

# Denotational Semantics for Gradual Typing in Synthetic Guarded Domain Theory

ERIC GIOVANNINI and MAX S. NEW

We develop a denotational semantics for a gradually typed language with effects that is adequate and proves the graduality theorem. The denotational semantics is constructed using *synthetic guarded domain theory* working in a type theory with a later modality and clock quantification. This provides a remarkably simple presentation of the semantics, where gradual types are interpreted as ordinary types in our ambient type theory equipped with an ordinary preorder structure to model the error ordering. This avoids the complexities of classical domain-theoretic models (New and Licata) or logical relations models using explicit step-indexing (New and Ahmed). In particular, we avoid a major technical complexity of New and Ahmed that requires two logical relations to prove the graduality theorem.

By working synthetically we can treat the domains in which gradual types are interpreted as if they were ordinary sets. This allows us to give a “naïve” presentation of gradual typing where each gradual type is modeled as a well-behaved subset of the universal domain used to model the dynamic type, and type precision is modeled as simply a subset relation.

## ACM Reference Format:

Eric Giovannini and Max S. New. 2018. Denotational Semantics for Gradual Typing in Synthetic Guarded Domain Theory. In *Woodstock '18: ACM Symposium on Neural Gaze Detection, June 03–05, 2018, Woodstock, NY*. ACM, New York, NY, USA, 12 pages. <https://doi.org/10.1145/1122445.1122456>

## 1 INTRODUCTION

### 1.1 Gradual Typing and Graduality

Gradual typing allows a language to have both statically-typed and dynamically-typed terms; the statically-typed terms are type checked at compile time, while type checking for the dynamically-typed terms is deferred to runtime.

Gradually-typed languages should satisfy two intuitive properties. First, the interaction between the static and dynamic components of the codebase should be safe – i.e., should preserve the guarantees made by the static types. In other words, in the static portions of the codebase, type soundness must be preserved. Second, gradual languages should support the smooth migration from dynamic typing to static typing, in that the programmer can initially leave off the typing annotations and provide them later without altering the meaning of the program.

Formally speaking, gradually typed languages should satisfy the *dynamic gradual guarantee*, originally defined by Siek, Vitousek, Cimini, and Boyland [8]. This property is also referred to as *graduality*, by analogy with parametricity. Intuitively, graduality says that in going from a dynamic to static style should not introduce changes in the meaning of the program. More specifically, making the types more precise by adding annotations will either result in the same behavior as the less precise program, or result in a type error.

---

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. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](mailto:permissions@acm.org).

Woodstock '18, June 03–05, 2018, Woodstock, NY

© 2018 Association for Computing Machinery.

ACM ISBN 978-1-4503-XXXX-X/18/06...\$15.00

<https://doi.org/10.1145/1122445.1122456>

## 1.2 Current Approaches

Current approaches to proving graduality include the methods of Abstracting Gradual Typing [4] and the formal tools of the Gradualier [3]. These allow the language developer to start with a statically typed language and derive a gradually typed language that satisfies the gradual guarantee. The downside is that not all gradually typed languages can be derived from these frameworks, and moreover, in both approaches the semantics is derived from the static type system as opposed to the alternative in which the semantics determines the type checking. Without a clear semantic interpretation of type dynamism, it becomes difficult to extend these techniques to new language features such as polymorphism.

New and Ahmed [7] have developed a semantic approach to specifying type dynamism in terms of *embedding-projection pairs*, which allows for a particularly elegant formulation of the gradual guarantee. Moreover, their axiomatic account of program equivalence allows for type-based reasoning about program equivalences. In this approach, a logical relation is constructed and shown to be sound with respect to the notion of observational approximation that specifies when one program is more precise than another. The downside of this approach is that each new language requires a different logical relation to prove graduality. Furthermore, the logical relations tend to be quite complicated due to a technical requirement known as *step-indexing*. As a result, developments using this approach tend to require vast effort, with the corresponding technical reports having 50+ pages of proofs.

An alternative approach, which we investigate in this paper, is provided by *synthetic guarded domain theory*. The techniques of synthetic guarded domain theory allow us to internalize the step-index reasoning normally required in logical relations proofs of graduality, ultimately allowing us to specify the logical relation in a manner that looks nearly identical to a typical, non-step-indexed logical relation.

In this paper, we report on work we have done to mechanize graduality proofs using SGDT techniques in Agda. Our hope in this work is that by mechanizing a graduality proof in a reusable way, we will provide a framework for other language designers to use to more easily and conveniently prove that their languages satisfy graduality.

## 1.3 Contributions

Our main contribution is a framework in Guarded Cubical Agda for proving graduality of a cast calculus. To demonstrate the feasibility and utility of our approach, we have used the framework to prove graduality for the simply-typed gradual lambda calculus. Along the way, we have developed an “intensional” theory of graduality that is of independent interest.

## 1.4 Proving Graduality in SGDT

**TODO:** This section should probably be moved to after the relevant background has been introduced.

In this paper, we will utilize SGDT techniques to prove graduality for a particularly simple gradually-typed cast calculus, the gradually-typed lambda calculus. This is just the usual simply-typed lambda calculus with a dynamic type  $\top$  such that  $A \sqsubseteq \top$  for all types  $A$ , as well as upcasts and downcasts between any types  $A$  and  $B$  such that  $A \sqsubseteq B$ . The complete definition will be provided in Section 3.

