

A PATH TO DOT

MARIANNA RAPOPORT

Formalizing Scala with Dependent Object Types

PhD Thesis

Department of Computer Science

Faculty of Mathematics

University of Waterloo

September 2019 –

DECLARATION

This thesis consists of material all of which I authored or co-authored: see Statement of Contributions included in the thesis. This is a true copy of the thesis, including any required final revisions, as accepted by my examiners.

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

Waterloo, Canada, September 2019

Marianna Rapoport

STATEMENT OF CONTRIBUTIONS

This thesis is a result of my collaboration with Ondřej Lhoták, Ifaz Kabir, Paul He, Lu Wang, and Yaoyu Zhao. The work in the thesis draws heavily on the material presented in the following publications:

- Rapoport, Marianna, Ifaz Kabir, Paul He, and Ondřej Lhoták (2017). “A simple soundness proof for dependent object types.” In: *PACMPL* 1.OOPSLA, 46:1–46:27.
- Rapoport, Marianna and Ondřej Lhoták (2017). “Mutable WadlerFest DOT.” In: *FTfP 2017*, 7:1–7:6.
- Rapoport, Marianna and Ondřej Lhoták (2019). “A Path To DOT: Formalizing Fully Path-Dependent Types.” In: *PACMPL* 1.OOPSLA.

ABSTRACT

The goal of my thesis is to enable formal reasoning about the Scala programming language. To that end I present a *core calculus* that formalizes Scala's

- *essential features* in a
- *type-safe* way and is
- *easy to extend* with more features.

I build on the Dependent Object Types (DOT) calculus that formalizes *path-dependent types*. My contributions are

- a *generalization* of DOT with types that depend on *paths of arbitrary length*,
- a *simple, extensible type-safety proof* for DOT, and
- an *extension* of DOT with *mutable references*.

The simple proof makes designing smaller extensions such as mutation straightforward, and larger extensions, such as full support for paths, approachable. Adding fully path-dependent types to DOT allows us to model the key feature of Scala's type and module system.

The calculi and proofs presented in this thesis are fully mechanized in Coq.

ACKNOWLEDGMENTS

To appear.

CONTENTS

Prologue

1	INTRODUCTION	3
1.1	The Scala Programming Language	3
1.1.1	Modularity Through Abstraction	4
1.1.2	Modularity Through Composition	5
1.2	The DOT Calculus	7
1.3	Limitations of DOT	9
1.4	This Thesis	9
1.4.1	Part I: A Simple Soundness Proof for DOT	9
1.4.2	Part II: A DOT With Mutable References	10
1.4.3	Part III: A DOT With Fully Path-Dependent Types	11
1.4.4	Contributions	13
2	BACKGROUND: THE DOT CALCULUS	15
2.1	DOT Abstract Syntax	15
2.2	DOT Operational Semantics	16
2.3	DOT Typing Rules	16
2.4	Example	18

I A SIMPLE SOUNDNESS PROOF FOR DOT

3	INTRODUCTION	23
4	BAD BOUNDS	27
5	THE SIMPLE DOT PROOF	29
5.1	Overview	29
5.2	Inert Typing Contexts	30
5.3	Tight Typing	31
5.4	Inversion of Tight Typing	35
5.5	Canonical Forms Lemmas	38
5.6	Progress, Preservation, and Soundness	41
5.7	Proof Structure and Extensions	43
5.7.1	Proof Structure	43
6	MODIFICATIONS OF THE CALCULUS	45
7	THE STRUGGLE FOR “GOOD” BOUNDS	47
8	RELATED WORK	49
8.1	DOT Soundness Proofs	49
8.2	History of Scala Calculi	50
8.3	Other Related Calculi	51
8.4	Type Checking Decidability	52
8.5	Syntactic vs. Semantic Proofs	53
9	SUMMARY	55

II CASE STUDY: MUTABLE DOT

10	INTRODUCTION	59
11	THE MUTABLE DOT CALCULUS	61
11.1	Mutable DOT Abstract Syntax	61
11.2	Mutable DOT Operational Semantics	62
11.3	Mutable DOT Typing Rules	63
11.4	Subtyping rules	65
12	TYPE SAFETY	67
13	DISCUSSION	71
13.1	Motivation for a heap of variables	71
13.2	Correctness of a heap of variables	72
13.3	Creating references	73
14	RELATED WORK	75
15	SUMMARY	77

III FULLY PATH-DEPENDENT TYPES

16	INTRODUCTION	81
17	CHALLENGES OF ADDING PATHS TO DOT	87
17.1	Path Limitations in DOT: A Minimal Example	87
17.2	Challenges of Adding Paths to DOT	88
17.2.1	Naive Path Extension Leads to Bad Bounds	89
18	MAIN IDEAS	91
18.1	Paths Instead of Variables	91
18.2	Paths as Identifiers	92
18.2.1	Variables are Identifiers in DOT	92
18.2.2	Paths are Identifiers in pDOT	93
18.2.3	Well-Typed Paths Don't Go Wrong	93
18.3	Path Replacement	94
18.4	Singleton Types	94
18.5	Distinguishing Fields and Methods	95
18.6	Precise Self Types	96
19	FROM DOT TO PDOT	99
19.1	Syntax	99
19.2	pDOT Typing Rules	99
19.2.1	From Variables to Paths	99
19.2.2	Object Typing	101
19.2.3	Path Alias Typing	102
19.2.4	Abstracting Over Field Types	104
19.3	Reduction Semantics	105
20	EXAMPLES	109
20.1	Class Encodings	109
20.2	Lists	111
20.3	Mutually Recursive Modules	112
20.4	Chaining methods with singleton types	113
21	TYPE SAFETY	115
21.1	Inert Types in pDOT	115

21.2	Proof Recipe for pDOT	117
21.2.1	Overview of Extended Proof Recipe	117
21.2.2	Typing Judgments for pDOT's Proof Recipe	119
21.2.3	Proof Recipe Lemmas	121
21.3	Typed-paths Environments	129
21.4	Canonical Forms in pDOT	130
21.4.1	Canonical Forms Proof	131
21.5	Value Typing	143
21.6	Type Soundness for pDOT	145
22	RELATED WORK	149
22.1	Early Class-based Scala Formalizations	149
22.2	DOT-like Calculi	150
22.2.1	Other Related Languages and Calculi	151
22.2.2	Decidability	152
23	CONCLUSION	155
	BIBLIOGRAPHY	157

LIST OF FIGURES

Figure 2.1	Abstract syntax of DOT	15
Figure 2.2	DOT reduction rules	16
Figure 2.3	DOT typing and subtyping rules (Amin, Grütter, et al., 2016)	17
Figure 5.1	Tight Typing Rules (Amin, Grütter, et al., 2016)	32
Figure 5.2	Simplified Precise Typing Rules based on (Amin, Grütter, et al., 2016)	33
Figure 5.3	Invertible typing rules	37
Figure 5.4	Dependencies between main lemmas in the proof. Gray nodes denote existing lemmas. White nodes denote lemmas that would need to be added if DOT were extended with a new type T and a new value v .	44
Figure 11.1	Abstract syntax of Mutable DOT (cf. DOT syntax in Figure 2.1)	61
Figure 11.2	Mutable DOT operational semantics	62
Figure 11.3	Mutable DOT typing rules	64
Figure 11.4	Mutable DOT subtyping rules	66
Figure 12.1	An instance of the dependency graph from Figure 5.4 showing the main lemmas in the Mutable DOT proof as an extension of the simple DOT proof (Part I). Gray nodes denote existing lemmas. White nodes denote Mutable DOT specific lemmas	69
Figure 13.1	Reduction sequence for example program	72
Figure 16.1	A simplified excerpt from the Dotty compiler in Scala. This code fragment cannot be expressed in DOT, as shown on the right	82
Figure 19.1	Abstract syntax of pDOT (cf. DOT syntax in Figure 2.1)	99
Figure 19.2	pDOT typing rules (cf. DOT typing in Figure 2.3)	100
Figure 19.3	pDOT subtyping rules (cf. DOT subtyping in Figure 2.3)	101
Figure 19.4	Replacement of a path p in a type by q	104
Figure 19.5	Operational semantics of pDOT	106
Figure 19.6	Value-environment path lookup	106
Figure 20.1	A covariant list library in pDOT	111
Figure 20.2	Mutually recursive types in a compiler package: fully-path-dependent types	112
Figure 20.3	Chaining method calls using singleton types	113

Figure 21.1	Precise typing in pDOT	119
Figure 21.2	Invertible-I typing in pDOT	122
Figure 21.3	Invertible-II typing in pDOT	123
Figure 21.4	Tight typing for pDOT	124
Figure 21.5	Typed-paths environments	129
Figure 21.6	The type lookup relation	144
Figure 21.7	An instance of the dependency graph from Figure 5.4 showing the main lemmas in the pDOT proof as an extension of the simple DOT proof (Part I). Gray nodes denote the pDOT lemmas that have similar analogues in the DOT proof. White nodes denote pDOT specific lemmas. We omit the Precise-II and Precise-I related lemmas as well as additional lemmas required to prove Corresponding Types and the conversions between tight and invertible-II, and between invertible-II and invertible-I typing.	148

LIST OF TABLES

Table 1.1	Scala calculi that precede DOT	8
Table 21.1	Auxiliary typing relations that make up the proof recipe of pDOT	118

ACRONYMS

ANF Administrative Normal Form

DOT Dependent Object Types calculus by Amin, Grütter, et al.,
(2016)

PROLOGUE

INTRODUCTION

Dependent Object Types (DOT) is intended to be a core calculus for modelling Scala. Its distinguishing feature is *path-dependent types* that refer to fields in objects that hold types rather than values. DOT was designed to serve as an extensible core calculus that could guide the design of future versions of Scala and help us understand the interactions of path-dependent types with other features. The goal of this thesis is to bridge the gap between DOT and Scala by making DOT more expressive and easier to work with.

The first shortcoming of DOT that this thesis addresses is its intricate soundness proof which makes seemingly simple extensions to the calculus complex and unpredictable. I propose a simple, modular, and extensible type-safety proof for DOT. I show how using the simple proof, extensions to the calculus such as mutable references become straightforward.

Second, this thesis presents pDOT, a generalization of DOT that makes it more expressive. DOT is designed to formalize Scala's module system that is based on path-dependent types, but the calculus actually lacks the ability to express a variety of valid Scala paths. As a result, we might be overlooking soundness issues in Scala that are caused by fully *path-dependent* types and can model only a restricted subset of type dependencies that are possible in Scala. In this thesis, I use the simple proof to generalize DOT to support path-dependent types on paths of arbitrary length, as well as *singleton types* to track path equality. I show that naive approaches to add paths to DOT make it inherently unsound, and present the necessary conditions for such a calculus to be sound. I discuss the key changes necessary to adapt the techniques of the DOT soundness proofs so that they can be applied to pDOT.

Path support in DOT allows us to express fully path- (as opposed to variable-) dependent types. Encoding path-dependent types, in turn, allows us to formalize the type dependencies that can occur in Scala, to have a sound theoretical foundation for its type and module system, and realize DOT's full potential for formalizing Scala-like calculi.

1.1 THE SCALA PROGRAMMING LANGUAGE

Scala is a complex programming language with a large number of features that allows one to program in a variety of styles. Higher-order functions, pattern matching, higher-kinded types, type inference, and immutable collections, among other features, make it possible to pro-

gram in a purely functional, immutable style that is reminiscent of Haskell's. Support for classes, traits, and inheritance allow one to also program in a traditionally object-oriented style, using class hierarchies and mutable state; and given Scala's JVM-interoperability, one can essentially write "Java programs" in Scala. Scala supports multiple inheritance through mixin composition; it has both nominal and structural typing. It has rich support for macros; implicit parameters, implicit conversions, and implicit return types; it has abstract type members and path-dependent types.

Which of all these features are essential if we want to formally reason about Scala? What features should be part of a core calculus for the language? According to the Dependent Object Types calculus, the central feature of Scala is *path-dependent types*. To understand why, and to arrive at the language features that are necessary to formalize Scala's path-dependent types, let us look at the design goals of Scala.

One of the design goals of the Scala programming language is to unify modules and objects, so that the same language constructs can be used to define the overall structure of a program as well as its implementation detail (Odersky and Zenger, 2005b). Scala achieves modularity in several ways:

1. Being *object-oriented* allows Scala to benefit from inheritance and aggregation. This allows one to easily extend programs with new data entities without modifying or duplicating existing code, and to easily reuse of large pieces of code.
2. Being *functional* allows Scala to use functions as the "glue" that composes smaller parts of a program together (Hughes, 1989). Scala functions can operate on algebraic data types (case classes); deconstruct them through pattern matching; and use higher-order functions to abstract over parts of a computation. By enabling easy, typesafe extensions of datatypes both with new *data variants* and *processors* Scala provides a solution to the well-known *expression problem* (Odersky and Zenger, 2005a).
3. Finally, Scala *unifies modules and objects*, so that the same language constructs can be used to specify the overall structure of a program as well as its implementation details. The unification of the module and term languages is witnessed by the following comparison with the ML module system:

object \leftrightarrow module
 class \leftrightarrow functor
 interface \leftrightarrow signature

1.1.1 Modularity Through Abstraction

To be composable, modules need to support abstraction. Functional programming abstracts over values and types using value and type

parameters. Object-oriented programming abstracts over class components by keeping them unimplemented. Scala provides abstraction both through parametrization and aggregation.

In most object-oriented languages, the only kind of class components that are allowed to be abstracted over by remaining unimplemented are *methods*. In Scala, abstraction over object members is also allowed for *values* and *types*.

Types that are members of classes and objects are called *abstract type members* (we will sometimes refer to them just as type members). An abstract type member is a type that is declared inside a class¹ but whose definition can be left unspecified. We can refine the declaration of abstract type members through *type bounds*. If the type member A of an object p should be a subtype of T, we can declare A with an upper bound: { type A <: T }. To make A a supertype of U, we declare it with a lower bound: { type A >: U }. To refer to A, we use the *path-dependent type* p.A — a type that depends on the object whose *path* is p = x.a₁. . . .a_n where x is the path’s receiver, and a_i are its fields.

abstract type members

type bounds

path-dependent types

1.1.2 Modularity Through Composition

To compose modules, Scala uses *mixin composition*, a form of multiple inheritance that uses a carefully designed *linearization* principle to avoid dispatch ambiguity. To abstract away from the rules of class linearization, we can focus on a simpler, commutative way of composing types called *type intersection*, which is not supported by Scala 2 (current Scala), but is already a feature of Dotty, the compiler for Scala 3 (Documentation, 2018a). A type intersection $T \wedge U$ is a type that has the members of both T and U.

type intersection

“Union” might seem like a better name than “intersection”. However, consider a class Dog that extends Animal. The set of Dogs is a subset of all Animals. Therefore, to denote the members that are both in Dog and Animal we would use intersection.

Intersection allows us to express *type refinement*; for example, intersecting the types { type A <: T } and { type A >: U } yields { type A >: U <: T }, a type declaration that has a lower and upper bound at the same time. In fact, { type A <: T } and { type A >: U } are just syntactic sugar for { type A >: Nothing <: T } and { type A >: U <: Any }, where Any and Nothing are the top and bottom of Scala’s subtyping lattice.

The combination of abstract type members, path-dependent types, and type intersection yields a surprisingly expressive type system. These concepts alone can encode a variety of features that are common in Scala and other programming languages:

- *parametric polymorphism* (generics):
 - a *parametrized class* class List[T] can be encoded using type members as class List { type T }, whereas *type application* List[Int] can be represented using refinement (type intersection): $\text{List} \wedge \{ \text{type T} = \text{Int} \}$;¹

¹ More complicated representations such as higher-kinded types require more complex encodings (Odersky, Martres, and Petrashko, 2016).

¹ The following text applies to both classes and *traits*, which are similar to Java’s interfaces but allow methods to have implementations.

- *type parameters* for polymorphic functions such as `def id[T](x: T) = x` can be encoded using type members as follows: `def id(y: { type T })(x: y.T) = x`
- *co- and contravariance* can be often encoded by translating declaration-site variance into use-site variance; for example, a covariant list

```
class List[+T] { def tail: List[T] }
```

can be expressed as

```
class List { type T; def tail: List ^ { type T <: this.T } }
```

(note the path-dependent type `this.T` which references the instance of `T` that belongs to `tail`'s enclosing `List` object) (Amin, Grütter, et al., 2016);

- *existential polymorphism*: type members can encode a subset of use cases for existential types, as shown by Norris, (2015);
- *family polymorphism* (Ernst, 2001): we can distinguish path-dependent types that belong to different objects (“families”) at compile time; for example, given two objects `tree` and `dag` of type `Graph` and a method `edge(g: Graph, n1: g.Node, n2: g.Node)`, a method call `edge(tree, treeNode, dagNode)` that is invoked on nodes belonging to different graphs will result in a type error;
- *subtyping hierarchies* between *recursively defined classes*: the general problem is to preserve the recursion between two classes `A2` and `B2` that inherit from two mutually recursive classes `A` and `B` respectively; as shown by Bruce, Odersky, and Wadler, (1998), such relationships require the use of additional types that are cumbersome and impractical to express using generics, leading to a quadratic increase in the number of necessary type parameters; however, it is easy to express this type of recursion/subtyping relationships using type members.

Abstract type members are thus a powerful feature that can encode a variety of useful programming concepts; but could they be too powerful? For example, how safe is it to allow type members to be declared with arbitrary lower and upper bounds? What if a type is declared with incompatible bounds, such as `{ type A >: Any <: Nothing }` or `{ type A >: Int <: String }` — should the type system try to detect such “bad bounds” and ban “nonsensical” type declarations? Furthermore, we reference an abstract type member through a path-dependent type `p.A`, which depends on the runtime value of `p`. What happens if `p` does not stop evaluating? Finally, will something break if we couple type members with other features? Which feature combinations are safe and which are not (Amin, Grütter, et al., 2016; Amin, Moors, and Odersky, 2012; Amin, Rompf, and Odersky, 2014; Moors, Piessens,

and Odersky, 2008; Odersky, Cremet, et al., 2003; Rompf and Amin, 2015, 2016a)?

To answer these questions we need to be able to formally reason about Scala’s type system. That is where Dependent Object Types come in.

