

A Very Modal Model of a Modern, Major, General Type System

Andrew W. Appel*

Princeton University
INRIA Rocquencourt
appel@princeton.edu

Paul-André Mellies

CNRS
Université Paris 7
mellies@pps.jussieu.fr

Christopher D. Richards*

Princeton University
richards@cs.princeton.edu

Jérôme Vouillon

CNRS
Université Paris 7
vouillon@pps.jussieu.fr

Abstract

We present a model of recursive and impredicatively quantified types with mutable references. We interpret in this model all of the type constructors needed for typed intermediate languages and typed assembly languages used for object-oriented and functional languages. We establish in this purely semantic fashion a soundness proof of the typing systems underlying these TILs and TALs—ensuring that every well-typed program is safe. The technique is generic, and applies to any small-step semantics including λ -calculus, labeled transition systems, and von Neumann machines. It is also simple, and reduces mainly to defining a Kripke semantics of the Gödel-Löb logic of provability. We have mechanically verified in Coq the soundness of our type system as applied to a von Neumann machine.

Categories and Subject Descriptors F.3.3 [Logics and Meanings of Programs]: Studies of Program Constructs—Type structure; F.3.2 [Logics and Meanings of Programs]: Semantics of Programming Languages; F.3.1 [Logics and Meanings of Programs]: Specifying and Verifying and Reasoning about Programs—Mechanical verification; F.4.1 [Mathematical Logic and Formal Languages]: Mathematical Logic—Modal logic

General Terms Languages, Theory

Keywords Recursive types, impredicative polymorphism, mutable references, Kripke models

1. Introduction

We wish to compile languages such as ML and Java into typed intermediate languages and typed assembly languages. These TILs and TALs are particularly difficult to design,

* Supported in part by NSF grants CCF-0540914 and CNS-0627650.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

To appear in *POPL'07*, January 17–19, 2007, Nice, France.
Copyright © 2007 ACM.

because in order to describe the program transformations applied in the course of compilation, they require a very rich and expressive type system, including

- intersection and union types;
- types for address arithmetic;
- types for mutable references, since the source languages have mutable references (or object fields);
- recursive types, in order to describe data structures loaded in memory such as lists and trees;
- type quantification, which is used to describe object and function-closure representations.

Not only that: type quantification should be *impredicative*, because the types which may be abstracted (i.e., put in private fields of objects or in function-closure environments) can be themselves object or function-closure (i.e., quantified) types. In particular, it is not possible to stratify the type system into levels k such that a quantification $\forall_k \alpha. \tau$ is always instantiated at level $< k$.

Putting all these type ingredients together in a low-level language is an intricate exercise. A formal proof of *soundness*—any well-typed program does not go wrong—is thus recommended for any type system for such TILs and TALs.

Contributions. In this paper we will establish the soundness of a type system including all the ingredients mentioned earlier, and more. Our semantics of impredicative references is an improvement over previous work because it removes a cumbersome restriction from the typing rules for quantifiers (see Section 13.1). Our approach is also novel in that we have decomposed the semantic types in such a way that the principles of *safety*, *induction over the future*, *type preservation*, and so on are each identifiable as a specific operator of the type system. Standard concepts such as (function) code pointers, subtyping, contractiveness, mutable and immutable references are constructed from these new primitives.

Overview. Although our type system could be proved sound syntactically—using progress and preservation—we will not use progress and preservation. Hence we do not need to formally list all our type constructors and typing rules—because we never do induction over them—but the

reader may examine Tables 1 and 2 and Figures 4–7. We will proceed in a purely semantic fashion, for two reasons at least. First, we find intuition easier to come by in this way: once types are interpreted in a relevant and felicitous way, it becomes much easier to combine or decompose them than if they remained ciphers rendered meaningful only by the typing rules.

Second, we enjoy the modularity and genericity of the semantic methods. We demonstrate this in the paper, with a nicely modularized proof of soundness which is applicable in principle to any calculus with a small-step semantics. We discuss here von Neumann machines and the call-by-value λ -calculus with references; our methods should also work well for labeled transition systems used in the intermediate representations of compilers.

To establish our soundness result, we construct a model of recursive and impredicatively quantified types with mutable references. We start from the idea of *approximation* which pervades all the semantic research in the area. If we type-check $v : \tau$ in order to guarantee safety of just the next k computation steps, we need only a k -approximation of the typing-judgment $v : \tau$. We express this idea here using a Kripke semantics whose possible-worlds accessibility relation R is well-founded: every path from a world w into the future must terminate. In this work, worlds characterize abstract computation states, giving an upper bound on the number of future computations steps and constraining the contents of memory. (We define the set of worlds precisely in Section 6.) We write

$$w \Vdash v : \tau$$

to mean that the value v has type τ in world w . This leads us to investigate a modal logic adapted to this Kripke semantics (Section 2). A judgment $\sigma \vdash \tau$ in this logic means that every value v of type σ at any world w has also type τ at that world w . Then (Section 4) a modal operator \triangleright (pronounced “later”) quantifies over all worlds (all times) *strictly* in the future. The strictness of \triangleright turns out to be a key technical insight, if we compare it to previous work [Ahm04] which used instead a construction analogous to the modal operator necessity \Box quantifying over worlds *now and* in the future. Indeed, the combination of a well-founded R with a strict modal operator \triangleright provides a clean induction principle to the logic, called the *Löb* rule,

$$\frac{\triangleright \tau \vdash \tau}{\vdash \tau},$$

where $\vdash \tau$ can be read as asserting the safety of the program described by τ . (Equation 16 gives a precise definition.)

The resulting Gödel-Löb logic is ideal for computations that can’t go on forever, and where information gets less and less precise as the computation goes on. Suppose that we want to type-check a program written for a von Neumann machine with arbitrary control flow (Section 11). Let τ give

the types of all the labels appearing in the program. We want to show that τ is an induction hypothesis for the safety of the program, at least as far as the end of the worlds. We first assume $\triangleright \tau$; that is, to a weaker, or “later” approximation, the type τ holds for any exit label to which we might jump in (at least) one step. Then we analyze each basic block in the program to show that the type τ (that is, the stronger approximation, without \triangleright) describes each entry label. From this we would like to conclude by induction that τ absolutely describes the program, with no hypothesis; this is precisely what the Löb rule enables.

Of course, some computations *should* go on forever. But the proof scheme above works for any *finite prefix* of the computation—so the program is safe for an arbitrary number of steps. We obtain in this way (Section 12) a proof of soundness of the type system for a von Neumann machine—of which we have machine-checked proofs¹ in Coq, using a straightforward shallow embedding of our model.

Related work. Our model unifies and generalizes several previous threads of research: the indexed model [AM01, AAV02, Ahm04], the orthogonality model [MV05], and the approximation modality [Nak01]. We will discuss some of these in Section 13.

2. Kripke semantics

Kripke models [Kri63] originated as models of modal logics; subsequently they have been shown useful for giving the semantics of type systems [MM91]. The main idea behind Kripke models is that truth (respectively, membership in a type) is not absolute but relative to some appropriate notion of world, or state, and that these states can be related precisely to each other.

In this section we will interpret worlds as characterizing abstract properties of the current state of computation. In particular, in a system with mutable references, each world contains a memory typing (Section 6); but the model can be more general than that, so we will start with an abstract characterization.

Models. A *model* is a triple (W, R, \Vdash) , where W is a nonempty set of worlds w ; an accessibility relation $R \subset W \times W$ describes when it is possible to move from one world to another; and \Vdash assigns truth values to typing judgments with respect to a given world, that is,

$$w \Vdash v : \tau. \tag{1}$$

Here $w \in W$ is a world, $v \in V$ is a *value* of our operational calculus, and $\tau \in \text{Type}$ is a type.

We apply our Kripke models to small-step operational semantics with a memory or store indexed by locations Loc ; but there are still many flavors of such semantics:

¹Available at <http://www.pps.jussieu.fr/~vouillon/smot/>

Notation	Eqn.	Description
<code>top</code>	2	Top type
<code>bot</code>	3	Bottom type
$\tau \wedge \sigma$	4	Intersection type
$\tau \vee \sigma$	5	Untagged union type
$\tau \Rightarrow \sigma$	6	Implication type
$\forall x:A.\tau$	7	Universal type
$\exists x:A.\tau$	8	Existential type
<code>!τ</code>	9	Configs where all values have τ
<code>?τ</code>	10	Configs where some value has τ
<code>just u</code>	13	Singleton type
<code>slot(j, τ)</code>	15	Von Neumann vector element type
<code>▷ τ</code>	18	“Later” modality
<code>rec F</code>	24	Recursive type
$l \mapsto \tau$	26	Configs where $\Psi(l) = \tau$
<code>plus(i, j, k)</code>	57	Asserts $i + j = k$
<code>greater(i, j)</code>	58	Asserts $i > j$
<code>safemem(m)</code>	65	Type of configs safe for memory m
<code>unit</code>	70	Unit type

Figure 1. Primitive types, operators

- In λ -calculus with mutable references, a *state* is a pair (e, m) where e is an expression and m is a *store*, a mapping from Loc to V . States evolve by a small-step relation $(e, m) \mapsto (e', m')$; if there is no such (e', m') then either e is a value v or (e, m) is stuck. A world w characterizes the store m .
- On a von Neumann machine, the state is a pair (v, m) , where $v \in V$ is a vector of Loc representing pointers in registers, including the program counter. The memory m is a mapping from Loc to Loc containing program and data. The small-step relation is $(v, m) \mapsto (v', m')$. A world w characterizes the contents of memory.
- In a labeled transition system such as is used in flow analysis, w characterizes the store and v contains both the local variables and the control-flow state.

Clearly, the same value v may or may not have type τ depending on the world w , that is, depending on the data structures in memory that v points to. Accordingly, we call a pair (w, v) a *configuration* (abbreviated “*config*”):

$$\text{Config} = W \times V,$$

and define a type $\tau \in \text{Type}$ as a set of configurations. Then,

$$(w, v) \in \tau \quad \text{and} \quad w \Vdash v : \tau$$

are two alternative notations expressing the same fact. We will show how our semantics connects the relation R between worlds and the relation \mapsto between states.