Our main theorem is the following:

**THEOREM 1.1 (GRADUALITY).** *If  $\cdot \vdash M \sqsubseteq N : \text{Nat}$ , then*

- (1) *If  $N = \top$ , then  $M = \top$*
- (2) *If  $N = 'n$ , then  $M = \top$  or  $M = 'n$*
- (3) *If  $M = V$ , then  $N = V$*

We also should be able to show that  $\mathbb{U}$ ,  $\text{zro}$ , and  $\text{suc } N$  are not equal.

Our first step toward proving graduality is to formulate an *intensional* gradual lambda calculus, which we call  $\text{Int-}\lambda C$ , in which the computation steps taken by a term are made explicit. The “normal” gradual lambda calculus for which we want to prove graduality will be called the *extensional* gradual lambda calculus, denoted  $\text{Ext-}\lambda C$ . We will define an erasure function  $\lfloor \cdot \rfloor : \text{Int-}\lambda C \rightarrow \text{Ext-}\lambda C$  which takes a program in the intensional lambda calculus and “forgets” the syntactic information about the steps to produce a term in the extensional calculus.

Every term  $M_e$  in  $\text{Ext-}\lambda C$  will have a corresponding program  $M_i$  in  $\text{Int-}\lambda C$  such that  $\lfloor M_i \rfloor = M_e$ . Moreover, we will show that if  $M_e \sqsubseteq_e M'_e$  in the extensional theory, then there exists terms  $M_i$  and  $M'_i$  such that  $\lfloor M_i \rfloor = M_e$ ,  $\lfloor M'_i \rfloor = M'_e$  and  $M_i \sqsubseteq_i M'_i$  in the intensional theory.

We formulate and prove an analogous graduality theorem for the intensional lambda calculus. We define an interpretation of the intensional lambda calculus into a model in which we prove various results. Using the observation above, given  $M_e \sqsubseteq_e M'_e : \text{Nat}$ , we can find intensional programs  $M_i$  and  $M'_i$  that erase to them and are such that  $M_i \sqsubseteq_i M'_i$ . We will then apply the intensional graduality theorem to  $M_i$  and  $M'_i$ , and translate the result back to  $M_e$  and  $M'_e$ .

## 1.5 Overview of Remainder of Paper

In Section 2, we provide technical background on gradually typed languages and on synthetic guarded domain theory. In Section 3, we introduce the gradually-typed cast calculus for which we will prove graduality. Important here are the notions of syntactic type precision and term precision. We introduce both the *extensional* gradual lambda calculus ( $\text{Ext-}\lambda C$ ) and the *intensional* gradual lambda calculus ( $\text{Int-}\lambda C$ ). In Section 4, we define several fundamental constructions internal to SGDT that will be needed when we give a denotational semantics to our intensional lambda calculus. This includes the notion of Predomains as well as the concept of EP-Pairs. In Section 5, we define the denotational semantics for the intensional gradually-typed lambda calculus using the domain theoretic constructions in the previous section. In Section 6, we outline in more detail the proof of graduality for the extensional gradual lambda calculus, which will make use of prove properties we prove about the intensional gradual lambda calculus. In Section 7, we discuss the benefits and drawbacks to our approach in comparison to the traditional step-indexing approach, as well as possibilities for future work.

## 2 TECHNICAL BACKGROUND

### 2.1 Gradual Typing

In a gradually-typed language, the mixing of static and dynamic code is seamless, in that the dynamically typed parts are checked at runtime. This type checking occurs at the elimination forms of the language (e.g., pattern matching, field reference, etc.). Gradual languages are generally elaborated to a *cast calculus*, in which the dynamic type checking is made explicit through the insertion of *type casts*.

In a cast calculus, there is a relation  $\sqsubseteq$  on types such that  $A \sqsubseteq B$  means that  $A$  is a *more precise* type than  $B$ . There a dynamic type  $?$  with the property that  $A \sqsubseteq ?$  for all  $A$ . If  $A \sqsubseteq B$ , a term  $M$  of type  $A$  may be *upcast* to  $B$ , written  $\langle B \prec A \rangle M$ , and a term  $N$  of type  $B$  may be *downcast* to  $A$ , written  $\langle A \preceq B \rangle N$ . Upcasts always succeed, while downcasts may fail at runtime. We also have a notion of *syntactic term precision*. If  $A \sqsubseteq B$ , and  $M$  and  $N$  are terms of type  $A$  and  $B$  respectively, we write  $M \sqsubseteq N : A \sqsubseteq B$  to mean that  $M$  is more precise than  $N$ , i.e.,  $M$  and  $N$  behave the same except that  $M$  may error more.

## 2.2 Difficulties in Prior Semantics

In this work, we compare our approach to proving graduality to the approach introduced by New and Ahmed [7] which constructs a step-indexed logical relations model and shows that this model is sound with respect to their notion of contextual error approximation.

Because the dynamic type is modeled as a non-well-founded recursive type, their logical relation needs to be parameterized by natural numbers to restore well-foundedness. This technique is known as a *step-indexed logical relation*. Reasoning about step-indexed logical relations can be tedious and error-prone, and there are some very subtle aspects that must be taken into account in the proofs. Figure ?? shows an example of a step-indexed logical relation for the gradually-typed lambda calculus.

In particular, the prior approach of New and Ahmed requires two separate logical relations for terms, one in which the steps of the left-hand term are counted, and another in which the steps of the right-hand term are counted. Then two terms  $M$  and  $N$  are related in the “combined” logical relation if they are related in both of the one-sided logical relations. Having two separate logical relations complicates the statement of the lemmas used to prove graduality, because any statement that involves a term stepping needs to take into account whether we are counting steps on the left or the right. Some of the differences can be abstracted over, but difficulties arise for properties as fundamental and seemingly straightforward as transitivity.