1.2 THE DOT CALCULUS

The Dependent Object Types (DOT) calculus is a sound formalization of “the essence of Scala”: abstract type members with arbitrary type bounds, path-dependent types, and type intersection. Rompf and Amin, (2015) and Amin, Grütter, et al., (2016) presented the first DOT soundness proofs based on a big- and small-step semantics after a decade-long search for a type-safe core calculus for Scala. Table 1.1 summarizes this search. The decreasing number of shaded boxes as we go from left to right illustrates that in order to be type-safe, the formalizations had to lose more and more Scala features — sound versions of DOT handle hardly more than the essential features of path-dependent types (displayed on black/dark gray background). The table also shows that the first two type-safe calculi, νObj and μDOT , did not handle arbitrary intersections of type declarations since they either exclude intersections or lower bounds, which further highlights the difficulty of formalizing all the type-member-related features.

This thesis focuses on the small-step version of DOT by Amin, Grütter, et al., (2016).

Discovering a sound DOT has already had the following practical benefits.

1. The formalization of type members enables a better understanding of the features that type members can translate to. For example, because DOT can encode Java’s wildcards (which are used to express variance in type parameters), DOT provides the first formalization of a type system that supports wildcards (Amin, Grütter, et al., 2016).
2. DOT has helped reveal feature combinations in Scala and Java that lead to unsoundness. Amin and Tate, (2016) showed how, using wildcards and null in Java, and type members and null in Scala, we can convert between arbitrary types without casting. The authors report that their examples are inspired by the insights gained from the DOT safety proof.
3. To provide a sound theoretical foundation for Scala, insights gained from DOT are guiding the design of the Dotty compiler:
 - Dotty aims to translate type parameters into type members and type application into type intersection because these features are formalized in DOT.

formalization of Java’s wildcards

discovery of new source of Scala and Java unsoundness

design guidelines for Dotty




Table 1.1: Scala calculi that precede DOT

- Intersection has been added to Dotty to faithfully mimic the DOT calculus and ensure soundness (Documentation, 2018a).
- Examples of how DOT influences decisions for representing higher-kinded types in the Dotty compiler are documented by Odersky, Martres, and Petrashko, (2016).
- Exposing unsoundness due to null is one of the motivations behind the ongoing Dotty work to discourage the use of null by forcing programmers to explicitly declare nullability through union types (Nieto, 2018).

*ruling out
nonsensical
dependent types can
lead to dead end*

4. Studying DOT has shown that trying to detect and rule out bad bounds during type-checking is impractical (Amin, Grütter, et al., 2016; Amin, Rompf, and Odersky, 2014). Instead, it is the task of the soundness proof to determine that the types of values that will appear at runtime will eventually be *narrowed* to meaningful types with “good bounds”. This realization has guided the decisions of how to approach bug reports such as Issue #1050 (Odersky, 2016).

Given the importance of having a formalization for Scala, we need to be able to extend DOT to bridge the gap between Scala and its core calculus. However, there are some problems.

1.3 LIMITATIONS OF DOT

To come up with a sound DOT, Amin, Grütter, et al., (2016) had to make two crucial decisions. These decisions were groundbreaking in that they enabled the first sound formalization of a difficult concept. But they also turned out to get in the way of making DOT an extensible core calculus.

The first decision relates to the soundness proof. To eliminate bad bounds such as $\{ \text{type } A >: \text{String} <: \text{Int} \}$, Amin et al. used the insight that if a type is inhabited it cannot have bad bounds since the value that inhabits the type is a proof that the type makes sense. To incorporate reasoning about inhabited types into the proof, whenever the proof depends on the absence of bad bounds, it couples types with values that inhabit those types. Since a large part of the proof requires meaningful information from types, this introduces complicated dependencies and reasoning into the proof. As a result, the proof becomes brittle: when designing extensions it is difficult to predict what will break.

The second decision directly affects the calculus: it is to restrict path-dependent types exclusively to variables. Type selections on variables of the form $x.A$ are DOT's way to underapproximate Scala's type selections on arbitrary *stable* paths $x.a_1 \dots a_n.A$, which must consist of immutable fields a_i . Restricting paths in path-dependent types to variables limits the expressivity of DOT and does not allow us to express a whole class of recursive dependencies between types. Conversely, generalizing from variables to paths reveals challenging soundness issues, as we will show in Part III of this thesis.

1.4 THIS THESIS

This thesis addresses the issues of DOT's complex soundness proof and incomplete modelling of path-dependent types, which limits DOT's applicability as a reusable core calculus for Scala.

1.4.1 Part I: A Simple Soundness Proof for DOT

The first step towards a reusable calculus is to simplify the DOT soundness proof.

Abstract type members allow DOT to express custom subtyping relationships within the calculus. For example, the recursive type

$$\mu(a: \{ \text{SmallFish}: \perp.. \top \} \wedge \\ \{ \text{MediumFish}: a.\text{SmallFish}..a.\text{BigFish} \} \wedge \\ \{ \text{BigFish}: \perp.. \top \})$$

contains three abstract type members. `SmallFish` and `BigFish` are fully abstract: they can be anything between \perp and \top , whereas `MediumFish`

must be a supertype of `SmallFish` and a subtype of `BigFish`. Since `SmallFish <: MediumFish <: BigFish`, by transitivity, `SmallFish` must be a subtype of `BigFish`, thus allowing us to encode fish size in the type system. Even though the definitions of `SmallFish` or `BigFish` say nothing about their size, when we instantiate these abstract type members, the type system will ensure that the declared subtyping relationships hold.

The challenge of the DOT soundness proof is to show that the type system will accept “meaningful” type declarations while ruling out nonsensical types such as `{StrangeFish: Piranha..Goldfish}`. The proof of Amin, Grütter, et al., (2016) addresses this problem by requiring that all typing contexts be inhabited with values: to show that a type is valid, it must be coupled with a value of that type.

The insight of requiring inhabited typing contexts allowed Amin, Grütter, et al., (2016) to present the first type-safe calculus with path-dependent types. The drawback of this proof is, however, that a significant part of it combines reasoning about types, variables, and values. The proof becomes not just intricate but brittle, making it hard to predict the effects of even the simplest changes.

To make the DOT soundness proof easier to work with, I present a simplified soundness proof with the following contributions:

- A *modular* proof that reasons about types, values, and operational semantics separately.
- The concept of *inert* typing contexts, a syntactic characterization of contexts that rule out any nonsensical subtyping that could be introduced by abstract type members.
- A simple *proof recipe* for deducing properties of terms from their types in full DOT while reasoning only in a restricted, intuitive environment free from the paradoxes caused by abstract type members. Multiple lemmas follow the same recipe, and following the recipe can facilitate the development of new lemmas needed in future extensions for DOT.
- A *Coq formalization* of this DOT soundness proof.

The result is a modular, simple proof of DOT type safety, which enables us to formalize other constituents of Scala, such as classes and inheritance, by a translation to the core features of DOT. The simple soundness proof was developed in collaboration with Ifaz Kabir, Paul He, and Ondřej Lhoták.

1.4.2 Part II: A DOT With Mutable References

To illustrate that the developed DOT soundness proof is indeed simple, I present a case study of extending DOT with support for ML-style

mutable references. The extension, called Mutable DOT, adds support for a heap; the ability to create, update, and dereference mutable references; and reference types. Proving Mutable DOT sound involved adding the following to the simple soundness proof:

- a new case to the definition of inert types: any reference type is inert;
- an additional case to the definitions of invertible typing for variables and invertible typing for values;
- two new canonical forms lemmas for values and variables of a reference type; these lemmas follow the proof recipe;
- the mutation-related cases to the proof-recipe lemmas that translate tight into invertible typing, and to the final progress and preservation lemmas.

Extending DOT with mutation reveals two aspects of working with DOT and the simple proof. On the one hand, it shows that if the change to the calculus involves mostly adding new forms of values and types then it can be fairly easy and straightforward to add new features to the calculus. In the case of mutation, the change involves adding a new type of values (locations) and types (reference types). Once we know what exactly we want to extend the calculus with, it is easy to extend the proof.

However, the challenge with DOT is usually to know how to redesign the calculus: to understand and predict what exactly the extended DOT should look like. Even in the case of mutation, a seemingly small and simple change, it was not obvious how to design the mutable store. Specifically, in order to maintain soundness, the mutable store had to map locations to *variables* as opposed to values, as would be more conventional. In the second part of this thesis I present the resulting calculus, show that our design decisions for Mutable DOT are sound, correct, and expressive, and explain in detail how to extend the simple proof to support mutation.

1.4.3 Part III: A DOT With Fully Path-Dependent Types

The final step towards a core calculus for Scala is to formalize path-dependent types in their full expressivity. The goal is to allow DOT to have paths of arbitrary length: such paths are admissible in Scala, but not expressible in DOT.

The DOT calculus is defined using Administrative Normal Form (ANF). ANF requires that function applications $x\ y$ and path selections $z.a$ always occur on variables: as we will show in Chapter 2, each of x , y , z must be either bound in a let expression or lambda abstraction, or be the self-variable of an object. Most of the time, this does not

limit expressivity, because we can chain sequences of let bindings to obtain the complex expressions that we need. The exception to that is path-dependent types.

Consider the following Scala object:

```
val o = {
  val a = new { type A = Any }
  type B = this.a.A
}
```

The definition of B cannot be expressed in DOT. B refers to type member A, which is nested inside the definition of another field a. Field accesses always constitute paths at least of length two, since we must reference the object (in this case o) that the fields belong to. As we will show in Chapter 17, any attempts at modifying the example to shorten the path to A involve flattening the structure of the whole program which disallows nesting, a fundamental feature of object-oriented programming. To model the relationships between classes that refer to each other from nested packages we thus need to use paths of arbitrary length.

Allowing full paths in DOT poses several challenges:

- | | |
|--|--|
| <i>avoid
non-terminating
paths</i> | 1. As we will show, to maintain soundness and avoid the <i>bad bounds</i> issue we must ensure that type selections occur only on normalizing (i.e. terminating, acyclic) paths. |
| <i>allow recursion</i> | 2. At the same time, the calculus should still permit <i>non-terminating paths</i> in general, since that is necessary to express recursive functions. |
| <i>track path equality</i> | 3. We must be able to track <i>equality</i> between paths that alias the same object, while distinguishing between paths that reference syntactically equal objects. |
| <i>proof engineering of
paths</i> | 4. Reasoning about paths of arbitrary length poses additional <i>proof-engineering</i> challenges. For example, the presence of paths introduces new possibilities for cyclic typing derivation which makes doing induction on typing difficult. Another problem is the use of paths where DOT uses variables: ANF allowed the original DOT proof to operate on objects and recursive types at their “outermost” level, while the pDOT proof needs to look inside an object’s nested field definitions, supporting recursive reasoning about the correspondences between objects and their types at the nested levels. |

In the last part of my thesis I present pDOT, a generalization of DOT with paths of arbitrary length. pDOT addresses the above challenges as follows:

- | | |
|--|--|
| <i>separate path lookup
from operational
semantics</i> | 1. To avoid unsoundness, the following considerations guided the design of pDOT: |
|--|--|

- To ensure that paths are terminating we define them to be *irreducible* in the operational semantics.
 - At the same time, we define a *path lookup* operation that retrieves the value that a path refers to in the runtime environment, and show that paths that have the types of values always reference a value of the same type.
 - We restrict how objects are defined and typed at nested levels while still allowing type abstractions that are even more expressive than DOT's.
2. To allow non-terminating paths pDOT supports cyclic references between paths which makes the calculus less restrictive and also easier to define. *allow cyclic paths*
 3. To track path aliasing in the type system we use *singleton types*. A singleton type $p.\text{type}$ is inhabited with only one value: p . If a path p has a singleton type $q.\text{type}$ it signals to the type system that p and q are aliases. The use of singleton types makes pDOT also the first calculus that formalizes that Scala feature. *singleton types track path equality*
 4. The pDOT type-safety proof handles arbitrary paths in places where DOT uses variables using various approaches. For example, the proof recipe is stratified into additional stages that allow it to avoid additional sources for cyclic typing derivations. The proof is presented in detail in Chapter 21. *proof is less "simple"*

1.4.4 Contributions

In summary, this thesis makes the following contributions:

1. A soundness proof that is *modular* because it decouples the reasoning of types and values from each other and *simple* because it presents a simple recipe that one can follow whenever they need to make sense of types. *Part I:
simple DOT proof*
2. A case study of adding mutable references to DOT that provides an extension of DOT with a fundamental Scala feature and shows how easy it is to extend the simple proof. *Part II:
DOT with mutation*
3. The pDOT calculus, a generalization of DOT that lifts the type-selection-on-variables restriction and supports paths of arbitrary length. pDOT fully formalizes path-dependent types; it can express the whole variety of class dependencies that are possible in Scala, and allows us to unleash the full expressive power of path-dependent types. *Part III:
DOT with full paths*

BACKGROUND: THE DOT CALCULUS

The development in this thesis follows the variant of the DOT calculus defined by Amin, Grütter, et al., (2016).

2.1 DOT ABSTRACT SYNTAX

We begin by describing the abstract syntax of the calculus. DOT defines two forms of *values* v :

- A *lambda abstraction* $\lambda(x: T) t$ is a function with parameter x of type T and a body consisting of the term t .
- An *object* of type T with definitions d is denoted as $v(x: T)d$. The body of the object consists of the definitions d , which are a collection of field and type member definitions, connected through the intersection operator. The field definition $\{a = t\}$ assigns a term t to a field labeled a , and the type definition $\{A = U\}$ defines the type label A as an alias for the type U . The object also explicitly declares a recursive *self*, or “this”, variable x . As a result, both T and d can refer to x .

A DOT *term* t is a variable x , value v , field selection $x.a$, function application $x y$, or let binding **let** $x = t$ **in** u . To keep the syntax simple, the DOT calculus uses administrative normal form (ANF); as a result, field selection and function application can involve only variables, not arbitrary terms.

A DOT *type* T can be one of the following:

- A *dependent function type* $\forall(x: S) T$ is the type of a function with a parameter x of type S , and with the return type T , which can refer to the parameter x .
- A *recursive type* $\mu(x: T)$ declares an object type T which can refer to its self-variable x .
- A *field declaration* $\{a: T\}$ states that the field labeled a has type T .
- A *type declaration* $\{A: S..T\}$ specifies that an abstract type member A is a subtype of T and a supertype of S .
- A *type projection* $x.A$ is the type assigned to the type member labelled A of the object x (ANF allows type projection only on variables).
- An *intersection type* $S \wedge T$ is the most general subtype of both S and T .

x, y, z	variable
a, b, c	term label
A, B, C	type label
t, u	term
v	value
d	definition
S, T, U	type

Metavariables used throughout the thesis

$t, u :=$
x
$x.a$
v
$x y$
let $x = t$ in u
$v :=$
$v(x: T)d$
$\lambda(x: T) t$
$d :=$
$\{a = t\}$
$\{A = T\}$
$d \wedge d'$
$S, T, U :=$
\top
\perp
$\{a: T\}$
$\{A: S..T\}$
$x.A$
$S \wedge T$
$\mu(x: T)$
$\forall(x: S) T$

Figure 2.1: Abstract syntax of DOT

- The *bottom* type \perp and the *top* type \top correspond to the bottom and top of the subtyping lattice, and are analogous to Scala’s Nothing and Any.

2.2 DOT OPERATIONAL SEMANTICS

the store γ maps
variables to values:

$$\gamma ::= \emptyset \\ |\gamma, x \mapsto v$$

¹ Evaluating fields strictly would require DOT to introduce a field initialization order which would complicate the calculus. DOT deliberately leaves initialization as an open question. For a DOT with constructors and strict fields, see Kabir and Lhoták, (2018).

The DOT reduction relation operates on terms whose free variables are bound in a runtime value environment, or *store*, γ that maps variables to values.

Variables and values are considered normal form, i.e. they are irreducible. In particular, objects $v(x: T)d$ are values, and the fields of an object are not evaluated until those fields are selected. Because in DOT fields are thus lazily evaluated, they are similar to Scala’s lazy vals which declare immutable, lazily evaluated values.¹ The reduction rules use ANF to ensure that before a term is used in a function application or field selection it is reduced to a value that is assigned to a variable through let bindings.

$$\frac{\gamma(x) = v(x: T) \dots \{a = t\} \dots}{\gamma \mid x.a \mapsto \gamma \mid t} \quad (\text{PROJ})$$

$$\frac{\gamma(x) = \lambda(z: T) t}{\gamma \mid x y \mapsto \gamma \mid t[y/z]} \quad (\text{APPLY})$$

$$\gamma \mid \text{let } x = y \text{ in } t \mapsto \gamma \mid t[y/x] \quad (\text{LET-VAR})$$

$$\gamma \mid \text{let } x = v \text{ in } t \mapsto \gamma, x \mapsto v \mid t \quad (\text{LET-VALUE})$$

$$\frac{\gamma \mid t \mapsto \gamma' \mid t'}{\gamma \mid \text{let } x = t \text{ in } u \mapsto \gamma' \mid \text{let } x = t' \text{ in } u} \quad (\text{CTX})$$

Figure 2.2: DOT reduction rules

We denote the reflexive, transitive closure of \mapsto as \mapsto^* .