3. Primitive type constructors

Our approach to the modeling of a type system is to define a few primitive type constructors—a type is just any set of

Notation	Eqn.	Description
$\tau \iff \sigma$	11	Type of configs in τ iff in σ
<code>teq(τ, σ)</code>	12	Asserts τ equal to σ
$u \circ \tau$	14	Asserts u has τ in current world
$\Box \tau$	19	Necessity modality
<code>validmem(m)</code>	27	Asserts m valid
<code>ref τ</code>	32	Mutable reference type
<code>offset(n, τ)</code>	59	Pointer arithmetic type
<code>pair(σ, τ)</code>	60	Pair type
<code>list τ</code>	61	List type
<code>iref τ</code>	62	Immutable reference type
<code>safe</code>	66	Type of safe configs
$\tau \rightarrow \sigma$	71	Function type
<code>codeptr(τ)</code>	72	Code pointer (continuation) type
Δ_p	75	Program code type
Γ_p	79	Program invariants type
<code>boundary l</code>	88	Asserts memory $> l$ unallocated

Figure 2. Synthesized types, operators

configurations—and then use these primitives to synthesize all the constructors needed for type-checking programs.

We can start with intersection types, union types, and their identities, as well as an implication type:

$$\text{top} \stackrel{\text{prim}}{=} \{(w, v) \mid \text{True}\} \quad (2)$$

$$\text{bot} \stackrel{\text{prim}}{=} \{\} \quad (3)$$

$$\tau \wedge \sigma \stackrel{\text{prim}}{=} \tau \cap \sigma \quad (4)$$

$$\tau \vee \sigma \stackrel{\text{prim}}{=} \tau \cup \sigma \quad (5)$$

$$\sigma \Rightarrow \tau \stackrel{\text{prim}}{=} \{(w, v) \mid (w, v) \in \sigma \Rightarrow (w, v) \in \tau\} \quad (6)$$

We write $c(\tau) \stackrel{\text{prim}}{=} \dots \tau \dots$ to indicate that the primitive type-constructor c is defined by the mathematical expression $\lambda \tau. \dots \tau \dots$ in the underlying logic. For our machine-checked proofs in Coq [B⁺98], the underlying logic is the Calculus of Inductive Constructions [PM93].

Quantification. Quantification over types is needed for polymorphism, data abstraction, and representation of closures. Besides, we are interested here in a λ -calculus or von Neumann machine equipped with a store (memory) of mutable references with address arithmetic. To reason about dataflow and addressability in such programs, our type system should also be able to quantify over locations $l \in \text{Loc}$ and memories $m \in \text{Mem}$. Our quantifiers are thus,

$$\forall x:A.\tau \stackrel{\text{prim}}{=} \bigcap_{a \in A} \tau[a/x] \quad (7)$$

$$\exists x:A.\tau \stackrel{\text{prim}}{=} \bigcup_{a \in A} \tau[a/x] \quad (8)$$

where A is either Type , Loc , or Mem .

Quantification over values in a world. Another kind of quantification is over values in the current² world:

$$! \tau \stackrel{\text{prim}}{=} \{(w, v) \mid \forall v'. (w, v') \in \tau\} \quad (9)$$

$$? \tau \stackrel{\text{prim}}{=} \{(w, v) \mid \exists v'. (w, v') \in \tau\} \quad (10)$$

The type $! \tau$ (“everywhere τ ”) means that every value in the current world has type τ ; and $? \tau$ (“somewhere τ ”) means that some value in the current world has type τ .

One strength of our approach is that we can define just a few primitive type constructors, and from those we can synthesize many more constructors of interest. We will use the notation $c(\tau) \stackrel{\text{syn}}{=} \dots \tau \dots$ to mean that the constructor c is synthesized from a combination $\dots \tau \dots$ of our type constructors, without dipping into the underlying logic.

$$\tau \iff \sigma \stackrel{\text{syn}}{=} \tau \Rightarrow \sigma \wedge \sigma \Rightarrow \tau \quad (11)$$

$$\text{teq}(\tau, \sigma) \stackrel{\text{syn}}{=} !(\tau \iff \sigma) \quad (12)$$

The judgment $w \Vdash v : \tau \iff \sigma$ means that (in world w) v has type τ iff it has type σ . In contrast, $w \Vdash v : \text{teq}(\tau, \sigma)$ means that every value in world w has type τ iff it has type σ ; that is, teq is our notion of type equality. In each world, the type $\text{teq}(\tau, \sigma)$ is equal to either top or bot; we call such types *world types* because they depend only on the world, not on the value. Henceforth if τ is a world type we will write $w \Vdash \tau$ to mean $w \Vdash v : \tau$ for an arbitrary v .

Vector values. As we have said, our type system applies to several computational models (e.g., λ -calculus, von Neumann) and in some of these we need to type several kinds of values. We have *locations* $l : \text{Loc}$ that index a mutable store m ; *storable values* $u : SV$ that are the range of m (contents of memory cells); and *values* $v : V$. We assume $\text{Loc} \subset SV$.

On a von Neumann machine, $SV = \text{Loc}$ and v is a vector of locations—one could think of a register-bank—indexed by a natural number j . That is, if v is a value then $v(j)$ is a Loc . In order to type locations, we choose an injective function $\vec{\cdot}$ from storable values to values, for instance $\vec{u} \stackrel{\text{def}}{=} \lambda j. u$. This way the same set of types can be used for all kinds of values.

In λ -calculus, $SV = V$ is the usual set of values in λ -calculus, we have $\text{Loc} \subsetneq SV$ by syntactic inclusion, and we take $\vec{u} \stackrel{\text{def}}{=} u$.

Singletons and slots. The singleton type $\text{just } u$ is the type of the single storable value u . Since $\text{Loc} \subset SV$ we can also write $\text{just } l$ for a location l .

$$\text{just } u \stackrel{\text{prim}}{=} \{(w, v) \mid v = \vec{u}\} \quad (13)$$

The construction $!(\text{just } u \Rightarrow \tau)$ asserts that the storable-value u has type τ in the current world, which we write $u \circ \tau$.

$$u \circ \tau \stackrel{\text{syn}}{=} !(\text{just } u \Rightarrow \tau) \quad (14)$$

²In the context of a definition that depends on a formal parameter $w \in W$, we refer to w as the “current” world.

We can define “ $\exists l : \text{Loc. just } l \wedge \tau(l)$ ”, the type of locations l which satisfy some property $\tau(l)$ (where $\tau(l)$ is a world type).

The type $\text{slot}(j, \tau)$ characterizes values v such that (informally) the j th slot has type τ .

$$\text{slot}(j, \tau) \stackrel{\text{prim}}{=} \{(w, v) \mid w \Vdash v(j) \circ \tau\} \quad (15)$$

To say that register 2 has the value 3 we write $\text{slot}(2, \text{just } 3)$.

4. Necessity and the modal operator \triangleright

In this section we will explain an operator that encapsulates the notion of induction; then we will demonstrate its application to recursive types (Section 5), recursive data structures (Sections 6, 8) and recursive functions and control-flow (Section 11).

Judgments. We introduce a notion of judgment and its semantic interpretation. Given two types σ and τ , we write

$$\sigma \vdash \tau \quad (16)$$

when the type σ is a subset of the type τ , or equivalently, when for every world w and for every value v ,

$$w \Vdash v : \sigma \quad \text{implies} \quad w \Vdash v : \tau.$$

More generally, we write $\sigma_1, \dots, \sigma_n \vdash \tau$ to mean $\sigma_1 \wedge \dots \wedge \sigma_n \vdash \tau$. We write $\top \vdash \tau$ to mean $\text{top} \vdash \tau$.

It is possible to reason about these judgments using a rich set of inference rules. These rules can be used to derive a larger number of type properties without having to expand the definition of type constructors and reason in the underlying logic. Some of these rules are given in Figures 4–7.

The “later” operator \triangleright . In any model (W, R, \Vdash) we require the accessibility relation R to be transitive and *well-founded*:

$$\begin{aligned} R \text{ well-founded} &\stackrel{\text{def}}{=} \\ &\text{from any world } w \\ &\text{there is no infinite path } w R w' R w'' R \dots \end{aligned} \quad (17)$$

Intuitively, $w R w'$ means that the world w' comes at a strictly later stage than the world w . This leads us to define the “later” operator:

$$\triangleright \tau \stackrel{\text{prim}}{=} \{(w, v) \mid \forall w'. w R w' \Rightarrow (w', v) \in \tau\} \quad (18)$$

That is, $w \Vdash v : \triangleright \tau$ precisely when $w' \Vdash v : \tau$ for any world w' coming strictly later than the world w .

The following properties follow easily.

Lemma 4.1.

1. \triangleright is monotone: if $\sigma \vdash \tau$, then $\triangleright \sigma \vdash \triangleright \tau$.
2. \triangleright distributes over intersection: $\triangleright \bigwedge \tau_i = \bigwedge \triangleright \tau_i$.

The necessity operator \Box . There is a related operator \Box , pronounced “necessarily”, which means “now and later” and is defined just that way:

$$\Box \tau \stackrel{\text{syn}}{=} \tau \wedge \triangleright \tau \quad (19)$$

The next proposition shows that the operator \Box is a *closure operator* on types, or what is called a *comonad* in category theory³—although in this paper we will not use the toolkit of category theory.

Lemma 4.2.

1. \Box is monotone: if $\sigma \vdash \tau$, then $\Box \sigma \vdash \Box \tau$.
2. For any type τ , $\Box \tau \vdash \tau$.
3. If $\Box \sigma \vdash \tau$, then $\Box \sigma \vdash \Box \tau$.
4. \Box distributes over intersection: $\Box \bigwedge \tau_i = \bigwedge \Box \tau_i$.

Necessary types. We are interested in types τ such that, once a value v has type τ in some world w , it has type τ in every future world. Such types are called *necessary* types:

$$\tau \text{ necessary} \stackrel{\text{def}}{=} \tau \vdash \triangleright \tau \quad (20)$$

Not all of our types are necessary. In particular, since the mutable store evolves from one world to the next, it will be possible to construct types (sets of configurations) such that $\tau \vdash \triangleright \tau$ does not hold. In practice, our datatypes (types for references, code pointers, ...), are necessary. However, some logical constructions on types do not preserve necessity: for instance, the fact that two types τ and τ' are necessary does not imply that the type $\tau \Rightarrow \tau'$ is necessary.

