of-granting-object).role-name*. This would be a meaningful description rather than a memory address. Identities would be preserved over a serialization and de-serialization round trip.

Descriptive identities can also be represented directly in the language as data. A non-self-contained entity class could be converted into a pure data class by replacing all references with materialized identity paths. Conversion functions could be automatically generated. The inverse operation (data to entity) is simply a function that takes a container object and an identity path of a nested entity and returns a reference.

With descriptive identities materialized as data, entities can be translated into pure data objects and can then be serialized. Additional functionality can be built on this as well. We've looked at boolean equality comparison, but you could also compare two objects and return the difference between them. With descriptive identity paths, you could return a list of new, deleted and modified entities. It's commonplace to *diff* two text files to see what has changed. Version control systems store a sequence of modifications to text and can re-create any version of a file by applying the recorded changes. This capability would also be useful in a distributed system. Replicas of data could be synchronized by sending only changes to the data. Many such systems like this exist, but not directly supported by a programming language.

If features like those in a version control system could be brought into an object-oriented language, then implementing an undo button would be a trivial task. You would be able to roll back the state of the program after an exception is caught.

Many details have been glossed over in this description, however with additional effort, the foundation of containers can be extended to become a richer data model with enhanced capabilities.

Although this work was not completed in this thesis, research related to roles was reviewed and we'll compare two of those works now.

In the work of Steimann [18], the connection is made between an object assigned to a variable and a role. Variables have purposes and when a variable refers to an object, it gives that object a role. The paper then argues that variables should be typed with interfaces specific to the role needed by the variable. Further, the design of interfaces should be role-centric, essentially making interfaces and roles the same thing. Although this paper is more focused on object oriented modeling, parallels can be drawn with our proposal. The recognition that the roles of an object can change throughout its lifetime is an important one. If we develop a system that links object identity to roles, then we


```

class Demo( v1 : Int, v2 : Int )
constructor Demo() {
  v1 = 1;
  v2 = 2;
}

method Main.test( i ref r1 : Int :: 'c, i ref r2 : Int :: 'c ) {
  ...
}

class Main ()
constructor Main() {
  var demo = Demo();
  self.test( demo.v1, demo.v2 ); // COMPILER ERROR! demo is not fixed!
  // The paths demo.v1 and demo.v2 clearly have the same container, but because
  // var demo is not declared as fixed, the container labels are 'unknown'
}

```

Figure 7.2: Data Flow Example

need to fully consider how to handle object identity when roles change. This will be a challenging problem.

An alternative approach to identity is proposed in [19] by Vaziri et al.. Here they introduce an explicit concept called a relation-type, where a programmer designates a set of immutable fields to act as that object’s identity. The class of an object is also considered part of its identity, so the fields only need to be unique with respect to their class. This model is similar to a relational database and makes identity a concrete part of every object. This system will inherit the difficulties databases have with unique key generation, such as the efficient allocation of unique keys across a distributed system. However, explicit identity nicely bypasses the problem of changing roles. Their relation-type model would work well to satisfy our goal of easy serialization.

7.3 Data Flow Analysis

We have the restriction that all container labels can only depend on fixed symbols. This restriction can be an annoyance in some cases, one of which is outlined in figure 7.2. Here, we have a method that requires both of its parameters to be within the same container. The code in Main clearly passed two references that are within the same container, but because the variable demo is not declared as fixed, the type system asserts that the container labels are unknown and the method call does not type check.

This issue could be fixed with data flow analysis. If a symbol is determined to be fixed within a region of code, then an automatic rewrite could be applied to inject a temporary reference that is fixed.

Chapter 8

Conclusion

In this thesis, we've taken the abstract idea of a container isolating its contents from the outside world and built two systems based on this idea. First, a broader implementation in Haskell that attempted to solve pragmatic issues such as reducing the amount of extra syntax needed using a container inferencing algorithm. By developing example code and test cases, we have a sense of what programming in a container language could be like. This system is too immature to evaluate its suitability as an implementation language for real object-oriented programs. However, in the code that has been written, no major problems have been found and simple data structures such as a linked list with containment were easy to develop.

The second system of operational semantics showed that our dependent type system is sound. Code using containers can be statically typed such that there will be no run-time violation of containers. This is an important result since it means that all container type information can be erased at run-time and a container system can run without any additional overhead compared to an equivalent system without containers.

Adding the ability to enforce self-containment creates a *kind* of object which is simple to reason about. There are a number of algorithms that are easy to implement when an object is self contained, but extremely tricky otherwise. Serialization, deep object copying and deep object equality comparisons can be built into the language. The opportunity exists to eliminate a significant amount of glue code by leveraging containers and self-contained objects.