2.3 DOT TYPING RULES

The DOT typing rules are presented in Figure 2.3. The rules ALL-I and {}-I give types to values. An object $v(x: T)d$ has the recursive type $\mu(x: T)$, where the types T must match, and T must summarize the definitions d following the definition typing rules in Figure 2.3. Note that due to DEF-TYP, each of the type declarations in an object must have equal lower and upper bounds (i.e. an object $v(x: \{A: S..U\}) \{A = T\}$ is only well-typed if $S = U = T$). The rules VAR, ALL-E, FLD-E, LET give types to the other four forms of terms, and are unsurprising. The recursion introduction REC-I, recursion

$\frac{\Gamma(x) = T}{\Gamma \vdash x: T} \quad (\text{VAR})$	$\frac{\Gamma \vdash t: T \quad \Gamma, x: T \vdash u: U \quad x \notin \text{fv}(U)}{\Gamma \vdash \text{let } x = t \text{ in } u: U} \quad (\text{LET})$	<i>Term typing</i> $\Gamma \vdash t: T$
$\frac{\Gamma, x: T \vdash t: U \quad x \notin \text{fv}(T)}{\Gamma \vdash \lambda(x: T) t: \forall(x: T) U} \quad (\text{ALL-I})$	$\frac{\Gamma \vdash x: T}{\Gamma \vdash x: \mu(x: T)} \quad (\text{REC-I})$	
$\frac{\Gamma \vdash x: \forall(z: S) T \quad \Gamma \vdash y: S}{\Gamma \vdash x y: T[y/z]} \quad (\text{ALL-E})$	$\frac{\Gamma \vdash x: \mu(z: T)}{\Gamma \vdash x: T[x/z]} \quad (\text{REC-E})$	
$\frac{\Gamma, x: T \vdash d: T}{\Gamma \vdash \nu(x: T) d: \mu(x: T)} \quad (\{\}-\text{I})$	$\frac{\Gamma \vdash x: T \quad \Gamma \vdash x: U}{\Gamma \vdash x: T \wedge U} \quad (\text{AND-I})$	
$\frac{\Gamma \vdash x: \{a: T\}}{\Gamma \vdash x.a: T} \quad (\text{FLD-E})$	$\frac{\Gamma \vdash t: T \quad \Gamma \vdash T <: U}{\Gamma \vdash t: U} \quad (\text{SUB})$	
$\frac{\Gamma \vdash t: U}{\Gamma \vdash \{a = t\}: \{a: U\}} \quad (\text{DEF-TRM})$	$\frac{\Gamma \vdash d_1: T_1 \quad \Gamma \vdash d_2: T_2 \quad \text{dom}(d_1), \text{dom}(d_2) \text{ disjoint}}{\Gamma \vdash d_1 \wedge d_2: T_1 \wedge T_2} \quad (\text{ANDDEF-I})$	<i>Definition typing</i> $\Gamma \vdash d: T$
$\Gamma \vdash \{A = T\}: \{A: T..T\} \quad (\text{DEF-TYP})$		
$\Gamma \vdash T <: \top \quad (\text{TOP})$	$\frac{\Gamma \vdash T <: U}{\Gamma \vdash \{a: T\} <: \{a: U\}} \quad (\text{FLD-<:-FLD})$	<i>Subtyping</i> $\Gamma \vdash T <: U$
$\Gamma \vdash \perp <: T \quad (\text{BOT})$		
$\Gamma \vdash T <: T \quad (\text{REFL})$	$\frac{\Gamma \vdash S_2 <: S_1 \quad \Gamma \vdash T_1 <: T_2}{\Gamma \vdash \{A: S_1..T_1\} <: \{A: S_2..T_2\}} \quad (\text{TYP-<:-TYP})$	
$\Gamma \vdash T \wedge U <: T \quad (\text{AND}_1\text{-<:})$		
$\Gamma \vdash T \wedge U <: U \quad (\text{AND}_2\text{-<:})$		
$\frac{\Gamma \vdash S <: T \quad \Gamma \vdash S <: U}{\Gamma \vdash S <: T \wedge U} \quad (<:-\text{AND})$	$\frac{\Gamma \vdash S_2 <: S_1 \quad \Gamma, x: S_2 \vdash T_1 <: T_2}{\Gamma \vdash \forall(x: S_1) T_1 <: \forall(x: S_2) T_2} \quad (\text{ALL-<:-ALL})$	
$\frac{\Gamma \vdash x: \{A: S..T\}}{\Gamma \vdash S <: x.A} \quad (<:-\text{SEL})$		
$\frac{\Gamma \vdash x: \{A: S..T\}}{\Gamma \vdash x.A <: T} \quad (\text{SEL-<:})$	$\frac{\Gamma \vdash S <: T \quad \Gamma \vdash T <: U}{\Gamma \vdash S <: U} \quad (\text{TRANS})$	

Figure 2.3: DOT typing and subtyping rules (Amin, Grütter, et al., 2016)

elimination REC-E, and intersection introduction AND-I rules apply only to variables, but the subsumption rule SUB applies to all terms.

The subtyping rules establish the top and bottom of the subtyping lattice (TOP, BOT), define reflexivity and transitivity (REFL, TRANS), and basic subtyping rules for intersection types (AND₁-<:, AND₂-<:, <:-AND). As is commonplace, dependent functions are covariant in the return type and contravariant in the parameter type ALL-<:-ALL. Field typing is covariant by the rule FLD-<:-FLD, whereas type member declarations are contravariant in the lower bound and covariant in the upper bound via TYP-<:-TYP. The most interesting rules that distinguish DOT are <:-SEL and SEL-<:, which introduce an object-dependent type $x.A$ and define subtyping between it and its bounds. As we will see, these rules are responsible for much of the complexity of the safety proof.

2.4 EXAMPLE

This section presents an example DOT program to give the reader a better intuition for the DOT calculus. This example will come up later when we talk about mutation in Part II.

Suppose we want to keep track of fish that live in aquariums. In Scala, we could write:

```
object AquariumModule {
  trait Aquarium {
    type Fish
    val fish : List[Fish]
  }

  def addFish(a: Aquarium)(f: a.Fish) =
    new Aquarium {
      type Fish = a.Fish
      val fish = a.fish :+ f
    }

  val piranhas = new Aquarium {
    type Fish = Piranha
    val fish = List.empty[Piranha]
  }

  val goldfish = new Aquarium {
    type Fish = Goldfish
    val fish = List.empty[Goldfish]
  }
}
```

This program lets us add a fish gf to the goldfish aquarium:

```
val gf: Goldfish = ...
addFish(goldfish, gf)
```

but it will result in a type error when trying to add gf to the piranha aquarium:

```
addFish(piranhas, gf) // expected: piranhas.Fish, actual: goldfish.Fish
```

The reason the goldfish is protected from the piranhas is that the type Fish is *path dependent*, i.e. specific to the run-time Aquarium object that the fish belongs to. This allows the addFish method to guarantee at compile time that an aquarium accepts only fish of type a.Fish.

As a first attempt to define Aquarium in DOT, we can make it an *intersection* of two types:

```
{ Aquarium = {Fish:  $\perp..T$ }  $\wedge$  {fish: List} }
```

The first type, $\{Fish: \perp..T\}$, declares a type member Fish with lower bound \perp (Nothing) and upper bound T (Any). The second type, $\{fish: List\}$, is a field declaration of type List that represents the list of fish in the aquarium. The type List is assumed to be defined in a library and contains a type member A for list elements.

A problem with the current Aquarium implementation is that it does not say that the type of elements in the fish list should be Fish. More specifically, the list elements should have the Fish type of the Aquarium *runtime object* to which the list belongs. To let the Aquarium type refer to its own runtime object a, we make Aquarium a *recursive type*:

```
{ Aquarium =  $\mu(a: \{Fish: \perp..T\} \wedge$   
                {fish: List  $\wedge$  {A: a.Fish .. a.Fish}}}
```

Here, to express that the type Fish should belong to the object a, we use the *type selection* a.Fish. The type a.Fish is then used as a refinement of List's element type A. In this way, the list can contain only the fish that are allowed in the aquarium a.

We can now define addFish as a function that takes an aquarium a and a fish f of type a.Fish, and creates a new aquarium a2:

```
{ addFish =  $\lambda(a: \text{aq.Aquarium}).\lambda(f: a.Fish).$   
     $\nu(a2: \text{aq.Aquarium} \wedge \{Fish: a.Fish .. a.Fish\}) \{$   
        Fish = a.Fish  
        fish = ... } }
```

The construct $\nu(x: T)d$ creates a new object of type T with a self-variable x and definitions d. In this case, the definitions are used to initialize the Fish type and fish list of the new aquarium. The Fish type is assigned a.Fish. The new fish list needs to append the fish f to the old a.fish list.

To be able to add an element to a list, we need access to an append method, which we will get from List. Suppose that the List type is defined in a collections library. It can be defined as a recursive type $\mu(\text{list}: \dots)$ that declares an element type A and an append function. The append function takes a parameter a of the element type list.A and returns a List of elements that are subtypes of a.A:

```
let collections =  $\nu(\text{col:}$ 
```

$$\{ \text{List} : \mu(\text{list} : (\{A : \perp..T\} \wedge \{\text{append} : \forall(a : \text{list} . A)(\text{col} . \text{List} \wedge \{A : \perp..a.A\}))) \} \dots$$

in ...

With an append method on Lists, we can fully implement the addFish method. The field a2.fish should be defined as a.fish.append(f). However, since DOT uses ANF, before performing any operations on terms, we have to bind the terms to variables:

```
fish = let oldFish = a.fish in
      let append = oldFish.append in
      append f
```

For better readability, we introduce the following DOT abbreviations (similar ones are used by Amin, Grütter, et al., (2016)):

$$\begin{aligned} \{A\} &\equiv \{A : \perp..T\} & t\ u &\equiv \text{let } x = t \text{ in} \\ \{A : T\} &\equiv \{A : T..T\} & \text{let } y = u \text{ in } x\ y \\ \{A <: T\} &\equiv \{A : \perp..T\} & t.L &\equiv \text{let } x = t \text{ in } x.L \\ \{D_1; D_2\} &\equiv \{D_1\} \wedge \{D_2\} & \nu(x)d : T &\equiv \nu(x : T)d \end{aligned}$$

where D_1, D_2 are declarations or definitions of either fields or types, and L is a label of a type or field.

With those abbreviations, the full aquarium program example looks as follows:

```
let collections =  $\nu(\text{col}) \{ \dots \}$ :
  { List :  $\mu(\text{list} : (A; \text{append} : \forall(a : \text{list} . A)(\text{col} . \text{List} ; A <: a.A)))$  }
in  $\nu(\text{aq}) \{$ 
  Aquarium =  $\mu(a : \{ \text{Fish}; \text{fish} : \{ \text{collections} . \text{List} ; A : a.\text{Fish} \} \})$ ;
  addFish =  $\lambda(a : \text{aq}.\text{Aquarium}).\lambda(f : a.\text{Fish}).$ 
     $\nu(a2) \{$ 
      Fish = a.Fish
      fish = a.fish.append f
    } : {aq.Aquarium; Fish: a.Fish}
  } : {Aquarium:  $\mu(a : \{ \text{Fish}; \text{fish} : \{ \text{collections} . \text{List} ; A : a.\text{Fish} \} \})$ ;
  addFish:  $\forall(a : \text{aq}.\text{Aquarium})\forall(f : a.\text{Fish})$ 
    {aq.Aquarium; Fish: a.Fish}}
```

Part I

A SIMPLE SOUNDNESS PROOF FOR DOT

INTRODUCTION

Scala’s type system is notoriously complex. It allows one to define abstract types as members of objects. To become less abstract and more interesting, these type members can be lower- and upper-bounded by other types. The type bounds can be arbitrary: there is no rule or restriction on what the relationship between a type member and its bounds should be. Furthermore, a type member A ’s type bounds can change when A ’s enclosing class is composed with other types, for example, through mixin composition. The type bounds can also depend on other objects. This includes A ’s own enclosing object, which allows A ’s bounds to recursively reference A itself.

If the role of a type-soundness proof is to ensure that we can safely rely on types to give us meaningful information about the values in a program then it is unsurprising that a soundness proof for Scala’s type system is hard. After all, it is the task of that proof to show that a type system with the above features is not too expressive for its own good: that it does not let us express anything we want, that it will not allow arbitrary subtyping relationships between types, and that instead it will provide us with meaningful type information about our programs.

This is why 2016 was an exciting year for those who desire a formalism to understand and reason about the unique features of Scala’s type system. Mechanized soundness results were published for the DOT calculus (Amin, Grütter, et al., 2016; Amin and Rompf, 2017; Rompf and Amin, 2016b). These proofs were the culmination of an elusive search that spanned more than ten years. The chief subtleties and paradoxes inherent in DOT and the Scala type system, which made the proof so challenging, were documented along the way (Amin, Moors, and Odersky, 2012; Amin, Rompf, and Odersky, 2014).

Since the DOT calculus exhibits such subtle and counterintuitive behaviour, and since the proofs are the result of such a long effort, it is to be expected that the proofs must be complicated. The calculus is dependently typed, so it is not surprising that the lemmas that make up the proofs reason about tricky relationships between types and values. In some contexts, the type system admits typings that seem just plain wrong, and give no hope for soundness, so it seems necessary to have lemmas that reason simultaneously about the intricate properties of values, types, and the environments that they inhabit.

A core calculus needs to be easy to extend. Some extensions of DOT are necessary even just to model essential Scala features. As a prominent example, types in Scala may depend on paths (e.g. $x.a_1.\dots.a_n.A$)

*we address the
problem of paths
in Part III*

but types in the existing DOT calculi can depend only directly on variables ($x.A$). Path-dependent types are needed to model essential features such as classes and traits (as members nested in objects and packages) and the famous cake pattern (Odersky and Zenger, 2005b). Another important Scala feature to be studied in DOT are implicit parameters. Moreover, language modifications and extensions are the *raison d'être* of a core calculus. DOT enables designers to experiment with exciting new features that can be added to Scala, to tweak them and reason about their properties before attempting to integrate them in the compiler with the complexity of the full Scala language.

The complexity of the proof is a hindrance to such extension and experimentation. Over the past ten years, DOT has been designed and re-designed to be just right, so that the brilliant lemmas that ensure its soundness hold and can be proven. When the DOT calculus is disrupted by a modification, it is difficult to predict which parts of the proof will be affected. Experimenting with modifications to DOT is difficult because each tweak requires many lemmas to be re-proven.

The goal of the first part of this thesis is a soundness proof that is simpler, more modular, and more intuitive. Such a proof separates the concepts of types, values, and operational semantics, and reasons about one concept at a time. Then, if a language extension modifies only one concept, such as typing, the necessary changes are localized to the parts of the proof that deal with types. We also aim to isolate most of the reasoning in a simpler system that is immune to the paradox of *bad bounds*, the key challenge that plagued the long search for a soundness proof. In this system, our reasoning can rely on intuitive notions from familiar object calculi without dependent object types (Pierce, 2002). The results of this reasoning are lifted to the full DOT type system by a single, simple theorem.

*the bad-bounds issue
is described
in Chapter 4*

*see Chapter 5 for the
simple DOT proof*

The main focus of our proof is on types. Dependent object types are the one feature that distinguishes DOT, so we aim to decouple that one feature, which mainly affects the static type system, from other concerns. We focus on proving the properties that one expects of types, and deliberately keep the proof independent of other aspects, such as operational semantics and runtime values, which are similar in DOT as in other object calculi. Of course, a soundness proof must eventually speak about execution and values, but once we have the necessary theory to reason about types, these other concerns can be handled separately, at the end of the proof, using standard proof techniques. Our final soundness theorem is stated for the small-step operational semantics given by Amin, Grütter, et al., (2016), but that is only the final conclusion; the theory that we develop about dependent object types would be equally applicable in a proof for a big-step operational semantics.

The power of DOT is also its curse. DOT empowers a program to define a domain-specific type system with a custom subtyping

lattice inside the existing Scala type system. This power has been used to encode in plain Scala expressive type systems that would otherwise require new languages to be designed. But this power also enables typing contexts that make no sense, in which types cannot be trusted and thus become meaningless. For example, a program could define typing contexts in which an object, which is not a function, nevertheless has a function type. Since such “crazy” contexts are possible, a soundness proof needs to consider them (but prove that they are harmless during execution).

Besides the general pursuit of modularity, the simplicity of our new proof depends on two main ingredients.

The first ingredient is *inert types* and *inert typing contexts*. The essential property of an inert type is that if all variables have inert types, then no unexpected subtyping relationships are possible, so types can be trusted, and none of the paradoxes are possible. An important part of the soundness proof is to ensure that a term cannot evaluate until the types of all its free variables have been narrowed to inert types.

*inertness is defined
in Section 5.2*

We define inertness as a concise, easily testable syntactic property of a type. The definition consists of only two non-recursive inference rules, so it can be easily inverted when it occurs in a proof. By contrast, existing DOT proofs achieve similar goals using properties that characterize types by the existence of values with specific relationships to those types. The benefit of our inertness property is that it involves only a type, not any values, and it is defined directly, not via existential quantification of some corresponding value.

The second ingredient is tight typing, a small restriction of the DOT typing rules with major consequences. We did not invent tight typing; it appears as a technical definition in the proof of Amin, Grütter, et al., (2016). Our contribution is to identify and demonstrate just how useful and important tight typing is to a simple proof. Amin, Grütter, et al., (2016) use tight typing in a collection of technical lemmas mixed with reasoning about other concerns, such as general typing (the full typing rules of DOT) and correspondences between values and types. In our proof, however, tight typing takes centre stage; it is the main actor that enables intuition and simplicity.

*tight typing is
discussed
in Section 5.3.*

Tight typing neutralizes the two DOT type rules that enable a program to define custom subtyping relationships. Tight typing immunizes the calculus: even if a typing context contains a type that is not inert, tight typing prevents it from doing any harm. The paradoxes that make it challenging to work with DOT disappear under tight typing. Without those two typing rules, the calculus behaves very differently, like object calculi without dependent object types, and our reasoning can rely on familiar properties that we are used to from these calculi.

Of course, DOT with tight typing is not at all the real DOT: it lacks the power to create customized type systems, and it is uninteresting;

Theorem 6 (\vdash to $\vdash_{\#}$)
is in Section 5.3

it is just another calculus with predictable behaviour. One simple theorem bridges the gap by showing that in inert contexts, tight typing has all the power of general typing. Therefore, all the reasoning that we do in the intuitive environment of tight typing applies to the full power of DOT. Even the proof of the theorem itself reasons entirely with tight typing, without having to deal with the paradoxes of general DOT typing, and without having to reason about relationships between types and values.

Combining these two ingredients, we contribute a unified general recipe that can be used whenever a proof about DOT needs to deduce information about a term from its type. Many of our lemmas follow this recipe. The first step of the recipe, which should be the first step of any reasoning about types in DOT, is to drop down from general typing to tight typing using the above theorem. The purpose of the remaining steps is to make inductive reasoning as easy and systematic as possible.