Since necessary types are much easier to reason about, we want a convenient way to construct them. This is precisely what the modal operators \triangleright and \Box achieve for us: they transform arbitrary types into necessary types. Let us state formally a few properties relating necessity to these two modal operators:

Lemma 4.3.

1. For every type τ , the types $\Box \tau$ and $\triangleright \tau$ are necessary,
2. A type τ is necessary iff $\tau = \Box \tau$, that is, $\vdash \text{teq}(\tau, \Box \tau)$.
3. The intersection of necessary types is necessary.

The Löb rule. Our hypothesis that the accessibility relation R is well-founded is reflected in our modal logic by the following induction principle:

Theorem 4.4 (Löb Rule). For every type τ ,

$$\frac{\triangleright \tau \vdash \tau}{\vdash \tau} \quad (21)$$

³ In fact, the \Box operator is the free comonad of the \triangleright operator, both understood as functors in the category of types ordered by inclusion. This fact may be proved directly, or deduced from our definition 20 of necessary types, and Section 9.4 of Barr and Wells [BW83].

Proof. As R is well-founded, it satisfies the following induction principle:

$$\forall P. (\forall w. (\forall w'. wRw' \Rightarrow P(w')) \Rightarrow P(w)) \Rightarrow \forall w. P(w).$$

By taking $P(w) = (w, v) \in \tau$, we get:

$$\forall v. (\forall w. (\forall w'. wRw' \Rightarrow (w', v) \in \tau) \Rightarrow (w, v) \in \tau) \Rightarrow \forall w. (w, v) \in \tau.$$

Finally, by distributing the quantification over v , we have:

$$(\forall w. v. (\forall w'. wRw' \Rightarrow (w', v) \in \tau) \Rightarrow (w, v) \in \tau) \Rightarrow \forall w. v. (w, v) \in \tau.$$

This is exactly what the rule 21 means. \square

Corollary 4.5 (Generalized Löb Rule).

$$\frac{\sigma, \triangleright \tau \vdash \tau \quad \sigma \text{ necessary}}{\sigma \vdash \tau}$$

Proof. By a short logical derivation:

$$\frac{\frac{\frac{\frac{\frac{\sigma, \triangleright \tau \vdash \tau}{\sigma, \triangleright (\sigma \wedge (\sigma \Rightarrow \tau)) \vdash \tau}}{\sigma, (\triangleright \sigma) \wedge \triangleright (\sigma \Rightarrow \tau) \vdash \tau}}{\sigma \wedge \triangleright \sigma, \triangleright (\sigma \Rightarrow \tau) \vdash \tau}}{\sigma, \triangleright (\sigma \Rightarrow \tau) \vdash \tau} \sigma \text{ necessary}}{\triangleright (\sigma \Rightarrow \tau) \vdash \sigma \Rightarrow \tau} \text{Theorem 4.4}}{\vdash \sigma \Rightarrow \tau} \sigma \text{ necessary}}{\sigma \vdash \tau} \square$$

5. Recursive types

Given a type operator $F \in \text{Type} \rightarrow \text{Type}$, we can construct the recursive type $\text{rec } F$. As in most models of recursive types, if F is *contractive* then $\text{rec } F$ is a fixed point of F . Informally, a contractive function is one such that if τ is approximately equal to σ , then $F\tau$ is more accurately equal to $F\sigma$. We can express contractiveness using the primitives of our type system—in fact, contractiveness and nonexpansiveness⁴ have very concise, intuitive definitions.

$$F \text{ contractive} \stackrel{\text{def}}{=} \triangleright \text{teq}(\tau, \sigma) \vdash \text{teq}(F\tau, F\sigma) \quad (22)$$

$$F \text{ nonexpansive} \stackrel{\text{def}}{=} \Box \text{teq}(\tau, \sigma) \vdash \text{teq}(F\tau, F\sigma) \quad (23)$$

The hypothesis $\triangleright \text{teq}(\tau, \sigma)$ means that (in the current world) τ is perhaps not yet equal to σ , but in every future world it will be. What is it about the future worlds that makes them accept $\tau = \sigma$? The answer is that the world-accessibility relation R is well founded, which informally means that in the current world there is a finite number k of computation

⁴ In a way now standard in type theory [MPS86], nonexpansive type constructors may be composed with contractive constructors to make contractive constructors.

steps remaining. Perhaps we can call a function that, given an argument of type τ would get stuck on the k th step, but given an argument of type σ executes safely for k steps; in this case $w \not\vdash v : \text{teq}(\tau, \sigma)$. But (by the well-foundedness of R), the next world has $< k$ steps of computation remaining to it, so it is less able to distinguish τ from σ . So F contractive means that if τ is indistinguishable from σ in every future world, then $F(\tau)$ is indistinguishable from $F(\sigma)$ in this world.

On the other hand, F nonexpansive means that if τ is indistinguishable from σ in this world and every future world, then $F(\tau)$ is indistinguishable from $F(\sigma)$ in this world. For example, we will later introduce the constructor `ref` which is contractive, and `offset` which is nonexpansive.

Given a type operator F , we define by induction on worlds the fixpoint candidate $\text{rec } F$.

$$\begin{aligned} \text{rec } F &\stackrel{\text{prim}}{=} \text{the unique set } \sigma \text{ such that} \\ (w, v) \in \sigma &\quad \text{iff} \\ (w, v) \in F(\{(w', v) \mid wRw' \wedge (w', v) \in \sigma\}). \end{aligned} \quad (24)$$

By the well-foundedness of R , this is defined for all type operators F —in order to find out whether $(w, v) \in \text{rec } F$ one only needs to consider configurations $(w', v') \in \text{rec } F$ where wRw' . This fixpoint candidate is easily defined as a structural `Fixpoint` function in `Coq`.

Lemma 5.1. *When F is contractive, the type $\text{rec } F$ is a fixpoint of F , that is,*

- $\text{rec } F \vdash F(\text{rec } F)$ *(unfolding)*
- $F(\text{rec } F) \vdash \text{rec } F$ *(folding)*

These two properties can also be phrased as a single assertion: $\vdash \text{teq}(\text{rec } F, F(\text{rec } F))$.

Lemma 5.2. *The fixpoint of a contractive operator F is unique: if $\vdash \text{teq}(\tau, F(\tau))$ and $\vdash \text{teq}(\sigma, F(\sigma))$, then $\vdash \text{teq}(\tau, \sigma)$.*

6. A Kripke semantics of stores

One advantage of our semantic approach is that we did not need to fully specify our Kripke model in order to interpret the logical part of our typing system (intersection and union types, quantification, necessity, recursive types). But since we are particularly interested here in languages with mutable references—we would like to include a notion of *reference type* in our system—this requires us to explain what Kripke model of worlds w and accessibility relation wRw' suitable for mutable references we have in mind.

Recall that, intuitively, the current world w shall be used for two purposes: (1) to ensure that the accessibility relation R is well founded, and (2) to constrain what values may be contained in memory locations. That is, a world $w = (n, \Psi)$ will be a pair of an index $n \in \mathbb{N}$ (which counts down as time advances) and a store typing Ψ . As one should expect, this store typing Ψ will be a partial function from

memory locations $l \in \text{Loc}$ to types. However, we will have to be extremely careful, and restrict ourselves to only some of these functions. Indeed, the mutually recursive equations in W and Type below have no solution (by a cardinality argument):

$$\begin{aligned} W &= \mathbb{N} \times (\text{Loc} \rightarrow \text{Type}) \\ \text{Type} &= \mathcal{P}(W \times V). \end{aligned} \quad (25)$$

The significance of \triangleright . Let us carry on our analysis. Informally, a memory $m \in \text{Mem}$ is well typed with respect to a store typing Ψ if for all locations l in the domain of Ψ , the value $m(l)$ has type $\Psi(l)$. Let us think ahead and imagine that the set of worlds W is already constructed. One could then define the world type $l \mapsto \tau$ which characterizes the worlds (n, Ψ) in which the memory-typing for l is type τ :

$$l \mapsto \tau \stackrel{\text{prim}}{=} \{(n, \Psi), v \mid \Psi(l) = \tau\} \quad (26)$$

Then, we may declare that a memory m is well-typed when the following assertion holds:

$$\begin{aligned} \text{validmem}(m) &\stackrel{\text{syn}}{=} \\ \forall l : \text{Loc}. \forall \tau : \text{Type}. l \mapsto \tau \Rightarrow \triangleright(m(l) \circ \tau) \end{aligned} \quad (27)$$

In this definition we write $\triangleright(m(l) \circ \tau)$. There is some value u in memory at address l , and we guarantee to every future world that $u \circ \tau$. We don't need to guarantee $u \circ \tau$ in the *current* world because it takes one step just to dereference $m(l)$, and in that step we move to a future world.

This use of the \triangleright operator rather than the \square operator is crucial in order to solve the cardinality issue. Indeed, for a configuration $((n, \Psi), v)$, only the configurations of index strictly less than n are then relevant in the type $\Psi(l)$.

The Kripke semantics. This informal discussion enables us to stratify our definition of worlds and types, and to start by defining “finitely stratified” worlds and types of rank n by mutual induction:

$$\begin{aligned} W_n &\stackrel{\text{def}}{=} \{n\} \times (\text{Loc} \rightarrow \text{Type}_n) \\ \text{Type}_n &\stackrel{\text{def}}{=} \mathcal{P}\left(\bigcup_{k < n} W_k \times V\right) \end{aligned} \quad (28)$$

The set of worlds is then defined as the union of finitely stratified worlds:

$$W \stackrel{\text{def}}{=} \bigcup_{n \in \mathbb{N}} W_n.$$

Finally, the set of types is defined⁵ according to the principles of Kripke semantics presented in Section 2:

$$\text{Type} \stackrel{\text{def}}{=} \mathcal{P}(W \times V).$$

⁵ **Remark.** The set Type is isomorphic to the set $\prod_{n \in \mathbb{N}} \mathcal{P}(W_n \times V)$ obtained as projective limit of the sequence of finitely stratified types $\{\text{Type}_n\}_{n \in \mathbb{N}}$, where $\text{Type}_n \cong \prod_{k < n} \mathcal{P}(W_k \times V)$.