Specifically, for transitivity, we would like to say that if  $M$  is related to  $N$  at index  $i$  and  $N$  is related to  $P$  at index  $i$ , then  $M$  is related to  $P$  at  $i$ . But this does not actually hold: we require that one of the two pairs of terms be related “at infinity”, i.e., that they are related at  $i$  for all  $i \in \mathbb{N}$ . Which pair is required to satisfy this depends on which logical relation we are considering, (i.e., is it counting steps on the left or on the right), and so any argument that uses transitivity needs to consider two cases, one where  $M$  and  $N$  must be shown to be related for all  $i$ , and another where  $N$  and  $P$  must be related for all  $i$ . This may not even be possible to show in some scenarios!

## 2.3 Synthetic Guarded Domain Theory

One way to avoid the tedious reasoning associated with step-indexing is to work axiomatically inside of a logical system that can reason about non-well-founded recursive constructions while abstracting away the specific details of step-indexing required if we were working analytically. The system that proves useful for this purpose is called *synthetic guarded domain theory*, or SGDT for short. We provide a brief overview here, but more details can be found in [2].

SGDT offers a synthetic approach to domain theory that allows for guarded recursion to be expressed syntactically via a type constructor  $\triangleright : \text{Type} \rightarrow \text{Type}$  (pronounced “later”). The use of a modality to express guarded recursion was introduced by Nakano [6]. Given a type  $A$ , the type  $\triangleright A$  represents an element of type  $A$  that is available one time step later. There is an operator  $\text{next} : A \rightarrow \triangleright A$  that “delays” an element available now to make it available later. We will use a tilde to denote a term of type  $\triangleright A$ , e.g.,  $\tilde{M}$ .

There is a *guarded fixpoint* operator

$$\text{fix} : \forall T, (\triangleright T \rightarrow T) \rightarrow T.$$

That is, to construct a term of type  $T$ , it suffices to assume that we have access to such a term “later” and use that to help us build a term “now”. This operator satisfies the axiom that  $\text{fix}f = f(\text{next}(\text{fix}f))$ . In particular, this axiom applies to propositions  $P : \text{Prop}$ ; proving a statement in this manner is known as Löb-induction.

The operators  $\triangleright$ ,  $\text{next}$ , and  $\text{fix}$  described above can be indexed by objects called *clocks*. A clock serves as a reference relative to which steps are counted. For instance, given a clock  $k$  and type  $T$ , the type

$$X_0 \xleftarrow{r_0^X} X_1 \xleftarrow{r_1^X} X_2 \xleftarrow{r_2^X} X_3 \xleftarrow{r_3^X} \dots$$

Fig. 1. An example of an object in the topos of trees.

$$\begin{array}{ccccccc} X_0 & \xleftarrow{r_0^X} & X_1 & \xleftarrow{r_1^X} & X_2 & \xleftarrow{r_2^X} & X_3 \xleftarrow{r_3^X} \dots \\ f_0 \downarrow & & f_1 \downarrow & & f_2 \downarrow & & f_3 \downarrow \\ Y_0 & \xleftarrow{r_0^Y} & Y_1 & \xleftarrow{r_1^Y} & Y_2 & \xleftarrow{r_2^Y} & Y_3 \xleftarrow{r_3^Y} \dots \end{array}$$

Fig. 2. An example of a morphism in the topos of trees.

$\triangleright^k T$  represents a value of type  $T$  one unit of time in the future according to clock  $k$ . If we only ever had one clock, then we would not need to bother defining this notion. However, the notion of *clock quantification* is crucial for encoding coinductive types using guarded recursion, an idea first introduced by Atkey and McBride [1].

**2.3.1 Ticked Cubical Type Theory.** In Ticked Cubical Type Theory [?], there is an additional sort called *ticks*. Given a clock  $k$ , a tick  $t : \text{tick}k$  serves as evidence that one unit of time has passed according to the clock  $k$ . The type  $\triangleright A$  is represented as a function from ticks of a clock  $k$  to  $A$ . The type  $A$  is allowed to depend on  $t$ , in which case we write  $\triangleright_t^k A$  to emphasize the dependence.

The rules for tick abstraction and application are similar to those of dependent  $\Pi$  types. In particular, if we have a term  $M$  of type  $\triangleright^k A$ , and we have available in the context a tick  $t' : \text{tick}k$ , then we can apply the tick to  $M$  to get a term  $M[t'] : A[t'/t]$ . We will also write tick application as  $M_t$ . Conversely, if in a context  $\Gamma, t : \text{tick}k$  we have that  $M$  has type  $A$ , then in context  $\Gamma$  we have  $\lambda t.M$  has type  $\triangleright A$ .

The statements in this paper have been formalized in a variant of Agda called Guarded Cubical Agda [?], an implementation of Clocked Cubical Type Theory.


**2.3.2 The Topos of Trees Model.** The topos of trees model provides a useful intuition for reasoning in SGGT [2]. This section presupposes knowledge of category theory and can be safely skipped without disrupting the continuity.

The topos of trees  $\mathcal{S}$  is the presheaf category  $\text{Set}^{\omega^o}$ . We assume a universe  $\mathcal{U}$  of types, with encodings for operations such as sum types (written as  $\hat{+}$ ). There is also an operator  $\widehat{\triangleright} : \triangleright \mathcal{U} \rightarrow \mathcal{U}$  such that  $\text{El}(\widehat{\triangleright}(\text{next}A)) = \triangleright \text{El}(A)$ , where  $\text{El}$  is the type corresponding to the code  $A$ .