CONTRIBUTIONS

This part of the thesis presents a simplified and extensible soundness proof for the DOT calculus. It contributes the following:

Chapter 6 describes
how to extend the
simple proof

- A *modular* and *extensible* proof that reasons about types, values, and operational semantics separately.
- The concept of *inert* typing contexts, a syntactic characterization of contexts that rule out any non-sensical subtyping that could be introduced by abstract type members.
- A simple *proof recipe* for deducing properties of terms from their types in full DOT while reasoning only in a restricted, intuitive environment free from the paradoxes caused by abstract type members. Multiple lemmas follow the same recipe, and following the recipe can facilitate the development of new lemmas needed in future extensions for DOT.

Coq proof:
[git.io/simple-dot-](https://git.io/simple-dot-proof)
proof

- A *Coq formalization* of the simple DOT soundness proof.

BAD BOUNDS

The type selection subtyping rules <:-SEL and SEL-< enable users to define a type system with a custom subtyping lattice. If a program defines a function $\lambda(x: \{A: S..U\}) t$, then t is typed in a context in which S is considered a subtype of U , because $S \text{<:-} x.A \text{<:-} U$. The soundness proof must ensure that such a user-defined subtyping lattice does not cause any harm, i.e., cannot cause a violation of type soundness of the overall calculus.

Let S be the object type $\{a: \top\}$ and U be the function type $\forall(z: \top) \top$. Then the following is a valid and well-typed DOT term:

$$\lambda(x: \{A: S..U\}) \text{let } y = v(y: S) \{a = y.a\} \text{ in } y \ y$$

How is this possible? The inner term $y \ y$ is a function application applying y to itself, but y is bound by the `let` to an object, not a function. How can y appear in a function application when it is not a function? This is possible because y has the object type S , and in the body of the lambda, we have the subtyping chain $S \text{<:-} x.A \text{<:-} U$. The declaration of the lambda asserts that $x.A$ is a supertype of S and a subtype of U , and therefore introduces the new custom subtyping relationship $S \text{<:-} U$. Inside the body of the lambda, the object type S is a subtype of the function type U , so since the object y has type S , it also has the function type U . The function application of object y to itself is therefore well-typed in this context.

This is crazy, the reader may be thinking. Indeed, in an environment in which subtyping can be arbitrarily redefined, types cannot be trusted. In particular, we cannot conclude from the fact that y has the function type S that it is indeed a function; actually, it is an object. The seemingly obvious fix is to require S to be a subtype of U when the parameter x of the lambda is declared to have type $\{A: S..U\}$. But as we will discuss in Chapter 7, this seemingly obvious fix does not work, and the struggle to try to make it work has caused much of the difficulty in the ten-year struggle for a DOT soundness proof.

How can DOT be sound then, when it is so crazy? After all, the function application $y \ y$ is well-typed but its evaluation gets stuck, because y is not a function, so how can DOT be sound? The key is that the DOT semantics is call-by-value. In order to invoke the body of the lambda, one must provide an argument value to pass for the parameter x . This value must contain a type assigned to A that is both a supertype of U and a subtype of S . If no such type exists, then no such argument value can exist, so the lambda cannot be called, so its body containing the crazy application $y \ y$ cannot ever be executed.

$$\frac{\Gamma \vdash x: \{A: S..T\}}{\Gamma \vdash S \text{<:-} x.A} \quad (\text{<:-SEL})$$

$$\frac{\Gamma \vdash x: \{A: S..T\}}{\Gamma \vdash x.A \text{<:-} T} \quad (\text{SEL-<})$$

$$S = \{a: \top\}$$

$$U = \forall(z: \top) \top$$

Therefore, this term is not a counterexample to the soundness of the DOT type system.

Why should DOT have such a strange feature? The ability to define a custom subtyping lattice turns out to be very useful. For example, we can define the term:

$$\lambda(x: \{A: \perp..T\} \wedge \{B: x.A..x.C\} \wedge \{C: \perp..T\}) t$$

In the body t of this lambda, we can make use of unspecified opaque types A , B , and C , making use of only the condition that $A <: B <: C$. We can use this feature to define arbitrary type systems within the language. For example, Scalas and Yoshida, (2016) have implemented session types, a feature that usually requires a custom-designed language, inside plain Scala. As another example, Osvald et al., (2016) used this ability to define a lattice of lifetimes within the Scala type system for categorizing values that cannot outlive different stack frames. Even the well-known Scala cake pattern (Odersky and Zenger, 2005b) is built using this feature.

To reconcile a custom subtyping lattice with a sound language, we only need to force the programmer to provide evidence that the custom lattice does not violate any familiar assumptions (e.g., it does not make object types subtypes of function types). This evidence takes the form of an argument value that must be passed to the lambda before the body that uses the custom type lattice can be allowed to execute. This value must be an object that provides existing types that satisfy the specified custom subtyping constraints. In our example, this is easy: it suffices to pass the same type, such as T , for all three type parameters, since $T <: T <: T$. However, the types are opaque: when checking the body of the lambda, the type checker cannot use the fact that $A = B = C = T$; the body must type-check even under only the assumptions that $A <: B <: C$.

Since DOT programs can exhibit unexpected subtyping lattices in some contexts, and since this is unavoidable, an essential feature of a soundness proof is to clearly distinguish contexts in which types can be trusted, because any custom subtyping relationships have been justified by actual type arguments, from contexts in which types cannot be trusted, because they could have been derived from arbitrary unjustified custom subtyping relationships. In Section 5.2, we will formally define this property that types can be trusted, and define a simple syntactic characterization of *inert* typing contexts that guarantee this property. In earlier DOT soundness proofs, the trusted types property was not precisely defined, and typing contexts in which there are no bad bounds were defined more indirectly, not in terms of the types themselves, but in terms of the existence of values having those types.

THE SIMPLE DOT PROOF

This chapter describes how inert contexts and three auxiliary typing judgments (tight, invertible, and precise typing) yield a simple proof recipe that can be used in the DOT soundness proof whenever one needs to make sense of types.

5.1 OVERVIEW

We will first outline the general recipe that we use to reason throughout the proof about the meaning of a type. The details of each step will be discussed in the following subsections. We present the overview on an example proof of Lemma 13, which will be introduced in Section 5.5, but the specific example is unimportant; most of the reasoning throughout the proof follows the same steps, through the same typing relations, in the same order, using the same reasoning techniques.

Usually, we know that some term has some type (e.g. $\Gamma \vdash x: \{a: T\}$), and we seek to interpret what the type tells us about the term, and to determine how the type of the term was derived. In this example, we seek more detailed information about x , for example that the typing context Γ assigns it an object type $\Gamma(x) = \mu(x: \dots \wedge \{a: T'\} \wedge \dots)$, or the shape of the value that it will hold at run time (e.g. an object $\nu(x: \dots \wedge \{a: T'\} \wedge \dots)(\dots \wedge \{a = t'\} \wedge \dots)$).

Each such derivation follows the same sequence of steps (although sometimes only a subsequence of the steps is necessary):

$$\begin{array}{c}
 \frac{\Gamma \vdash x: \{a: T\}}{\Gamma \vdash_{\#} x: \{a: T\}} \text{ THEOREM 6 } (\vdash \text{ TO } \vdash_{\#}) \\
 \frac{\Gamma \vdash_{\#} x: \{a: T\}}{\Gamma \vdash_i x: \{a: T\}} \text{ THEOREM 10 } (\vdash_{\#} \text{ TO } \vdash_i) \\
 \frac{\Gamma \vdash_i x: \{a: T'\} \quad \Gamma \vdash T' <: T}{\Gamma(x) = \mu(x: \dots \wedge \{a: T'\} \wedge \dots) \quad \Gamma \vdash T' <: T} \text{ INDUCTION ON } \vdash_i \\
 \frac{\Gamma(x) = \mu(x: \dots \wedge \{a: T'\} \wedge \dots) \quad \Gamma \vdash T' <: T}{\Gamma \vdash T' <: T} \text{ INVERSION OF } \vdash_i
 \end{array}$$

where we assume that Γ is inert. Although there are four steps, each individual step is quite simple. More importantly, each step is modular, independent of the other steps, and the proof techniques at each step are either directly reusable (theorems) or easily adaptable (induction) to proofs of properties other than this specific lemma.

The derivation starts with general typing ($\Gamma \vdash x: \{a: T\}$), the typing relation of the DOT calculus. The key property that makes reasoning possible is that the typing context Γ is inert. Inert contexts will be defined in Section 5.2. Inertness ensures that customized subtyping in

general typing
 $\Gamma \vdash t: T$

inert contexts

the program does not introduce unexpected subtyping relationships. If the context were not inert, any type could have been customized to have arbitrary subtypes and be inhabited by arbitrary terms, so it would be impossible to draw any conclusions about a term from its type.

Knowing that the typing context is inert, we apply Theorem 6 (\vdash to $\vdash_{\#}$) to get a tight typing ($\Gamma \vdash_{\#} x : \{a : T\}$), which will be discussed in Section 5.3. A tight typing is immune to any unexpected subtyping relationships that the program may have defined, so our reasoning can now rely on familiar intuitions about what types ought to mean about their terms.

However, the tight typing rules are not amenable to inductive proofs. Theorem 10 ($\vdash_{\#}$ to \vdash_i) gives invertible typing ($\Gamma \vdash_i x : \{a : T\}$), which is specifically designed to make inductive reasoning as easy as possible. Invertible typing will be discussed in Section 5.4.

By induction on invertible typing, we obtain a property of all of the precise types $\Gamma \vdash_i x : \{a : T'\}$ that could have caused x to have the general type $\{a : T\}$. Informally, the precise typing means that the type $\Gamma(x)$ given to x by the typing context is an object type containing a field a of type T' . We will present precise typing in Section 5.3. Precise typing is also amenable to straightforward induction proofs, so we can use one to obtain $\Gamma(x)$.

5.2 INERT TYPING CONTEXTS

Recall the function $\lambda(x : \{A : S..U\}) t$ that we discussed in Chapter 4. If the function appears in a context Γ , its body is type checked in an extended context $\Gamma, x : \{A : S..U\}$. The extended context adds a new subtyping relationship $\Gamma, x : \{A : S..U\} \vdash S <: U$ that might not have held in the original context Γ . In particular, the extended context could introduce a subtyping relationship that does not make sense, such as $\forall(x : S) T <: \mu(x : U)$, or $\top <: \perp$. To control such unpredictable contexts, we define the notion of inert typing contexts and inert types. Inert types are defined through *record types* – records whose type members have equal bounds.

record $\{a : T\}$
record $\{A : U..U\}$

DEFINITION 1 (Record Types). A record type is an intersection of types each of which is either a field declaration $\{a : T\}$ or a tight type declaration $\{A : U..U\}$.

record T record U

record $T \wedge U$

DEFINITION 2 (Inert Types). A type U is inert if it is either a function type or a recursive type $\mu(x : T)$ where T is a record type.

inert $\forall(x : T) U$

DEFINITION 3 (Inert contexts). A typing context Γ is inert if the type $\Gamma(x)$ that it assigns to each variable x is inert.

record T

inert $\mu(x : T)$

An inert typing context has the following useful property.

PROPERTY 4 (Inert Context Guarantee). *Let Γ be any inert typing context, t be a closed term and U be a closed type. If $\Gamma \vdash t : U$, then $\vdash t : U$.*

The significance of this property is that in an inert typing context, a term t does not have any “unexpected” types that it would not have in an empty typing context. For example, we can be sure that in an inert typing context, a function value will not have an object (recursive) type, and an object will not have a function type. Though we do not directly apply the property in the proof, it is useful for intuitive reasoning about typing and subtyping in inert typing contexts.

Every value has an inert type (as long as the value is well formed, i.e., as long as it has any type at all). This is because the two base typing rules for values, ALL-I and {}-I, and the definition typing rules that they depend on, always assign an inert type to the value. The converse is not true: not every inert type is inhabited by a value. For example, we cannot construct a value of type $\lambda(x : \top) \perp$.

Returning to the example, suppose now that the function is invoked with some value v bound to a variable y :

$$\text{let } y = v \text{ in } (\lambda(x : \{A : S..U\}) t) y.$$

Recall that the body t is typed with the assumption that $S <: U$. Type checking the overall term ensures that the argument y provides *evidence* for that assumption. Specifically, the value v has an inert type, so y has an inert type. The typing rule for function application requires subtyping between the argument and parameter types, so the type of y must have a member $\{A : T..T\}$ with $S <: T$ and $T <: U$. (The bounds T of the type member must be tight because the type is inert.) The type T that y provides is evidence that justifies the assumption $S <: U$ under which the body t of the function was type checked. During execution, when the function is called, all occurrences of x in the body t will be replaced by y before evaluation of the body begins. In general, the semantics ensures that before it begins evaluating a term (such as t), the term has a type in a context in which all non-inert types (such as the type of x) have been narrowed to inert types (such as the type of y).

5.3 TIGHT TYPING

Although inert contexts provide the assurance of Property 4 (Inert Context Guarantee), in our proofs, we often need to reason even in contexts that are not inert. Moreover, even when we know that a context is inert, it would be difficult to express the important consequences of the inert context in every proof that deals with the general DOT typing and subtyping rules.

Tight typing (Amin, Grütter, et al., 2016) is a slight restriction of general typing that can bridge the gap between the unpredictability

Tight term typing
 $\Gamma \vdash_{\#} t : T$

$$\begin{array}{c}
\frac{\Gamma(x) = T}{\Gamma \vdash_{\#} x : T} \quad (\text{VAR-}\#) \\
\\
\frac{\Gamma, x : T \vdash t : U \quad x \notin \text{fv}(T)}{\Gamma \vdash_{\#} \lambda(x : T) t : \forall(x : T) U} \quad (\text{ALL-I-}\#) \\
\\
\frac{\Gamma \vdash_{\#} x : \forall(z : S) T \quad \Gamma \vdash_{\#} y : S}{\Gamma \vdash_{\#} x y : T[y/z]} \quad (\text{ALL-E-}\#) \\
\\
\frac{\Gamma, x : T; d \vdash T :}{\Gamma \vdash_{\#} \nu(x : T) d : \mu(x : T)} \quad (\text{I-}\#) \\
\\
\frac{\Gamma \vdash_{\#} x : \{a : T\}}{\Gamma \vdash_{\#} x.a : T} \quad (\text{FLD-E-}\#) \\
\\
\frac{\Gamma \vdash_{\#} t : T \quad \Gamma \vdash_{\#} T <: U}{\Gamma \vdash_{\#} t : U} \quad (\text{SUB-}\#)
\end{array}$$

Tight subtyping
 $\Gamma \vdash_{\#} T <: U$

$$\begin{array}{c}
\Gamma \vdash_{\#} T <: \top \quad (\text{TOP-}\#) \\
\\
\Gamma \vdash_{\#} \perp <: T \quad (\text{BOT-}\#) \\
\\
\Gamma \vdash_{\#} T <: T \quad (\text{REFL-}\#) \\
\\
\frac{\Gamma \vdash_{\#} S <: T \quad \Gamma \vdash_{\#} T <: U}{\Gamma \vdash_{\#} S <: U} \quad (\text{TRANS-}\#) \\
\\
\Gamma \vdash_{\#} T \wedge U <: T \quad (\text{AND}_1\text{-<:-}\#) \\
\\
\Gamma \vdash_{\#} T \wedge U <: U \quad (\text{AND}_2\text{-<:-}\#) \\
\\
\frac{\Gamma \vdash_{\#} S <: T \quad \Gamma \vdash_{\#} S <: U}{\Gamma \vdash_{\#} S <: T \wedge U} \quad (\text{<:-AND-}\#) \\
\\
\frac{\Gamma \vdash_{\#} T <: U}{\Gamma \vdash_{\#} \{a : T\} <: \{a : U\}} \quad (\text{FLD-<:-FLD-}\#) \\
\\
\frac{\Gamma \vdash_{\#} S_2 <: S_1 \quad \Gamma \vdash_{\#} T_1 <: T_2}{\Gamma \vdash_{\#} \{A : S_1..T_1\} <: \{A : S_2..T_2\}} \quad (\text{TYP-<:-TYP-}\#) \\
\\
\frac{\Gamma \vdash_{\#} S_2 <: S_1 \quad \Gamma, x : S_2 \vdash T_1 <: T_2}{\Gamma \vdash_{\#} \forall(x : S_1) T_1 <: \forall(x : S_2) T_2} \quad (\text{ALL-<:-ALL-}\#)
\end{array}$$

Figure 5.1: Tight Typing Rules
(Amin, Grütter, et al., 2016)

$$\begin{array}{c}
\frac{\Gamma(x) = T}{\Gamma \vdash! x : T} \text{ (VAR-!)} \quad \frac{\Gamma \vdash! x : \mu(z : T)}{\Gamma \vdash! x : T[x/z]} \text{ (REC-E-!)} \quad \frac{\Gamma \vdash! x : T \wedge U}{\Gamma \vdash! x : T} \text{ (AND}_1\text{-E-!)} \quad \frac{\Gamma \vdash! x : T \wedge U}{\Gamma \vdash! x : U} \text{ (AND}_2\text{-E-!)} \\
\text{Precise variable typing} \\
\Gamma \vdash! x : T
\end{array}$$

$$\begin{array}{c}
\frac{\Gamma, x : T \vdash t : U \quad x \notin \text{fv}(T)}{\Gamma \vdash! \lambda(x : T) t : \forall(x : T) U} \text{ (ALL-I-!)} \quad \frac{\Gamma, x : T; d \vdash T :}{\Gamma \vdash! \nu(x : T) d : \mu(x : T)} \text{ ({}-I-!)} \\
\text{Precise value typing} \\
\Gamma \vdash! v : T
\end{array}$$

Figure 5.2: Simplified Precise Typing Rules based on (Amin, Grütter, et al., 2016)

of the general DOT typing rules in arbitrary typing contexts and the predictable assurances of Property 4 in inert typing contexts. The tight typing rules are presented in Figure 5.1. They are almost the same as the general DOT typing rules, except that the $\text{<:-SEL-}\#$ and $\text{SEL-<:-}\#$ rules have the restricted premise $\Gamma \vdash! x : \{A : T..T\}$, so they can be applied only when the bounds T of the type member A are tight. Precise typing, denoted $\vdash!$, is defined in Figure 5.2. The precise type of a variable x is the type $\Gamma(x)$ given to it by the typing context Γ , possibly decomposed using the elimination rules, so that if $\Gamma(x)$ is an object type such as $\mu(x : \dots \wedge \{A : T..T\} \wedge \dots)$, then x also has just the type member $\{A : T..T\}$ as a precise type. For values, precise typing applies only the base case rules ALL-I and {}-I from general typing. In premises of rules that extend the typing context (ALL-I-#, Let-#, {}-I-#), tight typing reverts to general typing in the extended context.