We have the inclusions

$$\text{Type}_0 \subset \dots \subset \text{Type}_n \subset \dots \subset \text{Type}$$

and a family of approximation functions

$$\lfloor _ \rfloor_n : \text{Type} \rightarrow \text{Type}_n$$

defined thus:

$$\lfloor \tau \rfloor_n \stackrel{\text{def}}{=} \{((k, \Psi), v) \in \tau \mid k < n\}. \quad (29)$$

These approximation functions are lifted pointwise to store typings Ψ such that

$$\begin{aligned} \text{dom} \lfloor \Psi \rfloor_n &= \text{dom} \Psi, \\ \lfloor \Psi \rfloor_n(l) &= \lfloor \Psi(l) \rfloor_n. \end{aligned} \quad (30)$$

Finally, the accessibility relation between worlds is defined as follows:

$$\begin{aligned} (n, \Psi) R(n', \Psi') &\stackrel{\text{def}}{=} \\ n > n' \wedge \forall l \in \text{dom} \Psi. \Psi'(l) &= \lfloor \Psi(l) \rfloor_{n'}. \end{aligned} \quad (31)$$

Intuitively, when moving to a future world, at least one step must be taken; the types of allocated locations must be approximately preserved; and new locations may be allocated.

This completes our definition of a Kripke model for mutable references. The two types $l \mapsto \tau$ and $\text{validmem}(m)$ which were used for earlier discussion, are then defined as in equations 26 and 27.

Reference types. We are ready now to define reference types. Intuitively, a reference of type τ is a pointer to a location l of type τ . Because of the inherent stratification of store typing, this is actually too strong a requirement. We need to replace it by the weaker but sufficient requirement that the type τ' of the location l coincides with the type τ in the future. Hence the following definition of reference types:

$$\text{ref } \tau \stackrel{\text{syn}}{=} \exists l : \text{Loc. just } l \wedge \exists \tau' : \text{Type. } (l \mapsto \tau' \wedge \triangleright \text{teq}(\tau, \tau')) \quad (32)$$

Note that this definition is given solely in terms of type constructors already defined. We have been arguing for the usefulness of a modal approach to type semantics, as well as a finer-grained decomposition of such semantics; definition 32 exemplifies both of these ideas in a nice way. To underscore the flexibility of this decomposition, we speculate that one could synthesize, without dipping back into the underlying logic, a replacement for “ $l \mapsto \tau'$ ” that would yield a region-enabled definition of the ref constructor, and thereby permit deallocation.

7. Type preservation

Suppose $v : \tau$ in some state, and then a memory location is updated—either to initialize a new reference or to update

an old one. One of the usual difficulties in a semantics of references is to ensure that $v : \tau$ in the new state. For example, if $\tau = \text{list}(\text{ref } \sigma)$ (with list defined in equation 61, below) how can we ensure that the store to memory has not modified one of the list cells or one of the references? In our framework the solution is (1) use R to constrain the evolution of worlds; (2) make sure that τ is *necessary*, that is, τ is preserved as we change worlds; and (3) use $\text{validmem}()$ to constrain the machine *states* to follow the evolution of *worlds*. Concretely, we demonstrate that, with our choice of accessibility relation, reference types are necessary and all interesting store operations (memory access, allocation and update) are permitted.

We use the following technical lemma several times.

Lemma 7.1 (Approximation and Equality). *A type and its approximation can be related as follows.*

$$(k, \Psi) \Vdash \triangleright \text{teq}(\tau, \lfloor \tau \rfloor_k)$$

We show that an allocated memory location remains allocated forever and that its type is preserved.

Lemma 7.2 (Store Typing Preservation).

$$l \mapsto \tau \vdash \triangleright \exists \sigma : \text{Type. } l \mapsto \sigma \quad (33)$$

$$l \mapsto \tau \vdash \triangleright \forall \sigma : \text{Type. } l \mapsto \sigma \Rightarrow \triangleright \text{teq}(\tau, \sigma) \quad (34)$$

Proof. We assume $(k, \Psi) \Vdash l \mapsto \tau$, that is, $\Psi(l) = \tau$. Let (k', Ψ') be a world such that $(k, \Psi) R(k', \Psi')$.

By definition 31, we have $\text{dom} \Psi \subset \text{dom} \Psi'$, and therefore, for some type σ , we have $\Psi'(l) = \sigma$, that is, $(k', \Psi') \Vdash l \mapsto \sigma$. This proves the first assertion.

Now, let σ be a type such that $(k', \Psi') \Vdash l \mapsto \sigma$, that is, $\Psi'(l) = \sigma$. By definition 31, $\sigma = \lfloor \tau \rfloor_{k'}$. Hence, by lemma 7.1, $(k', \Psi') \Vdash \triangleright \text{teq}(\tau, \sigma)$ as required to prove the second assertion. \square

In particular, if a storable value u can be stored at location l now, it can still be stored at this location in the future:

$$l \mapsto \tau \wedge \triangleright (u \circ \tau) \vdash \triangleright \forall \sigma : \text{Type. } l \mapsto \sigma \Rightarrow \triangleright (u \circ \sigma). \quad (35)$$

Another direct corollary of the lemma is that reference types are necessary.

We now show that the expected store operations are permitted. For instance, if the location l has type $\text{ref } \tau$, the storable-value u has type τ and the current memory m is valid, we should be able to store u at location l ; that is, the memory $m[l := u]$ should be valid in some world one step after the current world.

Lemma 7.3 (Unchanged memory). *If we have*

$$(k + 1, \Psi) \Vdash \text{validmem}(m) \quad (36)$$

then there exists a store typing Ψ' such that

$$(k + 1, \Psi) R(k, \Psi') \quad (37)$$

$$(k, \Psi') \Vdash \text{validmem}(m) \quad (38)$$

$$\text{dom } \Psi' = \text{dom } \Psi \quad (39)$$

Proof. We take $\Psi' = \lfloor \Psi \rfloor_k$. We clearly have assertions 37 and 39. We now prove assertion 38 using definition 27. Let l be a location and τ be a type such that $(k, \Psi') \Vdash l \mapsto \tau$. By definition 26, as $\text{dom } \Psi' = \text{dom } \Psi$, there exists τ' such that $(k+1, \Psi) \Vdash l \mapsto \tau'$. Furthermore, by assertion 36, $(k+1, \Psi) \Vdash \triangleright (m(l) \circ \tau')$. Hence, by assertion 35, $(k, \Psi') \Vdash \triangleright (m(l) \circ \tau)$ as wanted. \square

Lemma 7.4 (Memory update). *If the type τ is necessary and*

$$(k+1, \Psi) \Vdash \text{validmem}(m) \quad (40)$$

$$(k+1, \Psi) \Vdash (l \circ \text{ref } \tau) \quad (41)$$

$$(k+1, \Psi) \Vdash (u \circ \tau) \quad (42)$$

then there exists a store typing Ψ' such that

$$(k+1, \Psi) R(k, \Psi') \quad (43)$$

$$(k, \Psi') \Vdash \text{validmem}(m[l := u]) \quad (44)$$

$$\text{dom } \Psi' = \text{dom } \Psi \quad (45)$$

Proof. We take $\Psi' = \lfloor \Psi \rfloor_k$ as for the previous lemma 7.3. This ensures that $(k+1, \Psi) R(k, \Psi')$ and $\text{dom } \Psi' = \text{dom } \Psi$. We now prove assertion 44 using definition 27. Let l' be a location and τ' be a type such that $(k, \Psi') \Vdash l' \mapsto \tau'$. If $l' \neq l$, we can conclude as in lemma 7.3. So, we suppose that $l' = l$. From assumption 41 and by definition 32, there exists a type τ'' such that $(k+1, \Psi) \Vdash l \mapsto \tau''$ and $(k+1, \Psi) \Vdash \triangleright \text{teq}(\tau, \tau'')$. From assumption 42, using the fact that τ is necessary, we therefore get $(k+1, \Psi) \Vdash \triangleright (u \circ \tau'')$. Hence, by assertion 35, $(k, \Psi') \Vdash \triangleright (u \circ \tau')$ as wanted. \square

Lemma 7.5 (Memory allocation). *If the type τ is necessary, and*

$$(k+1, \Psi) \Vdash \text{validmem}(m) \quad (46)$$

$$(k+1, \Psi) \Vdash (u \circ \tau) \quad (47)$$

$$l \notin \text{dom } \Psi \quad (48)$$

then there exists a store typing Ψ' such that

$$(k+1, \Psi) R(k, \Psi') \quad (49)$$

$$(k, \Psi') \Vdash \text{validmem}(m[l := u]) \quad (50)$$

$$(k, \Psi') \Vdash (l \circ \text{ref } \tau) \quad (51)$$

$$\text{dom } \Psi' = \text{dom } \Psi \cup \{l\} \quad (52)$$

Proof. We take $\Psi' = \lfloor \Psi \rfloor_k \cup \{l \mapsto \lfloor \tau \rfloor_k\}$. This is well-defined as $l \notin \text{dom } \Psi$. Clearly, we have $(k+1, \Psi) R(k, \Psi')$ and $\text{dom } \Psi' = \text{dom } \Psi \cup \{l\}$. We now prove assertion 50 using definition 27. Let l' be a location and τ' be a type such that $(k, \Psi') \Vdash l' \mapsto \tau'$. If $l' \neq l$, we can conclude as in lemma 7.3. So, we suppose that $l' = l$. From assumption 47, using the fact that τ is necessary, we get $(k, \Psi') \Vdash \triangleright (u \circ \tau)$. Using lemma 7.1, this gives $(k, \Psi') \Vdash \triangleright (u \circ \lfloor \tau \rfloor_k)$. Thus, using the fact that $\tau' = \Psi'(l) = \lfloor \tau \rfloor_k$, we have $(k, \Psi') \Vdash \triangleright (u \circ \tau')$ as wanted. Finally, by lemma 7.1,

$(k, \Psi') \Vdash \triangleright \text{teq}(\tau, \lfloor \tau \rfloor_k)$, and by definition of Ψ' , $(k, \Psi') \Vdash l \mapsto \lfloor \tau \rfloor_k$. Therefore, we have assertion 51 by definition 32. \square

Lemma 7.6 (Memory access). *If*

$$w \Vdash (l \circ \text{ref } \tau) \quad (53)$$

$$w \Vdash \text{validmem}(m) \quad (54)$$

$$w R w' \quad (55)$$

then $w' \Vdash (m(l) \circ \tau)$ (56)

Proof. From assumption 53 there exists a type τ' such that $w \Vdash l \mapsto \tau'$ and $w \Vdash \triangleright \text{teq}(\tau, \tau')$. Using assumption 54 we get $w \Vdash \triangleright m(l) \circ \tau'$ and then $w \Vdash \triangleright m(l) \circ \tau$. Finally, by definition of \triangleright , $w' \Vdash (m(l) \circ \tau)$, as wanted. \square

With these lemmas, we can prove soundness of rules for type-checking store instructions for initializing new references or updating old ones (Appendix A).