An object  $X$  is a family  $\{X_i\}$  of sets indexed by natural numbers, along with restriction maps  $r_i^X : X_{i+1} \rightarrow X_i$  (see Figure 1). These should be thought of as “sets changing over time”, where  $X_i$  is the view of the set if we have  $i + 1$  time steps to reason about it. We can also think of an ongoing computation, with  $X_i$  representing the potential results of the computation after it has run for  $i + 1$  steps.

A morphism from  $\{X_i\}$  to  $\{Y_i\}$  is a family of functions  $f_i : X_i \rightarrow Y_i$  that commute with the restriction maps in the obvious way, that is,  $f_i \circ r_i^X = r_i^Y \circ f_{i+1}$  (see Figure 2).

The type operator  $\triangleright$  is defined on an object  $X$  by  $(\triangleright X)_0 = 1$  and  $(\triangleright X)_{i+1} = X_i$ . The restriction maps are given by  $r_0^{\triangleright X} = !$ , where  $!$  is the unique map into 1, and  $r_{i+1}^{\triangleright X} = r_i^X$ . The morphism  $\text{next}^X : X \rightarrow \triangleright X$  is defined pointwise by  $\text{next}_0^X = !$ , and  $\text{next}_{i+1}^X = r_i^X$ . It is easily checked that this satisfies the commutativity conditions required of a morphism in  $\mathcal{S}$ . Given a morphism  $f : \triangleright X \rightarrow X$ ,

Fig. 3. The guarded fixpoint of  $f$ .

i.e., a family of functions  $f_i: (\triangleright X)_i \rightarrow X_i$  that commute with the restrictions in the appropriate way, we define  $\text{fix}(f): 1 \rightarrow X$  pointwise by  $\text{fix}(f)_i = f_i \circ \dots \circ f_0$ . This can be visualized as a diagram in the category of sets as shown in Figure 3.

### 3 GTLC

Here we describe the syntax and typing for the gradually-typed lambda calculus. We also give the rules for syntactic type and term precision. We define two separate calculi: the normal gradually-typed lambda calculus, which we call the extensional lambda calculus (Ext- $\lambda C$ ), as well as an *intensional* lambda calculus (Int- $\lambda C$ ) whose syntax makes explicit the steps taken by a program.

#### 3.1 Syntax

Types  $A, B := \text{Nat}, ?, (A \Rightarrow B)$

Terms  $M, N := \mathbf{U}_A, \text{zro}, \text{suc } M, (\lambda x.M), (MN), (\langle B \prec_r A \rangle M), (\langle A \leftarrow B \rangle M)$

Contexts  $\Gamma := \cdot, (\Gamma, x : A)$

The typing rules are as expected, with a cast between  $A$  to  $B$  allowed only when  $A \sqsubseteq B$ .

$$\begin{array}{c}
 \frac{}{\Gamma \vdash \mathbf{U}_A : A} \quad \frac{}{\Gamma \vdash \text{zro} : \text{Nat}} \quad \frac{\Gamma \vdash M : \text{Nat}}{\Gamma \vdash \text{suc } M : \text{Nat}} \quad \frac{\Gamma, x : A \vdash M : B}{\Gamma \vdash \lambda x.M : A \Rightarrow B} \\
 \\
 \frac{\Gamma \vdash M : A \Rightarrow B \quad \Gamma \vdash N : A}{\Gamma \vdash MN : B} \quad \frac{A \sqsubseteq B \quad \Gamma \vdash M : A}{\Gamma \vdash \langle B \prec_r A \rangle M : B} \quad \frac{A \sqsubseteq B \quad \Gamma \vdash M : B}{\Gamma \vdash \langle A \leftarrow B \rangle M : A}
 \end{array}$$

The equational theory is also as expected, with  $\beta$  and  $\eta$  laws for function type.

#### 3.2 Type Precision

The rules for type precision are as follows:

$$\begin{array}{c}
 \frac{}{? \sqsubseteq ?} \quad \frac{}{\text{Nat} \sqsubseteq \text{Nat}}^{\text{NAT}} \quad \frac{}{\text{Nat} \sqsubseteq ?}^{\text{Inj}_{\text{NAT}}} \quad \frac{A_i \sqsubseteq B_i \quad A_o \sqsubseteq B_o}{(A_i \Rightarrow A_o) \sqsubseteq (B_i \Rightarrow B_o)} \Rightarrow \\
 \\
 \frac{(A_i \rightarrow A_o) \sqsubseteq (? \Rightarrow ?)}{(A_i \rightarrow A_o) \sqsubseteq ?}^{\text{Inj}_{\Rightarrow}}
 \end{array}$$

Note that as a consequence of this presentation of the type precision rules, we have that if  $A \sqsubseteq B$ , there is a unique precision derivation that witnesses this. As in previous work, we go a step farther and make these derivations first-class objects, known as *type precision derivations*. Specifically, for

every  $A \sqsubseteq B$ , we have a derivation  $d : A \sqsubseteq B$  that is constructed using the rules above. For instance, there is a derivation  $? : ? \sqsubseteq ?$ , and a derivation  $\text{Nat} : \text{Nat} \sqsubseteq \text{Nat}$ , and if  $d_i : A_i \sqsubseteq B_i$  and  $d_o : A_o \sqsubseteq B_o$ , then there is a derivation  $d_i \Rightarrow d_o : (A_i \Rightarrow A_o) \sqsubseteq (B_i \Rightarrow B_o)$ . Likewise for the remaining two rules. The benefit to making these derivations explicit in the syntax is that we can perform induction over them. Note also that for any type  $A$ , we use  $A$  to denote the reflexivity derivation that  $A \sqsubseteq A$ , i.e.,  $A : A \sqsubseteq A$ . Finally, observe that for type precision derivations  $d : A \sqsubseteq B$  and  $d' : B \sqsubseteq C$ , we can define their composition  $d' \circ d : A \sqsubseteq C$ . This will be used in the statement of transitivity of the term precision relation.