We observe two useful properties of tight typing that together combine to make it especially convenient for reasoning about DOT typing. The first property is that tight typing extends the benefits of Property 4 (Inert Context Guarantee) to *all* typing contexts, not only inert ones:

PROPERTY 5 (Tight Typing Guarantee). *Let Γ be any typing context, t be a closed term and U be a closed type. If $\Gamma \vdash_{\#} t : U$, then $\vdash_{\#} t : U$ and $\vdash t : U$.*

The general typing rules that enable DOT programs to define new user-defined subtyping relationships, <:-SEL and SEL-<:- , are restricted in tight typing to $\text{<:-SEL-}\#$ and $\text{SEL-<:-}\#$, which allow only to give an alias to an existing type, but not to introduce new subtyping between existing types.

Property 5 makes reasoning in tight typing easy: we never have to worry about unexpected custom subtyping relationships being introduced by the program, and we do not need to reason about whether we are in an inert typing context, because tight typing gives the guarantee in all contexts.

$$\begin{array}{c}
\frac{\Gamma \vdash x : \{A : S..T\}}{\Gamma \vdash S <: x.A <: T} \text{ (<:-SEL, SEL-<:-)} \\
\\
\frac{\Gamma \vdash! x : \{A : T..T\}}{\Gamma \vdash_{\#} T <: x.A <: T} \text{ (<:-SEL-}\#, \text{SEL-<:-}\#)
\end{array}$$

Although tight typing satisfies the desirable intuitive Property 5, it is not DOT. In particular, tight typing does not, in general, enable a program to use a custom-defined subtyping lattice that is the key feature of dependent object types. We would like the best of both worlds: to allow DOT programs to enjoy the full power of general typing, yet to reason about our proofs with the intuitive tight typing. For this, we need the second property of tight typing.

The second important property of tight typing is that in an inert typing context, tight typing is equivalent to general DOT typing:

THEOREM 6 (\vdash to $\vdash_{\#}$). *If Γ is an inert context, then $\Gamma \vdash t : T$ implies $\Gamma \vdash_{\#} t : T$, and $\Gamma \vdash S <: U$ implies $\Gamma \vdash_{\#} S <: U$.*

$$\frac{\text{inert } \Gamma \quad \Gamma \vdash t : T}{\Gamma \vdash_{\#} t : T} \quad (\vdash \text{ TO } \vdash_{\#})$$

The theorem statements in side notes as a quick overview or refresher and might exclude some details, such as existential quantifiers.

We delay giving the proof of the theorem until after some discussion.

These two properties motivate and justify our recommendation that tight typing should be at the core of all reasoning about the meaning of types in DOT. Tight typing is predictable, like the type systems of familiar calculi without dependent object types, yet in an inert typing context, it has the same power as general DOT typing. Therefore, every proof with a premise involving general typing and an inert typing context should immediately apply Theorem 6 (\vdash to $\vdash_{\#}$) to drop down into the intuitive environment of tight typing for the rest of the reasoning.

What if we do not have an inert context as a premise, and therefore cannot apply Theorem 6? In that case, we should not reason about the meanings of types at all. As we saw in Chapter 4, in such a context, a term could be given an arbitrary type by custom subtyping rules. Therefore, we cannot deduce anything about a term from its type, and it would be futile to try.

In summary, inert contexts, tight typing, and Theorem 6 that justifies reasoning in tight typing should be the cornerstones of any reasoning about the meaning of types in the DOT calculus.

How shall we prove Theorem 6, then? It is tempting to prove the theorem by trying to compare various properties of the tight and general typing *relations*, the closures of the tight and general typing *rules*. This approach was taken in the proof of Amin, Grütter, et al., (2016) for a related theorem (with the same conclusion but different premises). The typing relations are very different from each other (general typing is much more powerful), but the rules that give rise to them are quite similar. It is much easier, therefore, to instead show that the *rules* are equivalent in an inert context. The only rules in general typing missing from tight typing are the $<:-\text{SEL}$ and $\text{SEL}<:-$ rules. Our goal is therefore to replace these rules with a lemma:

LEMMA 7 ($\text{SEL}<:-\#$ Replacement). *If Γ is an inert context, then if $\Gamma \vdash_{\#} x : \{A : S..U\}$, then $\Gamma \vdash_{\#} S <: x.A$ and $\Gamma \vdash_{\#} x.A <: U$.*

$$\frac{\text{inert } \Gamma \quad \Gamma \vdash_{\#} x : \{A : S..U\}}{\Gamma \vdash_{\#} S <: x.A <: U} \quad (\text{SEL}<:-\# \text{ REPLACEMENT})$$

One nice property of this lemma is that it is stated entirely in terms of tight typing. Thus, to prove it, we can ignore the unpredictable

world of general typing, and work exclusively in the intuitive world of tight typing.

But how can we prove it? We would like to apply the $<:-\text{SEL-}\#$ and $\text{SEL-}<:-\#$ rules. Their premises are $\Gamma \vdash! x: \{A: T..T\}$. Therefore, we need to *invert* tight typing, to show the following:

LEMMA 8 ($\text{SEL-}<:-\#$ Premise). *If Γ is an inert context, then if $\Gamma \vdash_{\#} x: \{A: S..U\}$, then there exists a type T such that $\Gamma \vdash! x: \{A: T..T\}$, $\Gamma \vdash_{\#} S <: T$, and $\Gamma \vdash_{\#} T <: U$.*

$$\frac{\text{inert } \Gamma \quad \Gamma \vdash_{\#} x: \{A: S..U\}}{\Gamma \vdash! x: \{A: T..T\}} \quad (\text{SEL-}<:-\# \text{ PREMISE})$$

We will discuss how to invert tight typing to prove this lemma in Section 5.4.

Using Lemma 8, proving Lemma 7 ($\text{SEL-}<:-\#$ Replacement) is easy:

Proof of Lemma 7. Apply Lemma 8, then $<:-\text{SEL-}\#$ and $\text{SEL-}<:-\#$, to get $\Gamma \vdash_{\#} S <: T <: x.A <: T <: U$. The result follows by $\text{TRANS-}\#$. \square

Using Lemma 7, proving Theorem 6 (\vdash to $\vdash_{\#}$) is now also quite easy.

Proof of Theorem 6. The proof is by mutual induction on the tight typing and subtyping derivations of $\Gamma \vdash t: T$ and $\Gamma \vdash S <: U$. In general, for each rule of general typing, we invoke the corresponding rule of tight typing. The premises of the tight typing rules differ from those of the general typing rules in that they require tight typing in rules that do not extend the context. Since the unextended context is inert, the general premise implies the tight premise by the induction hypothesis. Premises that do extend the context use general typing, so nothing needs to be proven for them. The exception is the $<:-\text{SEL}$ and $\text{SEL-}<:-$ rules. Lemma 7 is an exact replacement for these rules, so we just apply it. Despite the long explanation, the proof in Coq is only two lines long. \square

5.4 INVERSION OF TIGHT TYPING

Although reasoning with tight typing is intuitive because it obeys Property 5 ($\text{Tight Typing Guarantee}$), we often need to invert the tight typing rules to prove properties such as Lemma 8, which we used in the proof of Lemma 7. More generally, we need to prove that if $\Gamma \vdash_{\#} x: T$, where T is of a certain form, then $\Gamma(x) = U$, and there is a certain relationship between T and U .

The obvious approach to proving such inversion properties is by induction on the derivation of the tight typing. This usually fails, however, because of cycles in the tight typing rules. Each language construct typically has both an introduction and an elimination rule, and the two form a cycle. For example, if $\Gamma \vdash_{\#} x: T$, then $\Gamma \vdash_{\#} x: \mu(x: T)$ by $\text{REC-I-}\#$, so again $\Gamma \vdash_{\#} x: T$ by $\text{REC-E-}\#$. Such cycles block inductive proofs because a proposition $\Gamma \vdash_{\#} x: T$ is justified by

$$\frac{\Gamma \vdash_{\#} x: T}{\Gamma \vdash_{\#} x: \mu(x: T)} \quad (\text{REC-I-}\#)$$

$$\frac{\Gamma \vdash_{\#} x: \mu(z: T)}{\Gamma \vdash_{\#} x: T[x/z]} \quad (\text{REC-E-}\#)$$

$\Gamma \vdash_{\#} x : \mu(x : T)$, which in turn is justified by the original proposition $\Gamma \vdash_{\#} x : T$. The solution is to define a set of acyclic, invertible rules on which induction is easy, and to prove that the invertible rules induce the same typing relation as the cyclic tight typing rules.

The construction of the invertible typing rules is simplified by two restrictions:

1. We only ever need to invert typing rules in inert typing contexts.
2. We only ever need to invert typings of variables and values, not of arbitrary terms.

In the invertible rules, we can thus exclude rules that cannot apply to variables or values, and rules that cannot apply to inert types or to types derived from inert types.

It remains to decide, when facing a cycle of two rules that introduce and eliminate a given language construct, which one of the two rules to remove and which one to keep in the acyclic, invertible rule set. In general, because a construct can be introduced an unbounded number of times in tight typing, we must keep the introduction rule. For example, if x has type T , then x also has type $\mu(y : \mu(y' : \mu(y'' : T)))$, and the invertible rules must generate this type. On the other hand, the base case of the typing rules for variables, the rule VAR-#, gives each variable x the type $\Gamma(x)$, which in an inert context is an inert type, and can therefore be a recursive type containing an intersection type. Since the tight typing rules eliminate the recursion and the intersection, the invertible rules must also eliminate them. It seems that we have reached a contradiction: the invertible rules must have both introduction and elimination rules for recursive and intersection types.

The solution is to split the invertible rules into two phases. The first phase of rules contains all the elimination rules. After all necessary eliminations have been performed, a second phase containing only introduction rules can then perform all necessary introductions. By splitting the rules into two phases, we ensure that no derivation can cycle between introductions and eliminations, so the rules are invertible. It turns out that we already have rules for the first phase: the precise typing rules introduced in Section 5.3 already contain all of the elimination rules that apply to variables and values, and eliminate from the type of a variable all constructs that can appear in an inert type.² To construct the invertible introduction rules, we propose the following recipe:

² Note that even the general DOT typing rules remove recursive and intersection types only from the types of variables, not values.

1. Start with the tight typing rules.
2. Inline the subsumption rule (inline the subtyping rules into the typing rules). This simplifies the construction, so we define only one relation instead of two separate typing and subtyping relations.

$\frac{\Gamma \vdash_{\#} x : T}{\Gamma \vdash_i x : T} \quad (\text{VAR-}i)$	$\frac{\Gamma \vdash_i x : \forall(z : S) T \quad \Gamma \vdash_{\#} S' <: S}{\Gamma, y : S' \vdash T <: T'} \quad (\text{ALL-}<:-\text{ALL-}i)$	<i>Invertible variable typing</i> $\Gamma \vdash_i x : T$
$\frac{\Gamma \vdash_i x : \{a : T\} \quad \Gamma \vdash_{\#} T <: U}{\Gamma \vdash_i x : \{a : U\}} \quad (\text{FLD-}<:-\text{FLD-}i)$	$\frac{\Gamma \vdash_i x : T \quad \Gamma \vdash_i x : U}{\Gamma \vdash_i x : T \wedge U} \quad (\text{AND-I-}i)$	
$\frac{\Gamma \vdash_{\#} T' <: T \quad \Gamma \vdash_{\#} U <: U'}{\Gamma \vdash_i x : \{A : T'..U'\}} \quad (\text{TYP-}<:-\text{TYP-}i)$	$\frac{\Gamma \vdash_i x : S \quad \Gamma \vdash_i y : \{A : S..S\}}{\Gamma \vdash_i x : y.A} \quad (\text{SEL-}i)$	
$\frac{\Gamma \vdash_i x : T}{\Gamma \vdash_i x : \mu(x : T)} \quad (\text{REC-I-}i)$	$\frac{\Gamma \vdash_i x : T}{\Gamma \vdash_i x : \top} \quad (\text{TOP-}i)$	
$\frac{\Gamma \vdash_{\#} v : T}{\Gamma \vdash_i v : T} \quad (\text{VAL-}iv)$	$\frac{\Gamma \vdash_i v : T \quad \Gamma \vdash_i v : U}{\Gamma \vdash_i v : T \wedge U} \quad (\text{AND-I-}iv)$	<i>Invertible value typing</i> $\Gamma \vdash_i v : T$
$\frac{\Gamma \vdash_i v : \forall(z : S) T \quad \Gamma \vdash_{\#} S' <: S}{\Gamma, y : S' \vdash T <: T'} \quad (\text{ALL-}<:-\text{ALL-}iv)$	$\frac{\Gamma \vdash_i v : S \quad \Gamma \vdash_i y : \{A : S..S\}}{\Gamma \vdash_i v : y.A} \quad (\text{SEL-}iv)$	
	$\frac{\Gamma \vdash_i v : T}{\Gamma \vdash_i v : \top} \quad (\text{TOP-}iv)$	

Figure 5.3: Invertible typing rules

3. Specialize the terms in all rules to variables and values, and remove all rules that cannot apply to variables or values.
4. Remove all elimination rules.
5. Remove all rules that cannot apply in an inert context. Specifically, this means the BOT-# rule, because it has $\Gamma \vdash_{\#} x : \perp$ as a premise, but this typing cannot be derived by any of the other remaining rules starting from an inert type given to a variable by the VAR-# rule or to a value by the ALL-I-# and {}-I-# rules.

By applying this recipe to the tight typing rules, we arrive at the invertible typing rules shown in Figure 5.3. We must now prove that the typing relation induced by the invertible typing rules is equal to the typing relation induced by the tight typing rules (restricted to inert

contexts and to variables and values). To do this, we need to first show that invertible typing is closed under tight subtyping.

$$\frac{\text{inert } \Gamma \quad \Gamma \vdash_i x : T}{\Gamma \vdash_{\#} T <: U}$$

$\Gamma \vdash_i x : U$
(INVERTIBLE <: CLOSURE)

LEMMA 9 (Invertible <: Closure). *If Γ is inert, $\Gamma \vdash_i x : T$, and $\Gamma \vdash_{\#} T <: U$ then $\Gamma \vdash_i t : U$.*

We can now prove that tight typing implies invertible typing:

$$\frac{\text{inert } \Gamma \quad \Gamma \vdash_{\#} t : T}{\Gamma \vdash_i t : T} \quad (\vdash_{\#} \text{ TO } \vdash_i)$$

THEOREM 10 ($\vdash_{\#}$ to \vdash_i). *If Γ is an inert context and $\Gamma \vdash_{\#} x : T$, then $\Gamma \vdash_i x : T$.*

Proof. The proof is by induction on the derivation of $\Gamma \vdash_{\#} x : T$. Although we said that induction on tight typing usually fails because the rules have cycles, in this specific case, the induction is quite straightforward because invertible typing is part of the induction hypothesis. The inductive cases for elimination rules, which would usually lead to cycles in the induction, are all discharged using the invertible typing in the induction hypothesis. \square

The soundness proof also has versions of the above two lemmas for values (instead of variables) but we omit them here because they are similar.

With this theorem, inversion proofs such as the proof of Lemma 8 (SEL-<:-# PREMISE) become easy inductions on the invertible typing rules:

$$\frac{\text{inert } \Gamma \quad \Gamma \vdash_{\#} x : \{A : S..U\}}{\Gamma \vdash_i x : \{A : T..T\}} \quad (\text{SEL-<:-# PREMISE})$$

Proof of Lemma 8.

$$\frac{\frac{\text{inert } \Gamma \quad \Gamma \vdash_{\#} x : \{A : S..U\}}{\text{inert } \Gamma \quad \Gamma \vdash_i x : \{A : S..U\}} \text{ THEOREM 10 } (\vdash_{\#} \text{ TO } \vdash_i)}{\text{inert } \Gamma \quad \Gamma \vdash_i x : \{A : T..T\} \quad \Gamma \vdash_{\#} S <: x.T <: U} \text{ INDUCTION ON } \vdash_{\#}$$

\square

We will see more lemmas that follow the same proof strategy in the next section.

5.5 CANONICAL FORMS LEMMAS

In general, soundness proofs require canonical-forms lemmas that show that if a value has a given type, then it is a particular form of value. Following our theme of a modular proof that deals with one concept at a time, we do most of our work at the level of types, following the same general recipe.

Because the DOT syntax enforces ANF, before a value can be used for anything interesting, it must first be assigned to a variable through a let expression. Suppose a variable x is bound to a value v by **let** $x = v$ **in** t and the variable x is used somewhere inside t . From the type U of the use of x , we would like to deduce the form of the value v .

We proceed in two steps. First, from a type U such that $\Gamma' \vdash x: U$, where Γ' is the typing context used to type the use of x occurring inside t , we follow the proof recipe to deduce the type $\Gamma'(x)$ given to x by the typing context. The typing context Γ' is constructed by the premises of the LET typing rule, which extends an existing typing context Γ to the typing context Γ' by adding a binding $(x: T)$. Here, T is some type such that $\Gamma \vdash v: T$. Therefore, $\Gamma'(x)$ is this T , and we have, in general, that $\Gamma \vdash v: \Gamma'(x)$ and thus also $\Gamma' \vdash v: \Gamma'(x)$.

For the second step, we know $\Gamma' \vdash v: T$, where the type T has been identified by the first step, and we wish to deduce the precise type of v , and thence invert the precise value typing rules to obtain the form of v .

The following lemmas instantiate these two steps, first for dependent function types, and then for field member types.