8. Data structures

To describe the memory layout of data structures and of the stored program, we need address arithmetic on locations (Loc). In this paper we will identify Loc with the integers. To reason about address arithmetic on Loc, we have arithmetic operators in the type system:

$$\text{plus}(i, j, k) \stackrel{\text{prim}}{=} \{(w, v) \mid i + j = k\} \quad (57)$$

$$\text{greater}(i, j) \stackrel{\text{prim}}{=} \{(w, v) \mid i > j\} \quad (58)$$

From these we can construct the offset type constructor, such that $l \circ \text{offset}(n, \tau)$ iff $l + n \circ \tau$. This is useful for record fields and program labels.

$$\text{offset}(n, \tau) \stackrel{\text{syn}}{=} \exists l: \text{Loc. just } l \wedge \exists l': \text{Loc. plus}(l, n, l') \wedge l' \circ \tau \quad (59)$$

Now records are just intersections of fields, and disjoint union types such as lists are easy to construct.

$$\text{pair}(\sigma, \tau) \stackrel{\text{syn}}{=} \text{offset}(0, \text{ref } \sigma) \wedge \text{offset}(1, \text{ref } \tau) \quad (60)$$

$$\text{list } \tau \stackrel{\text{syn}}{=} \text{rec } \alpha. \text{ just } 0 \vee \exists l: \text{Loc. just } l \wedge \text{greater}(l, 0) \wedge \text{pair}(\tau, \alpha) \quad (61)$$

That is, nil of a list is zero; the cons cell of a list must be at an address l , such that $l > 0$, and also must be a pair of an element τ and a list α .

Immutable references. An immutable reference to type τ can be modeled as a mutable reference to a singleton l , such that (in this world) l has type τ :

$$\text{iref } \tau \stackrel{\text{syn}}{=} \exists l: \text{Loc. ref}(\text{just } l) \wedge l \circ \tau \quad (62)$$

Henry Ford sold Model Ts in any color the customer wanted, so long as it was black. An immutable ref is a mutable ref into which the customer can store any value he wants, so long as it's the value that's already there.

9. Safe states in safe worlds

We wish to reason about function values and function application. A function of type $\tau \rightarrow \sigma$ (or, on a von Neumann machine, a code pointer of type $\text{codeptr}(\tau)$) is a value that is *safe* to apply to arguments of type τ . To synthesize function or code-pointer types, we must first model *safety* in our type system.

Our notion of “safe for k steps” is independent of our choice of calculus. It requires only that the calculus comes equipped with a small-step relation $\mapsto \subset \text{State} \times \text{State}$.

$$\text{safe}_0(s) \stackrel{\text{def}}{=} \text{True} \quad (63)$$

$$\text{safe}_{k+1}(s) \stackrel{\text{def}}{=} (\exists s'. s \mapsto s') \wedge \forall s'. s \mapsto s' \Rightarrow \text{safe}_k s' \quad (64)$$

For a calculus of mutable references, we assume that $\text{State} = V \times \text{Mem}$, that is, the state (v, m) on a von Neumann machine has register-bank v and memory m , or in a λ -calculus has expression v and store m .

The type $\text{safemem}(m)$ expresses the set of configurations (w, v) in which it is safe to run forward from state (v, m) . Note that the world w contains an “expiration date” n telling how many steps forward we may attempt to take. After n steps we declare victory, that is, we do not care whether the small-step relation might get stuck *after* the n th step.

$$\text{safemem}(m) \stackrel{\text{prim}}{=} \{(n, \Psi), v \mid \text{safe}_n(v, m)\} \quad (65)$$

$$\text{safe} \stackrel{\text{syn}}{=} \forall m: \text{Mem}. \text{validmem}(m) \Rightarrow \text{safemem}(m) \quad (66)$$

The judgment $w \Vdash v : \text{safe}$ is an important one. A memory m is “in” a world w whenever $w \Vdash \text{validmem}(m)$. A value v is safe in w when, for all m “in” w , the state (v, m) is safe for all the steps remaining in w . It is on the primitive constructor safe that we will build (in Section 11) the notion of codeptr , that is, continuation-pointer.

10. λ -calculus

The reader may skip this section without loss of comprehension of later sections. The main point of departure here is Ahmed’s thesis [Ahm04], especially Section 3.3.5, which gives constructions used in defining the type constructor for arrow types.

Our type system works well to classify the terms of systems of λ -calculus, such as the one given below. In fact, all of our previous definitions apply without modification; recall that for λ -calculus we have $SV = V$ and $\vec{u} \stackrel{\text{def}}{=} u$.

$$e ::= () \mid x \mid l \mid \lambda x. e \mid (e_1 e_2) \mid \text{new } e \mid !e \mid e_1 := e_2 \quad (67)$$

$$u, v ::= () \mid l \mid \lambda x. e \quad (68)$$

This is a λ -calculus extended with mutable references. It’s a contribution of our system that even in the presence of such

$\text{T-UNIT} \quad \frac{}{\Gamma \Vdash () : \text{unit}}$	$\text{T-VAR} \quad \frac{\Gamma(x) = \tau}{\Gamma \Vdash x : \tau}$	$\text{T-ABS} \quad \frac{\Gamma[x := \sigma] \Vdash e : \tau}{\Gamma \Vdash (\lambda x. e) : \sigma \rightarrow \tau}$
$\text{T-APP} \quad \frac{\Gamma \Vdash e_1 : \sigma \rightarrow \tau \quad \Gamma \Vdash e_2 : \sigma}{\Gamma \Vdash (e_1 e_2) : \tau}$	$\text{T-NEW} \quad \frac{\Gamma \Vdash e : \tau}{\Gamma \Vdash (\text{new } e) : \text{ref } \tau}$	
$\text{T-DEREF} \quad \frac{\Gamma \Vdash e : \text{ref } \tau}{\Gamma \Vdash (!e) : \tau}$	$\text{T-ASSIGN} \quad \frac{\Gamma \Vdash e_1 : \text{ref } \tau \quad \Gamma \Vdash e_2 : \tau}{\Gamma \Vdash (e_1 := e_2) : \text{unit}}$	

Figure 3. Typing rules specific to λ -calculus

references, impredicative type abstraction and type application (and so forth) are coercions with no operational significance; in particular, terms of quantified type have no explicit introduction or elimination forms, such as “pack” and “unpack”. Section 13.1 discusses how we achieve this.

Because we define our types semantically, we decline to give an inductive definition for type expressions. Moreover we treat λ -calculus à la Curry and omit type decorations from our terms. We use the standard small-step call-by-value operational semantics, and omit the rules for same.

Definitions of the primitive type constructors, and those for quantified, recursive, and reference types remain unchanged. It therefore remains to give semantics to the unit and arrow types; the typing judgments for closed terms, substitutions, and open terms; and to prove the typing lemmas (which would be rules in a syntactic progress-and-preservation proof) for abstraction and application.

First the typing judgment for closed terms.

$$e :_{(k, \Psi)} \tau \stackrel{\text{def}}{=} \forall j < k, m, m', e'. \quad (69)$$

$$(k, \Psi) \Vdash \text{validmem}(m)$$

$$\Rightarrow (m, e) \mapsto^j (m', e') \wedge \text{irred}(m', e')$$

$$\Rightarrow \exists \Psi'. (k, \Psi) R^=(k - j, \Psi')$$

$$\wedge (k - j, \Psi') \Vdash \text{validmem}(m')$$

$$\wedge (k - j, \Psi') \Vdash e' : \tau$$

The state (m, e) may step to (m', e') which is irreducible (either a value or *stuck*); but since $_ \Vdash e' : \tau$ only when e' is a value, then a well-typed term e cannot get stuck within k steps.

There are two new type constructors of interest,

$$\text{unit} \stackrel{\text{prim}}{=} \{(w, ()) \mid \text{True}\}, \text{ and} \quad (70)$$

$$\tau \rightarrow \sigma \stackrel{\text{syn}}{=} \triangleright (\tau \rightarrow' \sigma), \text{ where} \quad (71)$$

$$\tau \rightarrow' \sigma \stackrel{\text{prim}}{=} \{(w, \lambda x. e) \mid \forall w, v. (w, v) \in \tau \Rightarrow e[v/x] :_w \sigma\}.$$

Finally we define the general typing judgment \models in terms of ground substitutions γ :

$$\begin{aligned} \gamma :_w \Gamma &\stackrel{\text{def}}{=} \forall x \in \text{dom } \Gamma. \gamma(x) :_w \Gamma(x) \\ \Gamma \models e : \tau &\stackrel{\text{def}}{=} \forall w, \gamma. \gamma :_w \Gamma \Rightarrow \gamma(e) :_w \tau \end{aligned}$$

The expected typing rules (Figure 3) can be proved as lemmas from the model.