### 3.3 Term Precision

We allow for a *heterogeneous* term precision judgment on terms  $M$  of type  $A$  and  $N$  of type  $B$ , provided that  $A \sqsubseteq B$  holds.

In order to deal with open terms, we will need the notion of a type precision *context*, which we denote  $\Gamma^\sqsubseteq$ . This is similar to a normal context but instead of mapping variables to types, it maps variables  $x$  to related types  $A \sqsubseteq B$ , where  $x$  has type  $A$  in the left-hand term and  $B$  in the right-hand term. We may also write  $x : d$  where  $d : A \sqsubseteq B$  to indicate this. Another way of thinking of type precision contexts is as a zipped pair of contexts  $\Gamma^l, \Gamma^r$  with the same domain such that  $\Gamma^l(x) \sqsubseteq \Gamma^r(x)$  for each  $x$  in the domain.

The rules for term precision come in two forms. We first have the *congruence* rules, one for each term constructor. These assert that the term constructors respect term precision. The congruence rules are as follows:

$$\begin{array}{c}
\frac{\Gamma \vdash M : A}{\Gamma^\sqsubseteq \vdash \mathcal{U}_A \sqsubseteq_e M : A} \mathcal{U} \qquad \frac{d : A \sqsubseteq B \quad \Gamma^\sqsubseteq(x) = (A, B)}{\Gamma^\sqsubseteq \vdash x \sqsubseteq_e x : d} \text{VAR} \qquad \frac{}{\Gamma^\sqsubseteq \vdash \text{zro} \sqsubseteq_e \text{zro} : \text{Nat}} \text{ZRO} \\
\\
\frac{\Gamma^\sqsubseteq \vdash M \sqsubseteq_e N : \text{Nat}}{\Gamma^\sqsubseteq \vdash \text{suc } M \sqsubseteq_e \text{suc } N : \text{Nat}} \text{SUC} \\
\\
\frac{d_i : A_i \sqsubseteq B_i \quad d_o : A_o \sqsubseteq B_o \quad \Gamma^\sqsubseteq, x : d_i \vdash M \sqsubseteq_e N : d_o}{\Gamma^\sqsubseteq \vdash \lambda x. M \sqsubseteq_e \lambda x. N : (d_i \Rightarrow d_o)} \text{LAMBDA} \\
\\
\frac{\Gamma^\sqsubseteq \vdash M \sqsubseteq_e M' : (d_i \Rightarrow d_o) \quad \Gamma^\sqsubseteq \vdash N \sqsubseteq_e N' : d_i}{\Gamma^\sqsubseteq \vdash M N \sqsubseteq_e M' N' : d_o} \text{APP}
\end{array}$$

We then have additional equational axioms, including transitivity,  $\beta$  and  $\eta$  laws, and rules characterizing upcasts as least upper bounds, and downcasts as greatest lower bounds.

We write  $M \sqsupseteq N$  to mean that both  $M \sqsubseteq N$  and  $N \sqsubseteq M$ .

$$\begin{array}{c}
\frac{\Gamma^{\sqsubseteq} \vdash M \sqsubseteq_e N : d \quad \Gamma^{\sqsubseteq} \vdash N \sqsubseteq_e P : d'}{\Gamma^{\sqsubseteq} \vdash M \sqsubseteq_e P : d' \circ d} \text{TRANSITIVITY} \quad \frac{\Gamma, x : A_i \vdash M : A_o \quad \Gamma \vdash V : A_i}{\Gamma^{\sqsubseteq} \vdash (\lambda x.M) V \sqsubseteq_e M[V/x] : A_o} \beta \\
\\
\frac{\Gamma \vdash V : A_i \Rightarrow A_o}{\Gamma^{\sqsubseteq} \vdash \lambda x.(V x) \sqsubseteq_e V : A_i \Rightarrow A_o} \eta \quad \frac{d : A \sqsubseteq B \quad \Gamma \vdash M : A}{\Gamma^{\sqsubseteq} \vdash M \sqsubseteq_e \langle B \hookleftarrow A \rangle M : d} \text{UpR} \\
\\
\frac{d : A \sqsubseteq B \quad \Gamma^{\sqsubseteq} \vdash M \sqsubseteq_e N : d}{\Gamma^{\sqsubseteq} \vdash \langle B \hookleftarrow A \rangle M \sqsubseteq_e N : B} \text{UpL} \quad \frac{d : A \sqsubseteq B \quad \Gamma \vdash M : B}{\Gamma^{\sqsubseteq} \vdash \langle A \hookleftarrow B \rangle M \sqsubseteq_e M : d} \text{DnL} \\
\\
\frac{d : A \sqsubseteq B \quad \Gamma^{\sqsubseteq} \vdash M \sqsubseteq_e N : d}{\Gamma^{\sqsubseteq} \vdash M \sqsubseteq_e \langle A \hookleftarrow B \rangle N : A} \text{DnR}
\end{array}$$

The rules UpR, UpL, DnL, and DnR were introduced in [?] as a means of cleanly axiomatizing the intended behavior of casts in a way that doesn't depend on the specific constructs of the language. Intuitively, rule UpR says that the upcast of  $M$  is an upper bound for  $M$  in that  $M$  may error more, and UpL says that the upcast is the *least* such upper bound, in that it errors more than any other upper bound for  $M$ . Conversely, DnL says that the downcast of  $M$  is a lower bound, and DnR says that it is the *greatest* lower bound.