LEMMA 11 (\forall to $\Gamma(x)$).

$$\frac{\text{inert } \Gamma \quad \Gamma \vdash z: \forall(x: T) U}{\Gamma(z) = \forall(x: T') U' \quad \Gamma \vdash T <: T' \quad \Gamma, x: T \vdash U' <: U}$$

LEMMA 12 (\forall to λ).

$$\frac{\text{inert } \Gamma \quad \Gamma \vdash v: \forall(x: T) U}{v = \lambda(x: T') t \quad \Gamma \vdash T <: T' \quad \Gamma, x: T \vdash t: U}$$

LEMMA 13 (μ to $\Gamma(x)$).

$$\frac{\text{inert } \Gamma \quad \Gamma \vdash x: \{a: T\}}{\Gamma(x) = \mu(x: \dots \wedge \{a: T'\} \wedge \dots) \quad \Gamma \vdash T' <: T}$$

LEMMA 14 (μ to ν).

$$\frac{\text{inert } \Gamma \quad \Gamma \vdash v: \mu(x: S) \quad S = \dots \wedge \{a: T\} \wedge \dots}{v = \nu(x: S)(\dots \wedge \{a = t\} \wedge \dots) \quad \Gamma \vdash t: T}$$

The proofs of all of the lemmas follow the same general proof recipe that we introduced for Lemma 13 in Section 5.1.

Proof of Lemma 12 (\forall to λ).

$$\begin{array}{c}
\frac{\text{inert } \Gamma \quad \Gamma \vdash v: \forall(x: T) U}{\text{inert } \Gamma \quad \Gamma \vdash_{\#} v: \forall(x: T) U} \text{THEOREM 6 } (\vdash \text{ TO } \vdash_{\#}) \\
\frac{\text{inert } \Gamma \quad \Gamma \vdash_{\#} v: \forall(x: T) U}{\text{inert } \Gamma \quad \Gamma \vdash_i v: \forall(x: T) U} \text{THEOREM 10 } (\vdash_{\#} \text{ TO } \vdash_i) \\
\hline
\frac{\text{inert } \Gamma \quad \Gamma \vdash_i v: \forall(x: T') U' \quad \Gamma \vdash T <: T'}{\Gamma, x: T' \vdash U' <: U} \text{INDUCTION ON } \vdash_{\#} \\
\hline
\frac{v = \lambda(x: T') t \quad \Gamma, x: T' \vdash t: U' \quad \Gamma \vdash T <: T'}{\Gamma, x: T' \vdash U' <: U} \text{INVERSION OF ALL-I!} \\
\hline
\frac{v = \lambda(x: T') t \quad \Gamma, x: T \vdash t: U' \quad \Gamma \vdash T <: T'}{\Gamma, x: T \vdash U' <: U} \text{NARROWING} \\
\hline
\frac{v = \lambda(x: T') t \quad \Gamma, x: T \vdash t: U \quad \Gamma \vdash T <: T'}{\Gamma \vdash T <: T'} \text{SUB}
\end{array}$$

□

Proof of Lemma 11 (\forall to $\Gamma(x)$).

$$\begin{array}{c}
\frac{\text{inert } \Gamma \quad \Gamma \vdash x: \forall(y: T) U}{\text{inert } \Gamma \quad \Gamma \vdash_{\#} x: \forall(y: T) U} \text{THEOREM 6 } (\vdash \text{ TO } \vdash_{\#}) \\
\frac{\text{inert } \Gamma \quad \Gamma \vdash_{\#} x: \forall(y: T) U}{\text{inert } \Gamma \quad \Gamma \vdash_i x: \forall(y: T) U} \text{THEOREM 10 } (\vdash_{\#} \text{ TO } \vdash_i) \\
\hline
\frac{\Gamma \vdash_i x: \forall(y: T') U' \quad \Gamma \vdash T <: T' \quad \Gamma, y: T' \vdash U' <: U}{\Gamma \vdash_i x: \forall(y: T') U' \quad \Gamma \vdash T <: T' \quad \Gamma, y: T \vdash U' <: U} \text{ON } \vdash_i \\
\hline
\frac{\Gamma \vdash_i x: \forall(y: T') U' \quad \Gamma \vdash T <: T' \quad \Gamma, y: T \vdash U' <: U}{\Gamma(x) = \forall(y: T') U' \quad \Gamma \vdash T <: T' \quad \Gamma, y: T \vdash U' <: U} \text{NARROWING INDUCTION ON } \vdash_i
\end{array}$$

□

Proof of Lemma 13 (μ to $\Gamma(x)$).

$$\begin{array}{c}
\frac{\text{inert } \Gamma \quad \Gamma \vdash x: \{a: T\}}{\text{inert } \Gamma \quad \Gamma \vdash_{\#} x: \{a: T\}} \text{THEOREM 6 } (\vdash \text{ TO } \vdash_{\#}) \\
\frac{\text{inert } \Gamma \quad \Gamma \vdash_{\#} x: \{a: T\}}{\text{inert } \Gamma \quad \Gamma \vdash_i x: \{a: T\}} \text{THEOREM 10 } (\vdash_{\#} \text{ TO } \vdash_i) \\
\hline
\frac{\Gamma \vdash_i x: \{a: T'\} \quad \Gamma \vdash T' <: T}{\Gamma \vdash_i x: \{a: T'\} \quad \Gamma \vdash T' <: T} \text{INDUCTION ON } \vdash_{\#} \\
\hline
\frac{\Gamma(x) = \mu(x: \dots \wedge \{a: T'\} \wedge \dots) \quad \Gamma \vdash T' <: T}{\Gamma(x) = \mu(x: \dots \wedge \{a: T'\} \wedge \dots) \quad \Gamma \vdash T' <: T} \text{INDUCTION ON } \vdash_i
\end{array}$$

□

Proof of Lemma 14 (μ to v).

$$\begin{array}{c}
\frac{\text{inert } \Gamma \quad \Gamma \vdash v: \mu(x: S) \quad S = \dots \wedge \{a: T\} \wedge \dots}{\text{inert } \Gamma \quad \Gamma \vdash_{\#} v: \mu(x: S) \quad S = \dots \wedge \{a: T\} \wedge \dots} \text{THEOREM 6 } (\vdash \text{ TO } \vdash_{\#}) \\
\frac{\text{inert } \Gamma \quad \Gamma \vdash_{\#} v: \mu(x: S) \quad S = \dots \wedge \{a: T\} \wedge \dots}{\text{inert } \Gamma \quad \Gamma \vdash_i v: \mu(x: S) \quad S = \dots \wedge \{a: T\} \wedge \dots} \text{THEOREM 10 } (\vdash_{\#} \text{ TO } \vdash_i) \\
\hline
\frac{\text{inert } \Gamma \quad \Gamma \vdash_i v: \mu(x: S) \quad S = \dots \wedge \{a: T\} \wedge \dots}{\text{inert } \Gamma \quad \Gamma \vdash_i v: \mu(x: S) \quad S = \dots \wedge \{a: T\} \wedge \dots} \text{INDUCTION ON } \vdash_{\#} \\
\hline
\frac{\text{inert } \Gamma \quad \Gamma \vdash_i v: \mu(x: S) \quad S = \dots \wedge \{a: T\} \wedge \dots}{\text{inert } \Gamma \quad v = v(x: S)(\dots \wedge \{a: t\} \wedge \dots) \quad \Gamma \vdash t: T} \text{INVERSION OF } \{\}-\text{I!}
\end{array}$$

□

Since the return type of a dependent function type depends on the parameter type, this proof and the proof of Lemma 11 (\forall to $\Gamma(x)$) rely on a standard narrowing property, which states that making a typing

context more precise by substituting one of the types by its subtype preserves the typing and subtyping relations.

LEMMA 15 (Narrowing). *Suppose $\Gamma(x) = T$ and $\Gamma[x: T'] \vdash T' <: T$. Then $\Gamma \vdash t: U$ implies $\Gamma[x: T'] \vdash t: U$, and $\Gamma \vdash S <: U$ implies $\Gamma[x: T'] \vdash S <: U$.*

Narrowing is proved for DOT by Amin, Grütter, et al., (2016). The proof is standard, with no issues specific to DOT, by induction on the typing and subtyping rules.

Given the above lemmas we can infer the precise type of a variable given its general type, and the shape of a value given the value's type. What we need to do now is to establish a connection between variables and values so that we can prove the canonical-forms lemmas. The first step is to define a correspondence relation between value and type environments: since evaluation happens on pairs $\gamma \mid t$ of stores and terms we need to make sure that the typing environment in which we type a term corresponds to the term's runtime configuration γ .

DEFINITION 16 (Well-formedness). *A store $\gamma = (\overline{x_i, v_i})$ is well-formed with respect to a typing context $\Gamma = (\overline{x_i, T_i})$, denoted $\gamma: \Gamma$, if $\Gamma \vdash v_i: T_i$ for all $i = 1, \dots, n$.*

The following lemma states that any variable that has a type in the environment maps to a value of the same type in the store.

LEMMA 17 (Corresponding Types). *Suppose that a store γ is well-formed with respect to an environment Γ , and that Γ assigns type T to a variable x . Then there exists a value v that is assigned to x by the store γ , and v has the same type T as x .*

The final canonical-forms lemmas presented below immediately follow from Lemma 17 and Lemma 11 to Lemma 14.

LEMMA 18 (Canonical Forms for Objects). *Suppose that a store γ is well-formed with respect to an inert environment Γ , and that $\Gamma \vdash x: \{a: T\}$. Then the store γ assigns the object $v(x: U) \cdots \wedge \{a = t\} \wedge \dots$ to x , where t has type T .*

LEMMA 19 (Canonical Forms for Functions). *Suppose that a store γ is well-formed with respect to an inert environment Γ , and that $\Gamma \vdash x: \forall(x: T) U$. Then the store γ assigns the object $\lambda(x: T') t$ to x , where $\Gamma, x: T \vdash t: U$ and $\Gamma \vdash T <: T'$.*

$$\frac{\Gamma, x: T \vdash t: U \quad \Gamma, x: T' \vdash T' <: T}{\Gamma, x: T' \vdash t: U} \text{ (NARROWING)}$$

$$\frac{\emptyset: \emptyset \quad \gamma: \Gamma \quad \Gamma \vdash v: T}{\gamma, x \mapsto v: \Gamma, x: T}$$

$$\frac{\gamma: \Gamma \quad \Gamma(x) = T}{\gamma(x) = v \wedge \Gamma \vdash v: T} \text{ (CORRESP. TYPES)}$$

$$\frac{\gamma: \Gamma \quad \text{inert } \Gamma \quad \Gamma \vdash x: \{a: T\}}{\gamma(x) = v(x) \cdots \{a = t\} \cdots} \text{ (Can. Forms for } v)$$

$$\frac{\gamma: \Gamma \quad \text{inert } \Gamma \quad \Gamma \vdash x: \forall(x: T) U}{\gamma(x) = \lambda(x: T') t} \text{ (Can. Forms for } \lambda)$$

5.6 PROGRESS, PRESERVATION, AND SOUNDNESS

To express the final soundness theorems we need to introduce one additional lemma, the notion of a well-formed store, and a definition of normal forms.

The following lemma states that all well-typed values have an inert precise type.

$$\begin{array}{c}
\Gamma \vdash v : T \\
\hline
\Gamma \vdash_{\text{!}} v : T' \\
\text{inert } T' \\
\hline
\Gamma \vdash T' <: T \\
\text{(VALUE TYPING)}
\end{array}$$

LEMMA 20 (Value Typing). *If $\Gamma \vdash v : T$, then there exists an inert type T' such that $\Gamma \vdash_{\text{!}} v : T'$ and $\Gamma \vdash T' <: T$.*

Proof. The proof is by induction on the derivation of $\Gamma \vdash v : T$, and is short because only three typing rules apply to values: ALL-I, {}-I, and SUB. In the first two cases, the precise type of v coincides with the general type. The subsumption case is handled by using the induction hypothesis and transitivity of subtyping. Furthermore, the precise-typing judgment for values and definition-typing rules ensure that any precise type of a value is inert. \square

$$x \rightarrow \quad v \rightarrow$$

DEFINITION 21 (Normal Form). *A term t is in normal form, denoted $t \rightarrow$, if t is either a variable or a value.*

With these results, we can prove the standard progress theorem that every typable term is a normal form or reduces to some other term.

$$\begin{array}{c}
\gamma : \Gamma \quad \text{inert } \Gamma \\
\hline
\Gamma \vdash t : T \\
\hline
\gamma \mid t \mapsto \gamma' \mid t' \\
\vee t \rightarrow \\
\text{(PROGRESS)}
\end{array}$$

THEOREM 22 (Progress). *Let γ be a store and Γ an inert typing environment such that $\gamma : \Gamma$. If $\Gamma \vdash t : T$ then t is in normal form or there exists a term t' and a store γ' such that $\gamma \mid t \mapsto \gamma' \mid t'$.*

Proof. We proceed by induction on the derivation of $\Gamma \vdash t : T$, in each case finding a reduction rule that applies. The interesting cases are ALL-E and {}-E.

In the premises of ALL-E, variable x has type $\forall(z : S) T$. Lemma 11 (\forall to $\Gamma(x)$) tells us that Γ binds x to a compatible function type. $\gamma : \Gamma$ ensures that γ binds x to some value v , and that v has a compatible function type. Finally, Lemma 12 (\forall to λ) tells us the v is a lambda, so the APPLY reduction rule can be applied.

The {}-E case is similar, but using Lemma 13 (μ to $\Gamma(x)$) and Lemma 14 (μ to v) instead of Lemmas 11 and 12, respectively, and the PROJECT reduction rule instead of APPLY. \square

We can now prove the standard preservation theorem.

$$\begin{array}{c}
\gamma : \Gamma \quad \text{inert } \Gamma \\
\hline
\Gamma \vdash t : T \\
\hline
\gamma \mid t \mapsto \gamma' \mid t' \\
\hline
\Gamma' \vdash t' : T \\
\gamma' : \Gamma' \quad \text{inert } \Gamma' \\
\text{(PRESERVATION)}
\end{array}$$

THEOREM 23 (Preservation). *Let γ be a store and Γ an inert typing environment such that $\gamma : \Gamma$. If $\Gamma \vdash t : T$ and $\gamma \mid t \mapsto \gamma' \mid t'$ then there exists an inert context Γ' such that $\gamma' : \Gamma'$ and $\Gamma' \vdash u : T$.*

Proof. The proof proceeds by induction on the typing derivation. The interesting cases are ALL-E, FLD-E, and LET. In the first two cases we apply Lemmas 11 to 14 to obtain values with the necessary types like in the proof of Theorem 22 (Progress).

In the LET case, $t = \text{let } x = u \text{ in } u'$, where $\Gamma \vdash u : U$ and $\Gamma, x : U \vdash u' : T$ for some type U . We proceed by a case analysis on the shape of u . Here, the interesting case is when u is a value v . Since $\Gamma \vdash v : U$, by Lemma 20 (Value Typing), there exists an inert type U' such that $\Gamma \vdash_{\text{!}} v : U'$ and $\Gamma \vdash U' <: U$. We choose as our inert context $\Gamma' =$

$\Gamma, x: U'$. The term $\gamma \mid \text{let } x = v \text{ in } u'$ reduces to $\gamma, x \mapsto v \mid u'$, and we have $\gamma, x \mapsto v: \Gamma, x: U'$ as needed. Finally, we have to show that $\Gamma, x: U' \vdash u': T$, which follows $\Gamma, x: U \vdash u': T$ and $\Gamma \vdash U' <: U$ by Lemma 15 (Narrowing). \square

Using progress and preservation it is easy to prove the final DOT type soundness theorem.

DEFINITION 24 (Divergence). *A term t diverges, denoted $t \Uparrow$, if there exists an infinite reduction sequence*

$$\emptyset \mid t \mapsto \gamma_1 \mid t_1 \mapsto \dots \mapsto \gamma_n \mid t_n \mapsto \dots$$

THEOREM 25 (DOT Type Soundness). *If $\vdash t: T$ then either t diverges ($t \Uparrow$) or t reduces to a normal form t' , i.e. $\emptyset \mid t \mapsto^* \gamma \mid t'$, $t' \rightarrow$, and $\Gamma \vdash t': T$ for some Γ such that $\gamma: \Gamma$.*

$$\frac{\vdash t: T}{(\emptyset \mid t \mapsto^* \gamma \mid t' \wedge t' \rightarrow) \vee t \Uparrow} \text{ (DOT SOUNDNESS)}$$

5.7 PROOF STRUCTURE AND EXTENSIONS

This section summarizes the structure of the proof and discusses how the proof is affected by changes and extensions of the DOT calculus.

5.7.1 Proof Structure

The dependencies between the main lemmas in the proof are summarized in the diagram in Figure 5.4. The gray nodes and solid lines denote the lemmas in the simple proof for DOT. The white boxes and dotted lines correspond to changes needed to prove soundness of an example extension of the calculus that will be described in Chapter 6.

The final progress and preservation theorems depend on the four applications of the proof recipe to prove canonical forms for values and variables of object and function type (written in the figure as \forall to λ , μ to ν , \forall to $\Gamma(x)$, and μ to $\Gamma(x)$). Each application of the proof recipe uses Theorem 6 (\vdash to $\vdash_\#$) and Theorem 10 ($\vdash_\#$ to \vdash_i) to convert general typing to tight typing and then to invertible typing. Theorem 6 depends on Lemma 7 (SEL- $<:-\#$ Replacement) and Lemma 8 (SEL- $<:-\#$ Premise). Theorem 10 depends on a subtyping closure helper lemma. After using the theorems to obtain invertible typing, we invert the invertible typing in each of the four cases (see “invertible to precise” level in the dependency graph) to obtain either the type $\Gamma(x)$ assigned to a variable x by the typing context Γ or the form of the value v with the given type. The four light large boxes in the figure indicate the canonical-forms lemmas, and the three phases of the proof recipe (conversion of general to tight typing, tight to invertible typing, and inversion of invertible typing).