One of our major semantic innovations is the decomposition of function types into primitive notions such as safety. Equation 69 does not show this, because so far we have worked out this more modular decomposition only in the von Neumann setting. Therefore the next section explains this decomposition in terms of code pointers.

11. Machine instructions

Sections 2–9 apply to models of computation ranging from λ -calculus to von Neumann machines. Now we will discuss von Neumann machines specifically. A location of type $\text{codeptr}(\tau)$ is the address of a continuation to which it's safe to jump, providing one passes an argument of type τ . That is, if $l \circ \text{codeptr}(\tau)$ then one can safely set the program counter to l , as long as the registers satisfy τ .

$$\begin{aligned} \text{codeptr}(\tau) &\stackrel{\text{syn}}{=} \\ \exists l : \text{Loc}. \text{just } l \wedge \triangleright!(\text{slot}(\text{pc}, \text{just } l) \wedge \tau \Rightarrow \text{safe}) &\quad (72) \end{aligned}$$

At any time strictly in the future, whenever it is the case that the program counter = l and the (other) registers satisfy τ , then the state must be safe.⁶

A machine-instruction ι is a relation on machine states, $\iota \subset \text{State} \times \text{State}$. If the program counter in state s_1 points to a memory location containing a number n that codes for ι , then the machine will step to a state s_2 such that $\iota(s_1, s_2)$.

An example of such an instruction is $r3 \leftarrow m(r2 + 9)$, which we can write as

$$\begin{aligned} \text{LOAD}(3, 2, 9) &= \{((v_1, m_1), (v_2, m_2)) \mid \\ &v_2(\text{pc}) = v_1(\text{pc}) + 1 \wedge v_2(3) = m_1(v_1(2) + 9) \wedge \\ &\forall i \notin \{3, \text{pc}\}. v_2(i) = v_1(i) \wedge m_1 = m_2\} \end{aligned}$$

A number n codes for an instruction ι just when

$$\begin{aligned} n \text{ encodes } \iota &\stackrel{\text{def}}{=} \forall v_1, m_1, v_2, m_2. \\ m_1(v_1(\text{pc})) &= n \Rightarrow \\ (v_1, m_1) \rightsquigarrow (v_2, m_2) &\Leftrightarrow \iota((v_1, m_1), (v_2, m_2)) \end{aligned} \quad (73)$$

Programs. A von Neumann program p is a sequence of integers $p(0), \dots, p(n-1)$ such that each $p(i)$ encodes some instruction ι_i . For example, perhaps on some machine

7320 encodes $\text{LOAD}(3, 2, 0)$, we might have a program $p = (7320, 4231, 6007)$:

<i>Machine inst.</i>	ι	<i>Pseudocode</i>	(74)
7320	$\text{LOAD}(3, 2, 0)$	$r3 \leftarrow m(r2 + 0)$	
4231	$\text{ADDIMM}(2, 3, 1)$	$r2 \leftarrow r3 + 1$	
6007	$\text{JUMP}(7)$	$\text{pc} \leftarrow r7$	

Program-code types. For any program p , define

$$\Delta_p \stackrel{\text{syn}}{=} \bigwedge_{i \in \text{dom } p} \text{offset}(i, \text{ref}(\text{just } p(i))) \quad (75)$$

(The domain of p is finite; the $\bigwedge_{i \in \text{dom } p}$ is meta-notation for an expression with $|\text{dom } p|$ conjuncts.) For example, our sample program p (equation 74) has,

$$\begin{aligned} \Delta_0 &= \text{offset}(0, \text{ref}(\text{just } 7320)) \\ \Delta_1 &= \text{offset}(1, \text{ref}(\text{just } 4231)) \\ \Delta_2 &= \text{offset}(2, \text{ref}(\text{just } 6007)) \\ \Delta_p &= \Delta_0 \wedge \Delta_1 \wedge \Delta_2 \end{aligned}$$

Type-checking the example program p . Suppose that on entry to location l , register $r2$ is a ref int and that register $r7$ is a $\text{codeptr}(\text{slot}(2, \text{int}))$, that is, a continuation that requires an integer in register 2. Then the program dereferences the ref, adds 1, and continues by jumping to the address in $r7$. We can explain this as $l : \Gamma$ where

$$\Gamma = \text{codeptr}(\text{slot}(2, \text{ref int}) \wedge \text{slot}(7, \text{codeptr}(\text{slot}(2, \text{int}))))$$

Now we wish to prove that, in any world where this program is loaded at address l , then l is also a continuation as described by Γ . We write this claim as,

$$\Delta_p \vdash \Gamma$$

Tan [TA06, Tan05] gives a compositional logic for control flow in which such statements can be proved. He proves soundness of this logic using the indexed model. Our primitives can express Tan's semantics in a clean and simple way.

The form of an instruction-typing rule is $\Delta, \triangleright \Gamma \vdash \Gamma'$ where Δ describes an instruction in memory at location l , Γ describes the type of the continuation at location $l + 1$, and the conclusion Γ' claims the type of the continuation at l . Consider a $\text{LOAD}(i, j, k)$ at location l . A sound but naive typing rule is,

$$\frac{n \text{ encodes } \text{LOAD}(i, j, k) \quad i \neq \text{pc} \quad \tau \text{ necessary}}{\text{ref}(\text{just } n), \triangleright \text{offset}(1, \text{codeptr}(\text{slot}(i, \tau))) \vdash \text{codeptr}(\text{slot}(j, \text{offset}(k, \text{ref } \tau)))} \quad (76)$$

That is, if the load instruction is at address l , and it's safe to continue at $l + 1$ provided that register i has a value of type τ , then it's also safe to continue at l provided that $r_j + k$ has a value of type ref τ .

This is insufficiently general: it neglects the fact that all the registers other than i are not changed by the load

⁶ Why "strictly in the future"—why not right now? The codeptr does not require \square instead of \triangleright for two reasons: on the one hand, then it would not be contractive; and on the other hand, executing a jump to location l takes at least one step, so we don't need the stronger \square condition. Contractiveness is a useful property for codeptr , since it allows the recursive type rec α . $\text{codeptr}(\alpha)$.

instruction. Suppose we have a type ϕ that characterizes the type of several registers but leaves i unrestricted. We say, “ i is not in the domain of ϕ ,” that is,

$$\text{notindom}(i, \phi) \stackrel{\text{def}}{=} \forall w, v, u. (w, v) \in \phi \Rightarrow (w, v[i:=u]) \in \phi \quad (77)$$

A stronger load rule is,

$$\frac{\begin{array}{l} n \text{ encodes } \text{LOAD}(i, j, k) \quad i \neq \text{pc} \\ \tau \text{ necessary} \quad \phi \text{ necessary} \\ \text{notindom}(i, \phi) \quad \text{notindom}(\text{pc}, \phi) \end{array}}{\text{ref}(\text{just } n), \triangleright \text{offset}(1, \text{codeptr}(\phi \wedge \text{slot}(i, \tau))) \vdash \text{codeptr}(\phi \wedge \text{slot}(j, \text{offset}(k, \text{ref } \tau)))} \quad (78)$$

With respect to our sample program p (equation 74) we can write the invariants for locations $l + 0$ through $l + 2$ as,

$$\begin{aligned} \Gamma_0 &= \text{offset}(0, \text{codeptr}(\phi \wedge \text{slot}(2, \text{ref } \text{int}))) \\ \Gamma_1 &= \text{offset}(1, \text{codeptr}(\phi \wedge \text{slot}(3, \text{int}))) \\ \Gamma_2 &= \text{offset}(2, \text{codeptr}(\phi \wedge \text{slot}(2, \text{int}))) \\ \Gamma_p &= \Gamma_0 \wedge \Gamma_1 \wedge \Gamma_2 \\ &\text{where } \phi = \text{slot}(7, \text{codeptr}(\text{slot}(2, \text{int}))) \end{aligned} \quad (79)$$

Then we can use rules for LOAD, ADDIMM, and JUMP to show,

$$\begin{array}{l} \Delta_0, \triangleright \Gamma_1 \vdash \Gamma_0 \\ \Delta_1, \triangleright \Gamma_2 \vdash \Gamma_1 \\ \Delta_2 \quad \vdash \Gamma_2 \end{array} \quad (80)$$

11.1 Whole-program typing

The example program is not recursive and does not loop, so in our example, the proof for each instruction i needs only $\triangleright \Gamma_{i+1}$. However, our type system is applicable to programs with loops and recursion. Suppose i is a conditional branch to location j , then we need,

$$\Delta_i, \triangleright (\Gamma_j \wedge \Gamma_{i+1}) \vdash \Gamma_i.$$

The general case is $\Delta_i, \triangleright \Gamma_p \vdash \Gamma_i$.

Tan’s calculus for control flow [TA06, Tan05] manages such proofs in a modular way. Given proofs of this form for each i , we can combine them using rules for type intersection into,

$$\Delta_p, \triangleright \Gamma_p \vdash \Gamma_p \quad (81)$$

So, in general, we have a program p and its associated Δ_p . An oracle (perhaps a type-preserving compiler) provides a hypothesis Γ_p giving a codeptr type for each address. An automatic type-checker uses our rules—or syntax-directed specializations of our rules, as described by Wu [Wu05]—to prove $\Delta_p, \triangleright \Gamma_p \vdash \Gamma_p$. All ref types are necessary, and necessity is preserved by offset and intersection, therefore Δ_p is necessary. Using the generalized Löb rule (Corollary 4.5), we deduce that

$$\Delta_p \vdash \Gamma_p \quad (82)$$

This means, “in any world where the program code described by Δ_p is loaded, the loop-invariants Γ_p hold.”

The initial entry point. Let us assume that the “beginning of the program” is at offset 0 in a complete program, and suppose the hypothesis for offset 0 is codeptr(top), meaning that it is (claimed) safe to jump to the beginning of the program regardless of the contents of the registers. Then by the rules of type intersection we have $\Gamma_p \vdash \text{codeptr}(\text{top})$. Consequently (by equation 82) we have

$$\Delta_p \vdash \text{codeptr}(\text{top}). \quad (83)$$

Now we will show that this implies the program is safe.