A great deal of work would be required to turn this demonstration system into a practical programming language. This thesis has shown that core idea of containment

is sound, and demonstrated that containers can have practical benefits. We hope this work will motivate future research into containers.

References

- [1] Antonio Albano, Giorgio Ghelli, and Renzo Orsini. Fibonacci: A programming language for object databases. *VLDB J.*, 4(3):403–444, 1995.
- [2] Jonathan Aldrich, Valentin Kostadinov, and Craig Chambers. Alias annotations for program understanding. In Mamdouh Ibrahim and Satoshi Matsuoka, editors, *Proceedings of the 2002 ACM SIGPLAN Conference on Object-Oriented Programming Systems, Languages and Applications, OOPSLA 2002, Seattle, Washington, USA, November 4-8, 2002*, pages 311–330. ACM, 2002.
- [3] Lorenzo Bettini, Ferruccio Damiani, Kathrin Geilmann, and Jan Schäfer. Combining traits with boxes and ownership types in a java-like setting. *Sci. Comput. Program.*, 78(2):218–247, 2013.
- [4] Gavin M Bierman, MJ Parkinson, and AM Pitts. Mj: An imperative core calculus for java and java with effects. Technical report, University of Cambridge, Computer Laboratory, 2003.
- [5] Chandrasekhar Boyapati, Barbara Liskov, and Liuba Shrira. Ownership types for object encapsulation. In Alex Aiken and Greg Morrisett, editors, *Conference Record of POPL 2003: The 30th SIGPLAN-SIGACT Symposium on Principles of Programming Languages, New Orleans, Louisiana, USA, January 15-17, 2003*, pages 213–223. ACM, 2003.
- [6] John Boyland. Alias burying: Unique variables without destructive reads. *Softw. Pract. Exp.*, 31(6):533–553, 2001.
- [7] Nicholas Robert Cameron, Sophia Drossopoulou, and James Noble. Understanding ownership types with dependent types. In Dave Clarke, James Noble, and Tobias

- Wrigstad, editors, *Aliasing in Object-Oriented Programming. Types, Analysis and Verification*, volume 7850 of *Lecture Notes in Computer Science*, pages 84–108. Springer, 2013.
- [8] Dave Clarke, Johan Östlund, Ilya Sergey, and Tobias Wrigstad. Ownership types: A survey. In Dave Clarke, James Noble, and Tobias Wrigstad, editors, *Aliasing in Object-Oriented Programming. Types, Analysis and Verification*, volume 7850 of *Lecture Notes in Computer Science*, pages 15–58. Springer, 2013.
- [9] David G. Clarke, John Potter, and James Noble. Ownership types for flexible alias protection. In Bjørn N. Freeman-Benson and Craig Chambers, editors, *Proceedings of the 1998 ACM SIGPLAN Conference on Object-Oriented Programming Systems, Languages & Applications (OOPSLA '98), Vancouver, British Columbia, Canada, October 18-22, 1998*, pages 48–64. ACM, 1998.
- [10] W. Dietl and P. Müller. Ownership Type Systems and Dependent Classes. In *Foundations of Object-Oriented Languages (FOOL)*, January 2008.
- [11] Wei Huang, Werner Dietl, Ana Milanova, and Michael D. Ernst. Inference and checking of object ownership. In James Noble, editor, *ECOOP 2012 - Object-Oriented Programming - 26th European Conference, Beijing, China, June 11-16, 2012. Proceedings*, volume 7313 of *Lecture Notes in Computer Science*, pages 181–206. Springer, 2012.
- [12] Paley Li, Nicholas Cameron, and James Noble. Sheep cloning with ownership types. In *FOOL 2012: 19th International Workshop on Foundations of Object-Oriented Languages*, page 59, 2012.
- [13] Heather Miller, Philipp Haller, Eugene Burmako, and Martin Odersky. Instant pickles: generating object-oriented pickler combinators for fast and extensible serialization. In Antony L. Hosking, Patrick Th. Eugster, and Cristina V. Lopes, editors, *Proceedings of the 2013 ACM SIGPLAN International Conference on Object Oriented Programming Systems Languages & Applications, OOPSLA 2013, part of SPLASH 2013, Indianapolis, IN, USA, October 26-31, 2013*, pages 183–202. ACM, 2013.
- [14] Alan Mycroft and Janina Voigt. Notions of aliasing and ownership. In Dave Clarke, James Noble, and Tobias Wrigstad, editors, *Aliasing in Object-Oriented Programming. Types, Analysis and Verification*, volume 7850 of *Lecture Notes in Computer Science*, pages 59–83. Springer, 2013.