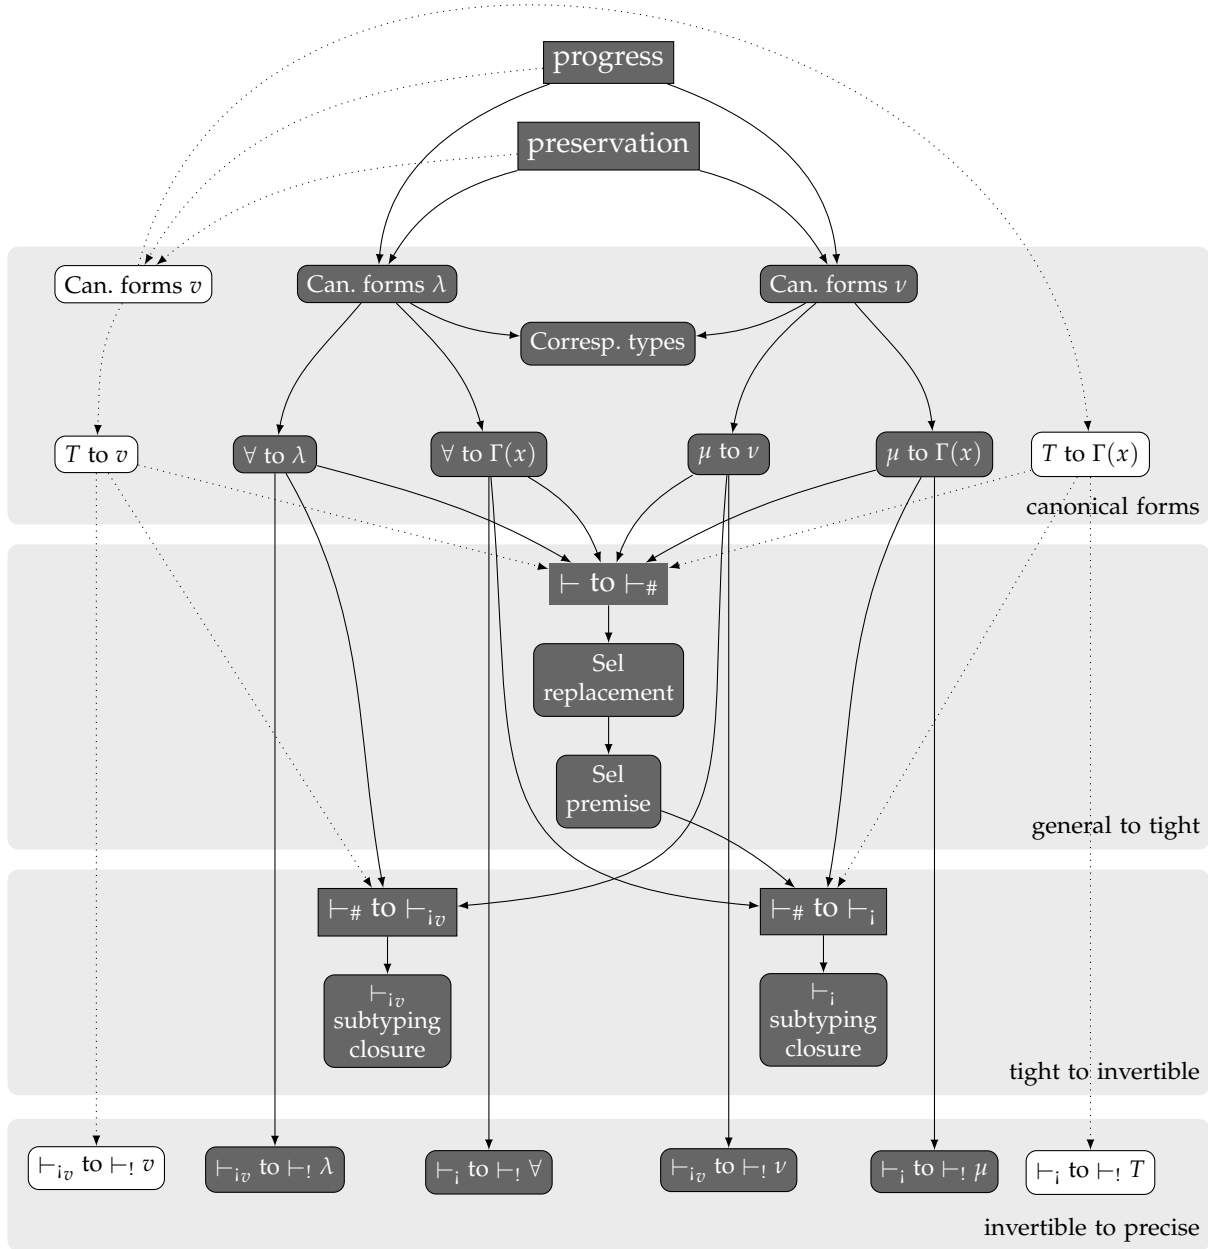


Figure 5.4: Dependencies between main lemmas in the proof. Gray nodes denote existing lemmas. White nodes denote lemmas that would need to be added if DOT were extended with a new type T and a new value v .

MODIFICATIONS OF THE CALCULUS

The most common expected extensions of a calculus are the addition of new forms of values and terms, of new forms of types and typing rules, and changes to the evaluation rules. Most extensions will change multiple aspects (e.g., add a new form of value and an associated type), but we discuss each change individually. In Part II, we will present a specific example of an extension that makes all of these kinds of changes.

The only part of our proof that deals with values are the two pairs of canonical forms lemmas in Section 5.5 and the final progress and preservation theorems. A new form of value will require an additional pair of canonical-forms lemmas. The lemma can follow the general recipe: it will apply Theorem 6 ($\vdash \text{to } \vdash_{\#}$) and Theorem 10 ($\vdash_{\#} \text{to } \vdash_{\downarrow}$). It does not need to reason with the general DOT typing rules, but only to invert the invertible typing obtained from Theorem 10. This last inversion step should be easy, because invertible typing is designed to be easily invertible. The addition of a new form of value is illustrated in the dependence graph by the two white nodes on the left side of the graph.

The only part of the proof that deals with terms in general are the final progress and preservation theorems. The only non-trivial change required when adding a new term is that if new reduction rules are added for the new term to the operational semantics, cases for the new reduction rules need to be added to the progress and preservation theorems. This is illustrated in the dependence graph by the dotted outlines of the nodes representing those two theorems.

Adding a new form of type is a more significant change. Given general typing rules for the new type, we must incorporate the changes into the tight, invertible, and precise typing rules. Tight typing differs from general typing only in its handling of abstract type members and type projections, so changes unrelated to those features can be incorporated directly into tight typing. A change involving abstract type members or type projections requires corresponding modifications to tight typing. Property 5 ([Tight Typing Guarantee](#)) gives a modular specification to guide the design of such modifications. Specifically, we know that as long as the modified tight typing rules satisfy the property and we can prove Theorem 6 ($\vdash \text{to } \vdash_{\#}$), then the proof recipe and the rest of the whole soundness proof will continue to hold without requiring non-trivial changes. To incorporate the modifications into invertible and precise typing, it suffices to follow the general recipe outlined in Section 5.4. Specifically, we must classify the new tight

typing rules as either introducing or eliminating a syntactic construct, and then add them to either invertible or precise typing, respectively. Adding new typing rules requires adding the corresponding cases to the proofs of Theorem 6 ($\vdash \text{to} \vdash_{\#}$) and Theorem 10 ($\vdash_{\#} \text{to} \vdash_i$). In the proof of Theorem 6 ($\vdash \text{to} \vdash_{\#}$), all cases except $<:-\text{SEL}$ and $\text{SEL}<:-$ are so simple that Coq discharges them automatically, so we do not add a dotted outline to the node to indicate new cases in the proof.

A change to the evaluation rules of the calculus does not affect any of the reasoning in Chapter 5 since this chapter is independent of any particular evaluation semantics. Only the final progress and preservation theorems are affected.

The next part of this thesis presents a specific extension of DOT by showing how to add mutable references to the calculus.

THE STRUGGLE FOR “GOOD” BOUNDS

A recurring theme in previous work on DOT has been the struggle to enforce “good” bounds. A type member declaration $\{A: S..U\}$ is considered to have “good” bounds if $S <: U$. If all type members could be forced to maintain “good” bounds, it would prevent an object of type $\mu(x: \{A: S..U\})$ from introducing a new, possibly non-sensical subtyping relationship $S <: U$ from $S <: x.A <: U$ and transitivity. Many of the challenges along the way to defining a sound DOT calculus arose from the negative interaction between “good” bounds and other properties, such as narrowing and transitivity. For example, although both $\{A: \perp.. \perp\}$ and $\{A: \top.. \top\}$ have “good” bounds, the narrowed type $\{A: \perp.. \perp\} \wedge \{A: \top.. \top\}$ causes trouble: in the function

$$\lambda(x: \{A: \perp.. \perp\} \wedge \{A: \top.. \top\}) t$$

the body t is type-checked in a typing context in which $\top <: x.A <: \perp$.

Not only do “good” bounds interact poorly with other desirable properties, but even defining precisely what “good” bounds are is surprisingly elusive. Informally, bounds are “good” if $S <: U$. But in what typing context should this subtyping relationship hold? In deciding whether the type $\mu(x: \{A: S..U\})$ should be allowable, it seems appropriate to respect the recursion implied by μ and use a context that includes x ; that is, to require that $\Gamma, x: \{A: S..U\} \vdash S <: U$. But this statement is always true regardless of the types S and U because it is self-justifying:

$$\Gamma, x: \{A: S..U\} \vdash S <: x.A <: U.$$

If we decide instead to exclude the self-reference x from the context used to decide whether $S <: U$, we exclude many desirable types from the definition of “good” bounds. For example, we consider “bad” the type

$$\mu(x: \{A: \perp.. \top\} \wedge \{B: x.A..x.C\} \wedge \{C: \perp.. \top\})$$

that innocently defines three type members with $A <: B <: C$, because $x.A$ cannot be a subtype of $x.C$ without x in the context. We also consider “bad” the following type that defines two type members $A <: B$ constrained to be function types:

$$\mu(x: \{A: \perp.. \forall(y: \perp) \top\} \wedge \{B: x.A.. \forall(y: \perp) \top\}).$$

Again, $x.A$ cannot be a subtype of $\forall(y: \perp) \top$ without x in the context. Finally, such a definition of “good” bounds restricts the applicability

of type aliases: the following type defines A and B as aliases for \top and \perp , respectively, but cannot use these aliases in the bounds of C because $x.B \not\leq x.A$ in a context without x :

$$\mu(x: \{A: \top.. \top\} \wedge \{B: \perp.. \perp\} \wedge \{C: x.B.. x.A\})$$

Although it would be possible to come up with some definition of “good” bounds that handles these specific examples, the definition of what was intended to be an obvious and intuitive concept would become very complicated, and other more sophisticated counterexamples would probably continue to exist. Thus, it appears that trying to enforce “good” bounds, and even trying to define what “good” bounds are, is a dead end.

By contrast, inert types obey a purely syntactic property that is easily defined and checked, without requiring a subtyping judgment in some typing context that would have to be specified. The property provided by an inert typing context can be stated precisely and formally (Property 4 ([Inert Context Guarantee](#))).

RELATED WORK

The work presented in this part of the thesis is based on a previously published publication of the simple DOT soundness proof (Rapoport, Kabir, et al., 2017).

8.1 DOT SOUNDNESS PROOFS

The work most closely related to ours is Amin, Grütter, et al., (2016), which defines and proves sound the variant of the DOT calculus for which we have developed our alternative soundness proof. That work also defines tight typing, though it does not use it as pervasively as our proof does.

A central notion of that proof is store correspondence, a relationship between typing contexts and stores of runtime values. A typing context Γ corresponds to a store s if for every variable x , $\Gamma \vdash! s(x) : \Gamma(x)$. Typing and subtyping in a context Γ that corresponds to some store s have similar predictable behaviour as they do in an inert context. Part of the proof consists of lemmas that relate internal details of values in stores with internal details of types in corresponding contexts. By contrast, the property of inert contexts is independent of values, so our proof does not depend on such lemmas.

Another central notion is “possible types”: if a typing context Γ corresponds to some store s , and s assigns to variable x the value v , then the possible types of the triple (Γ, x, v) include all types T such that $\Gamma \vdash x : T$. Possible types serve a similar purpose as our invertible typing rules, to facilitate induction proofs. Unlike invertible typing, possible types depend on the runtime value v of x . The possible types lemma relates general typing in a context with a corresponding store to possible types. It serves a similar purpose as our Theorem 10 ($\vdash_{\#} \text{ to } \vdash_i$) (which relates tight to invertible typing), but its proof is more complicated, because it depends on sublemmas that relate types to values in the context corresponding to the store, and on general typing.

Amin, Grütter, et al., (2016) also prove a similar result as Theorem 6 ($\vdash \text{ to } \vdash_{\#}$): the general to tight lemma states that in a context Γ for which there exists some corresponding runtime store s , general typing implies tight typing. We prove Theorem 10 ($\vdash_{\#} \text{ to } \vdash_i$) first, which makes proving Theorem 6 ($\vdash \text{ to } \vdash_{\#}$) easy. The proof of Amin, Grütter, et al., (2016) does the analogous steps in the opposite order: it proves the general to tight lemma first, and the possible types lemma afterwards, using the general to tight lemma in its proof. The proof of the general to tight lemma is thus complicated because it cannot make use

of possible types. Another complication is that the proof of the general to tight lemma, like the proof of the possible types lemma, depends on sublemmas that relate types to values in the context corresponding to the store.

Rompf and Amin, (2016b) define a variant of the DOT calculus with additional features, most significantly subtyping between recursive types. This adds significant complexity to the proof: Lemmas 6 to 11 are needed only because of this feature. However, subtyping between Scala’s types can be already modelled by subtyping between type members in DOT. Scala has nominal subtyping between classes and traits that are explicitly declared to be subtypes using an `extends` clause. A class or trait declaration in Scala corresponds in DOT to a type member definition in some package x that gives a label A to a recursive type. The recursive type is used to define the members of the class, and the recursion is necessary so that members of the class can refer to the object of the class `this`. A subclass B of A can be declared as the type member definition $B = x.A \wedge \mu (z: x.A \wedge T)$, where the type T describes the additional members that B adds to A . Then $x.B$ is a subtype of $x.A$, and given a variable of type $x.B$, it is possible to access both members that were declared in A and members that were added in B . This DOT encoding models the Scala subtyping between classes A and B without requiring subtyping between recursive types, and it can be expressed in the DOT of Amin, Grütter, et al., (2016).

Unlike Amin, Grütter, et al., (2016) and our proof, the proof of Rompf and Amin, (2016b) does not use tight typing, the typing relation that neutralizes the two type rules that enable a DOT program to introduce non-sensical subtyping relationships in a custom type system. Instead, the proof uses “precise subtyping”, a restriction of general subtyping to relationships whose derivation does not end in the transitivity rule.

8.2 HISTORY OF SCALA CALCULI

Odersky, Cremet, et al., (2003) introduce νObj , a calculus to formalize Scala’s path-dependent types. νObj includes abstract type members, classes, compound (non-commutative) mixin composition, and singleton types, among other features. However, the calculus lacks several essential Scala features, such as the ability to define custom lower bounds for type members, and has no top and bottom types. Additionally, νObj , unlike Scala, has classes as first-class values. νObj comes with a type soundness proof. The paper also shows that type checking for νObj is undecidable. Cremet et al., (2006) propose Featherweight Scala, which is similar to νObj , but without classes as first-class values. The paper shows that type inference in Featherweight Scala is decidable, but does not prove type safety. Scalina, introduced by

Moors, Piessens, and Odersky, (2008), presents a formalization for higher-kinded types in Scala, but also without a soundness proof.

Amin, Moors, and Odersky, (2012) present the first DOT. DOT has fewer syntax-level features than νObj : there are no classes, mixins, or inheritance. However, some of the previously missing crucial Scala features are now present. The calculus allows refinement of abstract type members through commutative intersections, combining nominal with structural typing. Type members can have custom lower and upper bounds, and the type system contains a bottom and top type. The paper comes without a type safety proof, but it explains the challenges and provides counterexamples to preservation. The paper shows how the environment narrowing property makes proving soundness complicated: replacing a type in the context with a more precise version can impose a new subtyping relationship, which could disagree with the existing ones.

Amin, Rompf, and Odersky, (2014) have the first mechanized soundness proof for μDOT , a simplified calculus that excludes refinements, intersections, and the bottom and top types, and uses big-step semantics. The paper proposes the idea to circumvent bad bounds by reasoning about types that correspond to runtime values.

Amin, Grütter, et al., (2016) and Rompf and Amin, (2016b) build on this store correspondence idea, to establish the first mechanized soundness proofs for DOT calculi with support for type intersection and refinement, and top and bottom types. The two calculi and soundness proofs were discussed in the previous section.

8.3 OTHER RELATED CALCULI

Path-dependent types were first introduced in the context of *family polymorphism* by Ernst, (2001). In family polymorphism, groups of types can form *families* that correspond to a specific object. Two types from the same class are considered incompatible if the types are associated with different runtime objects.

Family polymorphism is the foundation of *virtual classes*, which were introduced in the Beta programming language (Madsen and Møller-Pedersen, 1989) and further developed in *gbeta* (Ernst, 1999). Virtual classes are nested classes that can be extended or redefined (overridden), and are dynamically resolved through late binding. Family polymorphism allows for a fine-grained distinction between classes that have the same static path, yet belong to different runtime objects and can thus have different implementations.

Virtual classes were first formalized and proved type safe in the *vc* calculus (Ernst, Ostermann, and Cook, 2006). *vc* is a class-based, nominally-typed calculus with a big-step semantics. To create path-based types, the keyword *out* is used to refer to an enclosing object. With its support for classes, inheritance, and mutation of variables,

vc is more complex than DOT, whose purpose is to serve as a simple core calculus for Scala. Additionally, Scala has no support for virtual classes: the language does not allow class overriding, and its classes are resolved statically at compile time.

Tribe by Clarke et al., (2007) is a simpler, more general calculus inspired by vc. One of the main distinctions to vc is that variables, and not just enclosing objects (out), can be used as paths for path-dependent types. This makes the calculus more general, as it can express subtyping relationships between classes with arbitrary absolute paths. *Tribe* comes with a type-safety proof, which is based on a small-step semantics. Expanding paths to allow variables brings *Tribe* closer to DOT. However, the complexity of the type system, resulting from modelling classes and inheritance, and the modelling of virtual classes, which are not present in Scala, leaves DOT more suitable as a core calculus for Scala.

Amin and Rompf, (2017) offer a survey of mechanized soundness proofs for big-step, DOT-like calculi using definitional interpreters. The paper explores a family of calculi ranging from System F to System $D_{<:\triangleright}$ and general proof techniques that can be applied to this entire family. The paper discusses similarities and differences between System $D_{<:\triangleright}$ and DOT.