12. Soundness

A von Neumann program is *loaded at l* in a machine-state (r, m) if $r(\text{pc}) = l$ and for each $i \in \text{dom } p$, $m(l+i) = p(i)$.

Lemma 12.1. *If program p is loaded at l in machine-state (v, m) , then for any natural number k one can construct a world (k, Ψ) such that,*

$$(k, \Psi) \Vdash \text{validmem}(m) \quad (84)$$

$$(k, \Psi) \Vdash v : \text{slot}(\text{pc}, \Delta_p) \quad (85)$$

Proof. The type constructors used in Δ_p are very simple: just intersection, ref, and just. Therefore we can use a simple construction of Ψ :

$$\Psi = \{(l+i) \mapsto \lfloor \text{just } p(i) \rfloor_k \mid i \in \text{dom } p\}$$

We first check equation 84, that is, by equation 27, if $(k, \Psi) \Vdash l' \mapsto \tau$, then $(k, \Psi) \Vdash \triangleright (m(l') \circ \tau)$. By definition 26, when the hypothesis holds, there exists $i \in \text{dom } p$ such that $l' = l + i$ and $\tau = \lfloor \text{just } (p(i)) \rfloor_k$. The conclusion then becomes $(k, \Psi) \Vdash \triangleright (m(l+i) \circ \lfloor \text{just } (p(i)) \rfloor_k)$, that is, $(k, \Psi) \Vdash \triangleright (\text{just } (p(i)) \Rightarrow \lfloor \text{just } (p(i)) \rfloor_k)$. This is a consequence of lemma 7.1.

We now consider equation 85. Let $i \in \text{dom } p$ and v' be a value. We have

$$\begin{aligned} (k, \Psi) \Vdash l+i \mapsto \lfloor \text{just } p(i) \rfloor_k \\ (k, \Psi) \Vdash \triangleright \text{teq}(\text{just } p(i), \lfloor \text{just } p(i) \rfloor_k). \end{aligned}$$

by definition 26 and lemma 7.1. Therefore, by equation 32, $(k, \Psi) \Vdash l+i \circ \text{ref}(\text{just } p(i))$. Hence, by equation 59, $(k, \Psi) \Vdash l \circ \text{offset}(i, \text{ref}(\text{just } p(i)))$ for all $i \in \text{dom } p$. By intersection, we get $(k, \Psi) \Vdash l \circ \Delta_p$. We conclude by equation 15. \square

Theorem 12.2 (Soundness). *If program p is loaded at l in machine-state (v, m) , and $\Delta_p \vdash \text{codeptr}(\text{top})$, then state (v, m) is safe (for any number of steps k).*

Proof. By Lemma 12.1, construct the world $w = (k+1, \Psi)$. By Lemma 7.3, there exists a world $w' = (k, \Psi')$ such that wRw' and $w' \Vdash \text{validmem}(m)$. We have $w \Vdash v : \text{slot}(\text{pc}, \Delta_p)$, and thus $w \Vdash v : \text{slot}(\text{pc}, \text{codeptr}(\text{top}))$. As the program is loaded at l , we also have $w' \Vdash v : \text{slot}(\text{pc}, \text{just } l)$. By equation 72 we get $w' \Vdash v : \text{safe}$, so by equation 66 we have $w' \Vdash v : \text{safemem}(m)$. By equation 65, this means $\text{safe}_k(v, m)$. \square

We have established that our syntactic system is sound in the sense that every well typed program is safe. This follows from the fact that each individual typing rule is semantically valid in our model (from the definition of \vdash , equation 16). Even after we have proved Theorem 12.2 we may add more typing rules, as long as we derive them semantically.

13. Related work

13.1 Impredicative polymorphism with references

Here’s a type isomorphism that one would expect to hold, in a system with continuation types and polymorphism:

$$\text{codeptr}(\exists\alpha. F(\alpha)) \cong \forall\alpha. \text{codeptr}(F(\alpha)). \quad (86)$$

Let’s consider in particular a special case of the above,

$$\begin{aligned} \text{codeptr}(\exists\alpha:\text{Type}. \alpha \times A \times \text{codeptr}(B)) \\ \cong \forall\alpha:\text{Type}. \text{codeptr}(\alpha \times A \times \text{codeptr}(B)) \end{aligned} \quad (87)$$

where the left and right-hand sides are (de Morgan) dual ways of writing the CPS- and closure-converted type of a function $f : A \rightarrow B$. We’re happy to say that the isomorphisms 86 and 87 hold in our system. We’re also happy to say that in our system, none of the operations on values of quantified type—existential pack and unpack, type abstraction and application—require taking a step in the operational semantics. (As types are not manifest at runtime in a statically checked system, such steps would be effective no-ops, and therefore a nuisance.) To have both of these properties simultaneously is a technical advance over previous work [Ahm04], which presented two alternative models of quantifiers, each of which possessed one property but not the other.

We believe Ahmed’s set-theoretic model can be adjusted (by simplifying the model of quantified types) to eliminate this nuisance. However, it is not obvious how to directly formalize such a repaired model—or equally, the mathematics of Section 6—in a logic without dependent types, in particular Church’s higher-order logic (HOL). To our knowledge, no previous attempt to express equation 28 in HOL could do better than the solutionless 25. To avoid circularity in the HOL representation, Ahmed chooses the range of (HOL-encoded) Ψ to be a set of reified syntactic type expressions. This causes the failure of equation 87—in particular, the left side ($\text{codeptr}(\exists\alpha. \dots)$) has no useful semantics, while the right side ($\forall\alpha. \text{codeptr}(\dots)$) has the expected semantics. The Princeton FPCC compiler [CWF03] works around this problem by carefully using only the $\forall\text{codeptr}$ formulation of function-closures.

13.2 Approximation modality

The approximation modality is due to Nakano, who developed the idea in two papers. The first paper [Nak00] introduces a type system with recursive types and an approximation modality \bullet (to which our \triangleright is directly analogous),

which for example allows the derivation

$$\vdash \mathbf{Y} : (\bullet X \rightarrow X) \rightarrow X,$$

where \mathbf{Y} is the fixed-point combinator. The approximation modality allows, even in the presence of recursive types, a strongly normalizing calculus that may be sensibly used as a logic. The second paper [Nak01] generalizes the first, and makes precise the relationship to the intuitionistic version of the Gödel-Löb logic of provability.

Our work builds on and departs from Nakano’s in two principal ways. First, we use the approximation modality in calculi that are by design capable of nontermination. Second, our main contribution is to show how the approximation modality is not just a modality for recursion, but also a modality for mutable references, and in general a modality for realizing inductive properties of small-step programs.

13.3 XCAP

Shao et al.’s Certified Assembly Programming [YHS03, NS06] is based directly on the calculus of inductive constructions. Like our system, CAP/XCAP supports separate verification of code modules, permits impredicative quantification and (with extensions) mutable references, and is expressive enough to specify invariants for assembly code. They achieve impredicativity by means of a syntactic specification of validity rules for impredicative propositions. This method does not permit elimination rules for the quantifiers, but they show how to work around the lack of elimination rules when existentials appear at top level in Hoare rules, which is good enough for a TAL.

14. Conclusion

Previous work broke functions and data structures into simpler constructs—that is, existential quantifiers, address arithmetic, unions, intersections, all applied to first-order continuations and immutable or mutable references. But here we have shown (equations 32, 62, 72) that first-order continuations (codeptr) and references (iref , ref) can themselves be broken down into simpler concepts, each of which is a *type* operator: safe models program safety, \triangleright models induction over future steps, \square models type preservation, $!$ constructs logical statements about the current world. This decomposition allows simpler and finer reasoning about the interaction of instructions with types, and allows the easier construction of new types.

Our semantic model is powerful: it is the first model of references with impredicativity in which all the expected identities hold. It is versatile: it can express in a natural way all the types needed for typed assembly languages. It is modular: the vast majority of our definitions and lemmas apply equally to λ -calculus, von Neumann machines, and other useful abstract machines. It is compositional: there are a few simple primitives, from which everything else is constructed. It is sound: we have machine-checked proofs

ID $\tau \vdash \tau$	CUT $\frac{\tau_1 \vdash \tau_2 \quad \tau_2 \vdash \tau_3}{\tau_1 \vdash \tau_3}$	INTERSECTION-R $\frac{\sigma \vdash \tau_1 \quad \sigma \vdash \tau_2}{\sigma \vdash \tau_1 \wedge \tau_2}$	FORALL-R $\frac{\tau \vdash \sigma \quad x \text{ not free in } \tau}{\tau \vdash \forall x:A.\sigma}$	FORALL-L $\frac{a \in A}{\forall x:A.\tau \vdash \tau[a/x]}$
	INTERSECTION-L-1 $\tau_1 \wedge \tau_2 \vdash \tau_1$	INTERSECTION-L-2 $\tau_1 \wedge \tau_2 \vdash \tau_2$	EXISTS-L $\frac{\sigma \vdash \tau \quad x \text{ not free in } \tau}{\exists x:A.\sigma \vdash \tau}$	EXISTS-R $\frac{a \in A}{\tau[a/x] \vdash \exists x:A.\tau}$
UNION-L $\frac{\tau_1 \vdash \sigma \quad \tau_2 \vdash \sigma}{\tau_1 \vee \tau_2 \vdash \sigma}$	UNION-R-1 $\tau_1 \vdash \tau_1 \vee \tau_2$	UNION-R-2 $\tau_2 \vdash \tau_1 \vee \tau_2$	LATER-FORALL $\forall x:A. \triangleright \tau \vdash \triangleright \forall x:A.\tau$	
TOP-R $\tau \vdash \text{top}$	BOT-L $\text{bot} \vdash \tau$	IMPLIES-I $\frac{\sigma \wedge \tau_1 \vdash \tau_2}{\sigma \vdash \tau_1 \Rightarrow \tau_2}$	IMPLIES-E $\frac{\sigma \vdash \tau_1 \Rightarrow \tau_2}{\sigma \wedge \tau_1 \vdash \tau_2}$	
LATER-LIFT $\frac{\sigma \vdash \tau}{\triangleright \sigma \vdash \triangleright \tau}$	LATER-REP $\triangleright \tau \vdash \triangleright \triangleright \tau$	LATER-FIX $\frac{\triangleright \tau \vdash \tau}{\text{top} \vdash \tau}$	EVERYWHERE-LIFT $\frac{\sigma \vdash \tau}{!\sigma \vdash !\tau}$	EVERYWHERE-REP $!\tau \vdash !!\tau$
LATER-INTERSECTION $\triangleright \sigma \wedge \triangleright \tau \vdash \triangleright (\sigma \wedge \tau)$	LATER-TOP $\text{top} \vdash \triangleright \text{top}$		EVERYWHERE-L $!\tau \vdash \tau$	EVERYWHERE-INTERSECTION $!\sigma \wedge !\tau \vdash !(\sigma \wedge \tau)$