### 3.4 The Intensional Lambda Calculus

#### TODO: Subject to change!

Now that we have described the syntax along with the type and term precision judgments for Ext- $\lambda C$ , we can now do the same for Int- $\lambda C$ . One key difference between the two calculi is that we define Int- $\lambda C$  using the constructs available to us in the language of synthetic guarded domain theory, e.g., we use the  $\triangleright$  operator. Whereas when we defined the syntax of the extensional lambda calculus we were working in the category of sets, when we define the syntax of the intensional lambda calculus we will be working in the topos of trees.

More specifically, in Int- $\lambda C$ , we not only have normal terms, but also terms available “later”, which we denote by  $\tilde{M}$ . We have a term constructor  $\theta_s$  which takes a later-term and turns it into a term available now. The typing and precision rules for  $\theta_s$  involve the  $\triangleright$  operator, as shown below. Observe that  $\theta_s$  is a syntactic analogue to the  $\theta$  constructor of the lifting monad that we will define in the section on domain theoretic constructions (Section 4), but it is important to note that  $\theta_s(\tilde{M})$  is an actual term in Int- $\lambda C$ , whereas the  $\theta$  constructor is a purely semantic construction. These will be connected when we discuss the interpretation of Int- $\lambda C$  into the semantic model.

To better understand this situation, note that  $\lambda x.(\cdot)$  can be viewed as a function (at the level of the ambient type theory) from terms of type  $B$  under context  $\Gamma, x : A$  to terms of type  $A \Rightarrow B$  under context  $\Gamma$ . Similarly, we can view  $\theta_s(\cdot)$  as a function (in the ambient type theory, which is synthetic guarded domain theory) taking terms  $\tilde{M}$  of type  $A$  available *later* to terms of type  $A$  available *now*.

Notice that the notion of a “term available later” does not need to be part of the syntax of the intensional lambda calculus, because this can be expressed in the ambient theory. Similarly, we do not need a syntactic form to delay a term, because we can simply use next.

$$\text{Terms } M, N := \mathcal{U}_A, \dots, \theta_s(\tilde{M})$$



$$\frac{\triangleright_t (\Gamma \vdash M_t : A)}{\Gamma \vdash \theta_s M : A}$$

$$\frac{\triangleright_t (\Gamma^\Xi \vdash M_t \sqsubseteq_i N_t : d)}{\Gamma^\Xi \vdash \theta_s M \sqsubseteq_i \theta_s N : d}$$

Recall that  $\triangleright_t$  is a dependent form of  $\triangleright$  where the argument is allowed to mention  $t$ . In particular, here we apply the tick  $t$  to the later-terms  $M$  and  $N$  to get “now”-terms  $M_t$  and  $N_t$ .

Formally speaking, the term precision relation must be defined as a guarded fixpoint, i.e., we assume that the function is defined later, and use it to construct a definition that is defined “now”. This involves applying a tick to the later-function to shift it to a now-function. Indeed, this is what we do in the formal Agda development, but in this paper, we will elide these issues as they are not relevant to the main ideas.

## 4 DOMAIN-THEORETIC CONSTRUCTIONS

In this section, we discuss the fundamental objects of the model into which we will embed the intensional lambda calculus and inequational theory. It is important to remember that the constructions in this section are entirely independent of the syntax described in the previous section; the notions defined here exist in their own right as purely mathematical constructs. In the next section, we will link the syntax and semantics via an interpretation function.

### 4.1 The Lift Monad

When thinking about how to model intensional gradually-typed programs, we must consider their possible behaviors. On the one hand, we have *failure*: a program may fail at run-time because of a type error. In addition to this, a program may “think”, i.e., take a step of computation. If a program thinks forever, then it never returns a value, so we can think of the idea of thinking as a way of intensionally modelling *partiality*.

With this in mind, we can describe a semantic object that models these behaviors: a monad for embedding computations that has cases for failure and “thinking”. Previous work has studied such a construct in the setting of the latter only, called the lift monad [5]; here, we add the additional effect of failure.

For a type  $A$ , we define the *lift monad with failure*  $L_{\mathcal{U}}A$ , which we will just call the *lift monad*, as the following datatype:

$$\begin{aligned} L_{\mathcal{U}}A &:= \\ \eta &: A \rightarrow L_{\mathcal{U}}A \\ \mathcal{U} &: L_{\mathcal{U}}A \\ \theta &: \triangleright (L_{\mathcal{U}}A) \rightarrow L_{\mathcal{U}}A \end{aligned}$$

Formally, the lift monad  $L_{\mathcal{U}}A$  is defined as the solution to the guarded recursive type equation

$$L_{\mathcal{U}}A \cong A + 1 + \triangleright L_{\mathcal{U}}A.$$

An element of  $L_{\mathcal{U}}A$  should be viewed as a computation that can either (1) return a value (via  $\eta$ ), (2) raise an error and stop (via  $\mathcal{U}$ ), or (3) think for a step (via  $\theta$ ). Notice there is a computation  $\text{fix}\theta$  of type  $L_{\mathcal{U}}A$ . This represents a computation that thinks forever and never returns a value.

Since we claimed that  $L_{\mathcal{U}}A$  is a monad, we need to define the monadic operations and show that they respect the monadic laws. The return is just  $\eta$ , and extend is defined via guarded recursion by cases on the input. Verifying that the monadic laws hold requires Löb-induction and is straightforward.

The lift monad has the following universal property. Let  $f$  be a function from  $A$  to  $B$ , where  $B$  is a  $\triangleright$ -algebra, i.e., there is  $\theta_B: \triangleright B \rightarrow B$ . Further suppose that  $B$  is also an “error-algebra”, that is, an algebra of the constant functor  $1: \text{Type} \rightarrow \text{Type}$  mapping all types to Unit. This latter statement amounts to saying that there is a map  $\text{Unit} \rightarrow B$ , so  $B$  has a distinguished “error element”  $\mathcal{U}_B: B$ .

Then there is a unique homomorphism of algebras  $f': L_{\mathcal{U}}A \rightarrow B$  such that  $f' \circ \eta = f$ . The function  $f'(l)$  is defined via guarded fixpoint by cases on  $l$ . In the  $\mathcal{U}$  case, we simply return  $\mathcal{U}_B$ . In the  $\theta(\tilde{l})$  case, we will return

$$\theta_B(\lambda t. (f'_t \tilde{l}_t)).$$

Recalling that  $f'$  is a guarded fixpoint, it is available “later” and by applying the tick we get a function we can apply “now”; for the argument, we apply the tick to  $\tilde{l}$  to get a term of type  $L_{\mathcal{U}}A$ .

**4.1.1 Model-Theoretic Description.** We can describe the lift monad in the topos of trees model as follows.

## 4.2 Predomains

The next important construction is that of a *predomain*. A predomain is intended to model the notion of error ordering that we want terms to have. Thus, we define a predomain  $A$  as a partially-ordered set, which consists of a type which we denote  $\langle A \rangle$  and a reflexive, transitive, and antisymmetric relation  $\leq_P$  on  $A$ .

For each type we want to represent, we define a predomain for the corresponding semantic type. For instance, we define a predomain for natural numbers, a predomain for the dynamic type, a predomain for functions, and a predomain for the lift monad. We describe each of these below.

We define monotone functions between predomain as expected. Given predomains  $A$  and  $B$ , we write  $f: A \rightarrow_m B$  to indicate that  $f$  is a monotone function from  $A$  to  $B$ , i.e, for all  $a_1 \leq_A a_2$ , we have  $f(a_1) \leq_B f(a_2)$ .

- There is a predomain  $\text{Nat}$  for natural numbers, where the ordering is equality.
- There is a predomain  $\text{Dyn}$  to represent the dynamic type. The underlying type for this predomain is defined by guarded fixpoint to be such that  $\langle \text{Dyn} \rangle \cong \mathbb{N} + \triangleright (\langle \text{Dyn} \rangle \rightarrow_m \langle \text{Dyn} \rangle)$ . This definition is valid because the occurrences of  $\text{Dyn}$  are guarded by the  $\triangleright$ . The ordering is defined via guarded recursion by cases on the argument, using the ordering on  $\mathbb{N}$  and the ordering on monotone functions described below.
- For a predomain  $A$ , there is a predomain  $L_{\mathcal{U}}A$  that is the “lift” of  $A$  using the lift monad. We use the same notation for  $L_{\mathcal{U}}A$  when  $A$  is a type and  $A$  is a predomain, since the context should make clear which one we are referring to. The underlying type of  $L_{\mathcal{U}}A$  is simply  $L_{\mathcal{U}}\langle A \rangle$ , i.e., the lift of the underlying type of  $A$ . The ordering on  $L_{\mathcal{U}}A$  is the “lock-step error-ordering” which we describe in 4.3.
- For predomains  $A_i$  and  $A_o$ , we form the predomain of monotone functions from  $A_i$  to  $A_o$ , which we denote by  $A_i \Rightarrow A_o$ . The ordering is such that  $f$  is below  $g$  if for all  $a$  in  $\langle A_i \rangle$ , we have  $f(a)$  is below  $g(a)$  in the ordering for  $A_o$ .

### 4.3 Lock-step Error Ordering

As mentioned, the ordering on the lift of a predomain  $A$  is called the *lock-step error-ordering*, the idea being that two computations  $l$  and  $l'$  are related if they are in lock-step with regard to their intensional behavior, up to  $l$  erroring. Formally, we define this ordering as follows:

- $\eta x \lesssim \eta y$  if  $x \leq_A y$ .
- $\mathbb{U} \lesssim l$  for all  $l$
- $\theta \tilde{r} \lesssim \theta \tilde{r}'$  if  $\triangleright_t (\tilde{r}_t \lesssim \tilde{r}'_t)$

We also define a heterogeneous version of this ordering between the lifts of two different predomains  $A$  and  $B$ , parameterized by a relation  $R$  between  $A$  and  $B$ .

### 4.4 Weak Bisimilarity Relation

We also define another ordering on  $L_{\mathbb{U}}A$ , called “weak bisimilarity”, written  $l \approx l'$ . Intuitively, we say  $l \approx l'$  if they are equivalent “up to delays”. We introduce the notation  $x \sim_A y$  to mean  $x \leq_A y$  and  $y \leq_A x$ .