Figure 4. Gödel-Löb logic

JUST-UNIQUE $\text{just } l \wedge \text{just } l' \vdash \tau[l] \Rightarrow \tau[l']$	LATER-JUST $\text{just } l \vdash \triangleright \text{just } l$
LATER-JUST-2 $\text{just } l \Rightarrow \triangleright \tau \vdash \triangleright (\text{just } l \Rightarrow \tau)$	SOMEWHERE-JUST $\text{top} \vdash ? \text{just } l$
SLOT-LIFT $\frac{\sigma \vdash \tau}{\text{slot}(j, \sigma) \vdash \text{slot}(j, \tau)}$	LATER-SLOT $\text{slot}(j, \triangleright \tau) \vdash \triangleright \text{slot}(j, \tau)$

Figure 5. Just and slot

for the von Neumann version. It is operational: the statement of the soundness theorem guarantees directly an untyped notion of safety. Finally, it is *semantic*: the modal operator \triangleright “later” explains the essence of approximation over the remaining steps of computation, in a way that allows us to better understand old types and better invent new ones.

Future work. Using this concise and easily representable model, we expect that machine-checked soundness proofs of usable TALs will be substantially smaller than those attempted to date [Cra03, App01].

Acknowledgments

We thank James Leifer, Amal Ahmed, and Zhong Shao for comments on earlier drafts of the paper.

Figure 6. Quantification	
EVERYWHERE-L $!\tau \vdash \tau$	EVERYWHERE-INTERSECTION $!\sigma \wedge !\tau \vdash !(\sigma \wedge \tau)$
EVERYWHERE-TOP $\text{top} \vdash !\text{top}$	LATER-EVERYWHERE $!\triangleright \tau \vdash \triangleright !\tau$
EVERYWHERE-LATER $\triangleright !\tau \vdash \triangleright !\tau$	EVERYWHERE-IMPLIES $!(\sigma \Rightarrow \tau) \vdash ?\sigma \Rightarrow ?\tau$
SOMEWHERE-L $\frac{\sigma \vdash !\tau}{?\sigma \vdash !\tau}$	EVERYWHERE-SOMEWHERE $? \tau \vdash !? \tau$
LATER-SOMEWHERE $? \triangleright \tau \vdash \triangleright ? \tau$	

Figure 7. Quantification over values in a world

References

- [AAV02] Amal Ahmed, Andrew W. Appel, and Roberto Virga. A stratified semantics of general references embeddable in higher-order logic. In *17th Annual IEEE Symp. on Logic in Computer Science*, pp. 75–86, June 2002.
- [AF00] Andrew W. Appel and Amy P. Felty. A semantic model of types and machine instructions for proof-carrying code. In *POPL '00: 27th ACM SIGPLAN-SIGACT Symp. on Principles of Programming Languages*, pp. 243–253, Jan. 2000.
- [Ahm04] Amal J. Ahmed. *Semantics of Types for Mutable State*. PhD thesis, Princeton University, Princeton, NJ, Nov. 2004. Tech Report TR-713-04.
- [AM01] Andrew W. Appel and David McAllester. An indexed model of recursive types for foundational proof-carrying code. *ACM Trans. on Programming Languages and Systems*, 23(5):657–683, Sept. 2001.

- [App01] Andrew W. Appel. Foundational proof-carrying code. In *Symp. on Logic in Computer Science (LICS '01)*, pp. 247–258. IEEE, 2001.
- [B⁺98] Bruno Barras et al. The Coq Proof Assistant reference manual. Technical report, INRIA, 1998.
- [BW83] Michael Barr and Charles Wells. *Toposes, Triples and Theories*. Grundlehren der math. Wissenschaften. Springer Verlag, 1983. Reprint 12 in *Theory and Applications of Category*, //www.emis.de/journals/TAC/.
- [Cra03] Karl Cray. Toward a foundational typed assembly language. In *POPL '03: 30th ACM Symp. on Principles of Programming Languages*, pp. 198–212, 2003.
- [CWF03] Juan Chen, Dinghao Wu, Andrew W. Appel, and Hai Fang. A provably sound TAL for back-end optimization. In *PLDI '03: Proc. 2003 ACM SIGPLAN Conference on Programming Language Design and Implementation*, pp. 208–219, June 2003.
- [Kri63] Saul A. Kripke. Semantical considerations on modal logic. In *Proceedings of a Colloquium: Modal and Many Valued Logics*, vol. 16, pp. 83–94, 1963.
- [MM91] John C. Mitchell and Eugenio Moggi. Kripke-style models for typed lambda calculus. *Annals of Pure and Applied Logic*, 1991.
- [MPS86] David MacQueen, Gordon Plotkin, and Ravi Sethi. An ideal model for recursive polymorphic types. *Information and Computation*, 71(1/2):95–130, 1986.
- [MV05] Paul-André Melliès and Jérôme Vouillon. Recursive polymorphic types and parametricity in an operational framework. In *IEEE Symp. on Logic in Computer Science (LICS '05)*, 2005.
- [Nak00] Hiroshi Nakano. A modality for recursion. In *LICS '00: 15th Annual IEEE Symp. on Logic in Computer Science*, pp. 255–266. IEEE Computer Society Press, 2000.
- [Nak01] Hiroshi Nakano. Fixed-point logic with the approximation modality and its Kripke completeness. In *Theoretical Aspects of Computer Software*, vol. 2215 of *LNCS*, pp. 165–182. Springer, 2001.
- [NS06] Zhaozhong Ni and Zhong Shao. Certified assembly programming with embedded code pointers. In *33rd ACM Symp. on Principles of Programming Languages*, pp. 320–333. ACM Press, Jan. 2006.
- [PM93] Christine Paulin-Mohring. Inductive definitions in the system Coq; rules and properties. In M. Bezem and J. F. Groote, editors, *Proc. International Conference on Typed Lambda Calculi and Applications*, vol. 664, pp. 328–345. Springer Verlag Lecture Notes in Computer Science, 1993.
- [TA06] Gang Tan and Andrew W. Appel. A compositional logic for control flow. In *7th International Conference on Verification, Model Checking and Abstract Interpretation (VMCAI'06)*, pp. 80–94, Jan. 2006.
- [Tan05] Gang Tan. *A Compositional Logic for Control Flow and its Application to Foundational Proof-Carrying Code*. PhD thesis, Princeton University, Princeton, NJ, Aug. 2005. Tech Report CS-TR-731-05.

[Wu05] Dinghao Wu. *Interfacing Compilers, Proof Checkers, and Proofs for Foundational Proof-Carrying Code*. PhD thesis, Princeton University, Aug. 2005.

[YHS03] Dachuan Yu, Nadeem A. Hamid, and Zhong Shao. Building certified libraries for PCC: Dynamic storage allocation. In *Proc. 2003 European Symp. on Programming (ESOP'03)*, April 2003.

A. Allocation and initialization

Our type system can express the invariants of data structures that manage a region calculus (a typed malloc/free system, or a garbage collector). Such type systems are not new [AF00, CWF03, Ahm04, Wu05]. Here we briefly illustrate how our new type system can express the necessary invariants just using the primitives we have already defined.

The simplest possible such data structure is an integer variable that defines the boundary between allocated and unallocated memory. We define $\text{boundary}(l)$ to mean that every address $\geq l$ is unallocated, i.e. not in the domain of Ψ :

$$\text{boundary } l \stackrel{\text{syn}}{=} \forall l': \text{Loc}. \forall \tau: \text{Type}. l' \mapsto \tau \Rightarrow \text{greater}(l, l') \quad (88)$$

Now we give a rule for type-checking $\text{STORE}(i, 0, j)$ (i.e., $m(r(i) + 0) \leftarrow r(j)$) when it serves to initialize a new reference:

$$\frac{\begin{array}{l} n \text{ encodes } \text{STORE}(i, 0, j) \\ \tau \text{ necessary} \quad \phi \text{ necessary} \quad \text{notindom}(\text{pc}, \phi) \\ P = \phi \wedge \exists l: \text{Loc}. \text{boundary } l \wedge \text{slot}(i, \text{just } l) \wedge \text{slot}(j, \tau) \\ Q = \phi \wedge \exists l: \text{Loc}. \text{boundary}(l + 1) \wedge \text{slot}(i, \text{just } l \wedge \text{ref } \tau) \end{array}}{\text{ref}(\text{just } n), \triangleright \text{offset}(1, \text{codeptr}(Q)) \vdash \text{codeptr}(P)} \quad (89)$$

Before the instruction, register i points at the boundary. The instruction does not modify i , but does modify the boundary, in that the domain of Ψ in the new world has moved.

The same store instruction, when it is used for the different purpose of updating an existing ref, has this typing rule:

$$\frac{\begin{array}{l} n \text{ encodes } \text{STORE}(i, 0, j) \quad B = \text{boundary } l \\ \tau \text{ necessary} \quad \phi \text{ necessary} \quad \text{notindom}(\text{pc}, \phi) \\ P = B \wedge \phi \wedge \text{slot}(i, \text{ref } \tau) \wedge \text{slot}(j, \tau) \end{array}}{\text{ref}(\text{just } n), \triangleright \text{offset}(1, \text{codeptr}(B \wedge \phi)) \vdash \text{codeptr}(P)} \quad (90)$$

Of course, both of these rules for STORE must be proved as derived lemmas from the definition in our model. By proving every typing rule as a derived lemma, we thus prove soundness of the rules.