The weak bisimilarity relation is defined by guarded fixpoint as follows:

$$\begin{aligned}
 &\mathbb{U} \approx \mathbb{U} \\
 &\eta x \approx \eta y \text{ if } x \sim_A y \\
 &\theta \tilde{x} \approx \theta \tilde{y} \text{ if } \triangleright_t (\tilde{x}_t \approx \tilde{y}_t) \\
 &\theta \tilde{x} \approx \mathbb{U} \text{ if } \theta \tilde{x} = \delta^n(\mathbb{U}) \text{ for some } n \\
 &\theta \tilde{x} \approx \eta y \text{ if } (\theta \tilde{x} = \delta^n(\eta x)) \text{ for some } n \text{ and } x : \langle A \rangle \text{ such that } x \sim_A y \\
 &\mathbb{U} \approx \theta \tilde{y} \text{ if } \theta \tilde{y} = \delta^n(\mathbb{U}) \text{ for some } n \\
 &\eta x \approx \theta \tilde{y} \text{ if } (\theta \tilde{y} = \delta^n(\eta y)) \text{ for some } n \text{ and } y : \langle A \rangle \text{ such that } x \sim_A y
 \end{aligned}$$

### 4.5 EP-Pairs

Here, we adapt the notion of embedding-projection pair as used to study gradual typing ([7]) to the setting of intensional denotational semantics.

## 5 SEMANTICS

### 5.1 Relational Semantics

5.1.1 *Types as Predomains.*

5.1.2 *Terms as Monotone Functions.*

5.1.3 *Type Precision as EP-Pairs.*

5.1.4 *Term Precision via the Lock-Step Error Ordering.*

### 5.2 Logical Relations Semantics

## 6 GRADUALITY

The main theorem we would like to prove is the following:

**THEOREM 6.1 (GRADUALITY).** *If  $\cdot \vdash M \sqsubseteq N : \text{Nat}$ , then*

- (1) *If  $N = \mathbb{U}$ , then  $M = \mathbb{U}$*
- (2) *If  $N = 'n$ , then  $M = \mathbb{U}$  or  $M = 'n$*
- (3) *If  $M = V$ , then  $N = V$*

## 6.1 Outline

## 6.2 Extensional to Intensional

## 6.3 Intensional Results

## 6.4 Adequacy

## 6.5 Putting it Together

## 7 DISCUSSION

### 7.1 Synthetic Ordering

While the use of synthetic guarded domain theory allows us to very conveniently work with non-well-founded recursive constructions while abstracting away the precise details of step-indexing, we do work with the error ordering in a mostly analytic fashion in that gradual types are interpreted as sets equipped with an ordering relation, and all terms must be proven to be monotone. It is possible that a combination of synthetic guarded domain theory with *directed* type theory would allow for an a synthetic treatment of the error ordering as well.

## REFERENCES

- [1] Robert Atkey and Conor McBride. 2013. Productive Coprogramming with Guarded Recursion. In *Proceedings of the 18th ACM SIGPLAN International Conference on Functional Programming* (Boston, Massachusetts, USA) (ICFP '13). Association for Computing Machinery, New York, NY, USA, 197–208. <https://doi.org/10.1145/2500365.2500597>
- [2] Lars Birkedal, Rasmus Ejlers Mogelberg, Jan Schwinghammer, and Kristian Stovring. 2011. First Steps in Synthetic Guarded Domain Theory: Step-Indexing in the Topos of Trees. In *2011 IEEE 26th Annual Symposium on Logic in Computer Science*. 55–64. <https://doi.org/10.1109/LICS.2011.16>
- [3] Matteo Cimini and Jeremy G. Siek. 2016. The Gradualizer: A Methodology and Algorithm for Generating Gradual Type Systems. In *Proceedings of the 43rd Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages* (St. Petersburg, FL, USA) (POPL '16). Association for Computing Machinery, New York, NY, USA, 443–455. <https://doi.org/10.1145/2837614.2837632>
- [4] Ronald Garcia, Alison M. Clark, and Éric Tanter. 2016. Abstracting Gradual Typing. In *Proceedings of the 43rd Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages* (St. Petersburg, FL, USA) (POPL '16). Association for Computing Machinery, New York, NY, USA, 429–442. <https://doi.org/10.1145/2837614.2837670>
- [5] Rasmus Ejlers Møgelberg and Marco Paviotti. 2016. Denotational Semantics of Recursive Types in Synthetic Guarded Domain Theory. In *Proceedings of the 31st Annual ACM/IEEE Symposium on Logic in Computer Science* (New York, NY, USA) (LICS '16). Association for Computing Machinery, New York, NY, USA, 317–326. <https://doi.org/10.1145/2933575.2934516>
- [6] H. Nakano. 2000. A modality for recursion. In *Proceedings Fifteenth Annual IEEE Symposium on Logic in Computer Science* (Cat. No.99CB36332). 255–266. <https://doi.org/10.1109/LICS.2000.855774>
- [7] Max S. New and Amal Ahmed. 2018. Graduality from Embedding-Projection Pairs. *Proc. ACM Program. Lang.* 2, ICFP, Article 73 (jul 2018), 30 pages. <https://doi.org/10.1145/3236768>
- [8] Jeremy G. Siek, Michael M. Vitousek, Matteo Cimini, and John Tang Boyland. 2015. Refined Criteria for Gradual Typing. In *1st Summit on Advances in Programming Languages (SNAPL 2015)* (Leibniz International Proceedings in Informatics (LIPIcs), Vol. 32), Thomas Ball, Rastislav Bodik, Shriram Krishnamurthi, Benjamin S. Lerner, and Greg Morrisett (Eds.). Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, Dagstuhl, Germany, 274–293. <https://doi.org/10.4230/LIPIcs.SNAPL.2015.274>