8.4 TYPE CHECKING DECIDABILITY

Odersky, Cremet, et al., (2003) proved that there exists no algorithm that can decide if a typing judgment in the νObj calculus is well-typed. To carry out the proof, νObj includes first-class classes, which are not present in Scala. The proof works via a reduction from System $F_{<:}$, which was shown to be undecidable (Pierce, 1992a). To come up with a decidable type system for Scala, Cremet et al., (2006) developed the *Featherweight Scala* calculus and proved its type system decidable; however, the formalization was not proved sound and excluded unique lower and upper bounds for the subtyping lattice and lower bounds for type members. By formalizing a large subset of Scala, νObj and *Featherweight Scala* significantly differ from DOT-like calculi whose purpose is to serve as small extensible core calculi.

The DOT calculus is widely conjectured to have undecidable type-checking because it includes the features of $F_{<:}$, for which typechecking is undecidable (Pierce, 1992b). Rompf and Amin, (2016b) give a mapping from $F_{<:}$ to $D_{<:}$, a simpler calculus than DOT, and prove that if the $F_{<:}$ term is typeable then so is the $D_{<:}$ term. However, the mapping does not preserve typeability in the only-if direction which is required to prove undecidability of typing: Hu and Lhoták, (2019) present an example of a typeable $D_{<:}$ term whose translation to $F_{<:}$ does not have a type. Hu and Lhoták then present an alternative

mapping between $F_{<}$ and $D_{<}$ and prove undecidability of $D_{<}$ in Agda.

To investigate what it would take to make typechecking a DOT-like calculus algorithmic, Nieto, (2017) presents a decidable type-checking algorithm for a subset of $D_{<}$ with restricted subtyping rules. The paper discusses how bad bounds complicate algorithmic typing of DOT-like calculi, and how the Scala compiler avoids this issue by dropping subtyping transitivity altogether. Hu and Lhoták, (2019) define *kernel* $D_{<}$, a calculus equivalent to Nieto’s subset of $D_{<}$ and formalize its decidability proof in Coq. The paper also presents a mechanized proof for a version of *kernel* $D_{<}$ that allows comparing the parameter types of function types for equality, adding significant expressivity to the calculus.

8.5 SYNTACTIC VS. SEMANTIC PROOFS

The DOT papers that establish type safety based on a small-step operational semantics, including this work, use the *syntactic* approach to proving type safety of Wright and Felleisen, (1994).³ In a syntactic proof, we develop a set of inductive type rules and prove *progress* and *preservation* theorems for our operational semantics. Progress says that any well-typed, closed term is either a value or can take a step. Preservation says that reduction preserves the types of terms. Together, progress and preservation imply type safety: well-typed programs do not get stuck.

The syntax-based approach treats types as syntactic constructs without directly expressing their meaning. By contrast, *semantic* approaches to type soundness are based on establishing a set-theoretic semantics for types, where types are explicitly defined in terms of the values they represent (Appel and Felty, 2000). Instead of using type rules as axioms, a semantic proof allows us to prove type rules as derived lemmas. This leads to a system in which we only need to trust the operational semantics and type definitions, instead of relying on the meaningfulness of the type system, whose definition is often larger and more difficult to reason about. In addition, semantic proofs tend to be more reusable and are easier to scale (Bell et al., 2008). Unfortunately, encoding types as sets of values proves often difficult and involves more sophisticated mathematics.

Creating a semantic type-soundness proof for DOT remains an open problem. Wang and Rompf, (2017) propose a proof of strong normalization using for $D_{<}$, a restricted version of DOT. The paper presents the first semantic encoding of a subset of DOT’s types, and the normalization proof is a semantic one that is based on logical relations. At the moment, Paolo Giarrusso is actively working on a semantic type soundness proof for the full DOT calculus in Iris that is based on logical relations (Giarrusso et al., 2019).

³ Big-step DOT proofs also use a syntactic approach but based on definitional interpreters (Reynolds, 1998).

SUMMARY

DOT (Amin, Grütter, et al., 2016) is the result of a long effort to develop a core calculus for Scala. Now that there is a sound version of the calculus, we would like to extend it with other Scala features, such as classes, mixin composition, side effects, implicit parameters, etc. DOT can be also used as a platform for developing new language features and for fixing Scala's soundness issues (Amin and Tate, 2016). But these applications are hindered by the complexity of the existing soundness proofs, which interleave reasoning about variables, types, and runtime values, and their complex interactions.

We have presented a simplified soundness proof for the DOT calculus, formalized in Coq. The proof separates the reasoning about types, typing contexts, and values from each other. The proof depends on the insight of inert typing contexts, a syntactic characterization of contexts that rule out any non-sensical subtyping that could be introduced by abstract type members. The central lemmas of the proof follow a general proof recipe for deducing properties of terms from their types in full DOT while reasoning only in a restricted, intuitive environment free from the paradoxes caused by abstract type members. The same recipe can be followed to prove similar lemmas when the calculus is modified or extended. The result is a simple, modular proof that is well suited for developing extensions.

Part II

CASE STUDY: MUTABLE DOT

INTRODUCTION

DOT models the key components of the Scala type system such as type members, path-dependent types, and subtyping. However, the calculus is still lacking some fundamental Scala features, one of which is mutation.

Without mutation, it is difficult to model mutable variables and fields, or to reason about side effects in general. Interestingly, mutation is even necessary to model a sound class initialization order for immutable fields, which are mutated once when they are initialized (Kabir and Lhoták, 2018). Formalizing Scala initialization order would require reasoning about overwriting of class members that were initialized with null, which is not directly possible in DOT.

This chapter presents the Mutable DOT calculus, an extension to DOT with typed mutable references. To that end, we augment the calculus with a mutable heap and the possibility to create, update, and dereference mutable memory cells, or locations. To model mutable variables (vars), one can create a heap location and store immutable variables (vals) in it (immutable variables are already modelled in DOT). For example, a Scala object

```
object O {
  val x = 1
  var y = 2
}
```

can be represented in mutable-DOT pseudocode as follows:

```
new {this: {x: Int} ∧ {y: Ref Int}}
  {x = 1} ∧ {y = ref 2 Int}
```

An unusual characteristic of our heap implementation is that it maps locations to variables instead of values. This design choice is induced by DOT's type system, which disallows subtyping between recursive types. We show how, as a result, storing values on the heap would significantly limit the expressiveness of our calculus, and explain the correctness of storing variables on the heap.

CONTRIBUTIONS

This part of the thesis presents the following contributions:

- an operational semantics and type system for *Mutable DOT*, an extension of the DOT calculus with mutable references;
- a *mechanized type safety proof* in Coq, in the form of an extension of the original DOT proof;

see Chapter 11

Coq proof:
<https://git.io/fjlq6>,
 see Chapter 12

see Chapter [13](#) – a *discussion* of the Mutable DOT design choices and examples.

THE MUTABLE DOT CALCULUS

In this chapter we present the Mutable DOT calculus as an extension of DOT.

11.1 MUTABLE DOT ABSTRACT SYNTAX

To support mutation, we augment the DOT syntax with *references* that point to *mutable memory cells*, or *locations*, as shown in Figure 11.1.

Locations are a new kind of value that is added to the syntax, and are denoted as l . The syntax comes with three new terms to support the following reference operations:

- $\text{ref } x \ T$ *creates* a new reference of type T in the heap and initializes it with the variable x . Section 13.3 explains why reference expressions need to contain a declared type T .
- $!x$ *reads* the contents of a reference x .
- $x := y$ *updates* the contents of a reference x with the variable y .

The operations that create, read, and update references operate on variables, not arbitrary terms, in order to preserve ANF.

Newly-created references become *locations*, or memory addresses, denoted as l . Locations are stored in the *heap*, denoted as σ .

The heap is a *map* from locations to *variables*. This differs from the common definition of a heap which maps locations to values. We discuss the motivation for this design choice in Section 13.1. In order to preserve the commonly expected intuitive behaviour of a heap, we must be sure that while a variable is in the heap, it does not go out of scope or change its value. We show this in Section 13.2.

Updating a heap σ that contains a mapping $l \mapsto x$ with a new mapping $l \mapsto y$ overwrites x with y :

$$(\sigma[l \mapsto x])(l') = \begin{cases} x & \text{if } l = l' \\ \sigma(l') & \text{otherwise.} \end{cases}$$

Locations are typed with the reference type $\text{Ref } T$. The underlying type T indicates that the location stores variables of type T .

```

 $t, u :=$ 
   $x$ 
   $v$ 
   $x.a$ 
   $xy$ 
  let  $x = t$  in  $u$ 
  ref  $x \ T$ 
   $!x$ 
   $x := y$ 
 $v :=$ 
   $v(x : T)d$ 
   $\lambda(x : T) t$ 
   $l$ 
 $d :=$ 
   $\{a = t\}$ 
   $\{A = T\}$ 
   $d \wedge d'$ 
 $S, T, U :=$ 
   $\top$ 
   $\perp$ 
   $\{a : T\}$ 
   $\{A : S..T\}$ 
   $x.A$ 
   $S \wedge T$ 
   $\mu(x : T)$ 
   $\forall(x : S) T$ 
  Ref  $T$ 

```

Figure 11.1: Abstract syntax of Mutable DOT (cf. DOT syntax in Figure 2.1)

$$\begin{array}{c}
\frac{\gamma(x) = \nu(x: T) \dots \{a = t\} \dots}{\sigma \mid \gamma \mid x.a \mapsto \sigma \mid \gamma \mid t} \quad (\text{PROJECT}) \\
\\
\frac{\gamma(x) = \lambda(z: T) t}{\sigma \mid \gamma \mid x y \mapsto \sigma \mid \gamma \mid t[y/z]} \quad (\text{APPLY}) \\
\\
\sigma \mid \gamma \mid \text{let } x = y \text{ in } t \mapsto \sigma \mid \gamma \mid t[y/x] \quad (\text{LET-VAR}) \\
\\
\sigma \mid \gamma \mid \text{let } x = v \text{ in } t \mapsto \sigma \mid \gamma, x \mapsto v \mid t \quad (\text{LET-VALUE}) \\
\\
\frac{\sigma \mid \gamma \mid t \mapsto \sigma' \mid \gamma' \mid t'}{\sigma \mid \gamma \mid \text{let } x = t \text{ in } u \mapsto \sigma' \mid \gamma' \mid \text{let } x = t' \text{ in } u} \quad (\text{CTX}) \\
\\
\frac{l \notin \text{dom}(\sigma)}{\sigma \mid \gamma \mid \text{ref } x T \mapsto \sigma[l \mapsto x] \mid \gamma \mid l} \quad (\text{REF}) \\
\\
\frac{\gamma(x) = l}{\sigma \mid \gamma \mid x := y \mapsto \sigma[l \mapsto x] \mid \gamma \mid y} \quad (\text{HEAP}) \\
\\
\frac{\gamma(x) = l \quad \sigma(l) = y}{\sigma \mid \gamma \mid !x \mapsto \sigma \mid \gamma \mid y} \quad (\text{DEREF})
\end{array}$$

Figure 11.2: Mutable DOT operational semantics

To write concise Mutable DOT programs, we extend the abbreviations from Section 2.4 with the following rules:

$$\begin{aligned}
\text{ref } t T &\equiv \text{let } x = t \text{ in ref } x T \\
t := u &\equiv \text{let } x = t \text{ in let } y = u \text{ in } x := y \\
!t &\equiv \text{let } x = t \text{ in } !x \\
t; u &\equiv \text{let } x = t \text{ in } u
\end{aligned}$$

11.2 MUTABLE DOT OPERATIONAL SEMANTICS

Since the meaning of a Mutable DOT term depends on the heap contents, we represent a program state as a triple $\sigma \mid \gamma \mid t$, denoting a term t that can point to memory contents in the heap σ and whose variables are stored in the store γ .

The new reduction semantics is shown in Figure 11.2:

- A newly created reference $\text{ref } x T$ reduces to a fresh location with an updated heap that maps l to x (REF).
- Dereferencing a variable using $!x$ is possible if x is bound to a location l by a let expression. If so, $!x$ reduces to $\sigma(l)$, the variable stored at location l (DEREF).

*the heap σ maps
locations to variables:*

$$\begin{aligned}
\sigma &:= \emptyset \\
&\mid \sigma[l \mapsto x]
\end{aligned}$$

- Similarly, if x is bound to l by a `let`, then the assignment operation $x := y$ updates the heap at location l with the variable y (HEAP).

Programs written in the Mutable DOT calculus generally do not contain explicit location values in the original program text. Locations are included as values in the Mutable DOT syntax only because terms such as `ref x T` will evaluate to fresh locations during reduction.

The remaining rules are the DOT evaluation rules, with the only change that they pass along a heap.

11.3 MUTABLE DOT TYPING RULES

The Mutable DOT typing rules, depicted in Figure 11.3, depend on a *heap typing* Σ in addition to a type environment Γ . A heap typing maps locations to the types of the variables that they store. The heap typing spares us the need to re-typecheck locations and allows to typecheck cyclic references (Pierce, 2002).

As an example, the following Mutable DOT program cannot be easily typechecked without an explicit heap typing (using only the runtime heap and the type environment):

$$\begin{aligned} & \text{let } \text{id} = \lambda(x : \top) x \quad \text{in} \\ p = & \text{let } r = \text{ref id } (\top \rightarrow \top) \quad \text{in} \\ & \text{let } \text{id}' = \lambda(x : \top) (!r) x \quad \text{in} \\ & r := \text{id}' \end{aligned}$$

Starting with an empty heap, after two reduction steps we get

$$\emptyset \mid p \longmapsto^* \{l \rightarrow \text{id}'\} \mid p',$$

where

$$\begin{aligned} & \text{let } \text{id} = \lambda(x : \top) x \quad \text{in} \\ p' = & \text{let } r = \boxed{l} \quad \text{in} \\ & \text{let } \text{id}' = \lambda(x : \top) (!r) x \quad \text{in} \\ & \boxed{\text{id}'} \end{aligned}$$

We would see by looking into the heap that to typecheck the location l , we needed to typecheck `id'`. `id'` depends on `r`, which in turn refers to the location l , creating a cyclic dependency.

We therefore augment our typing rules with a heap typing, allowing us to typecheck each location once and for all, at the time of a reference creation. In the example, we would know that l is mapped to $(\top \rightarrow \top)$ from the `let`-binding of `r` and remember this typing in Σ . To express that a term t has type T under the type environment Γ and heap typing Σ , we write $\Gamma, \Sigma \vdash t : T$.

Tight term typing
 $\Gamma, \Sigma \vdash t : T$

$$\begin{array}{c}
\frac{\Gamma(x) = T}{\Gamma, \Sigma \vdash x : T} \quad (\text{VAR}) \qquad \frac{\Gamma, \Sigma \vdash x : T}{\Gamma, \Sigma \vdash x : \mu(x : T)} \quad (\text{REC-I}) \\
\\
\frac{\Sigma(l) = T}{\Gamma, \Sigma \vdash l : \text{Ref } T} \quad (\text{LOC}) \qquad \frac{\Gamma, \Sigma \vdash x : \mu(x : T)}{\Gamma, \Sigma \vdash x : T} \quad (\text{REC-E}) \\
\\
\frac{\Gamma, x : T, \Sigma \vdash t : U \quad x \notin \text{fv}(T)}{\Gamma, \Sigma \vdash \lambda(x : T) t : \forall(x : T) U} \quad (\text{ALL-I}) \qquad \frac{\Gamma, \Sigma \vdash x : T \quad \Gamma, \Sigma \vdash x : U}{\Gamma, \Sigma \vdash x : T \wedge U} \quad (\&\text{-I}) \\
\\
\frac{\Gamma, \Sigma \vdash x : \forall(z : S) T \quad \Gamma, \Sigma \vdash y : S}{\Gamma, \Sigma \vdash x y : T[y/z]} \quad (\text{ALL-E}) \qquad \frac{\Gamma, \Sigma \vdash t : T \quad \Gamma, \Sigma \vdash T <: U}{\Gamma, \Sigma \vdash t : U} \quad (\text{SUB}) \\
\\
\frac{\Gamma, x : T, \Sigma \vdash d : T}{\Gamma, \Sigma \vdash \nu(x : T) d : \mu(x : T)} \quad (\{\}-\text{I}) \qquad \frac{\Gamma, \Sigma \vdash x : T}{\Gamma, \Sigma \vdash \text{ref } x T : \text{Ref } T} \quad (\text{REF-I}) \\
\\
\frac{\Gamma, \Sigma \vdash x : \{a : T\}}{\Gamma, \Sigma \vdash x.a : T} \quad (\{\}-\text{E}) \qquad \frac{\Gamma, \Sigma \vdash x : \text{Ref } T}{\Gamma, \Sigma \vdash !x : T} \quad (\text{REF-E}) \\
\\
\frac{\Gamma, \Sigma \vdash t : T \quad \Gamma, x : T, \Sigma \vdash u : U \quad x \notin \text{fv}(U)}{\Gamma, \Sigma \vdash \text{let } x = t \text{ in } u : U} \quad (\text{LET}) \qquad \frac{\Gamma, \Sigma \vdash x : \text{Ref } T \quad \Gamma, \Sigma \vdash y : T}{\Gamma, \Sigma \vdash x := y : T} \quad (\text{ASGN})
\end{array}$$

Definition typing
 $\Gamma, \Sigma \vdash d : T$

$$\begin{array}{c}
\frac{\Gamma, \Sigma \vdash t : T}{\Gamma, \Sigma \vdash \{a = t\} : \{a : T\}} \quad (\text{FLD-I}) \qquad \frac{\Gamma, \Sigma \vdash d_1 : T_1 \quad \Gamma, \Sigma \vdash d_1 : T_2 \quad \text{dom}(d_1), \text{dom}(d_2) \text{ disjoint}}{\Gamma, \Sigma \vdash d_1 \wedge d_2 : T_1 \wedge T_2} \quad (\text{ANDDEF-I}) \\
\\
\Gamma, \Sigma \vdash \{A = T\} : \{A : T..T\} \quad (\text{TYP-I})
\end{array}$$

Figure 11.3: Mutable DOT typing rules

The typing rules for Mutable DOT are shown in Figure 11.3. The DOT rules are intact except that all typing derivations carry a heap typing. The new rules related to mutable references are as follows:

- We typecheck locations by looking them up in the heap typing. If, according to Σ , a location l stores a variable of type T , then l has type $\text{Ref } T$ (LOC).
- A newly created reference $\text{ref } x \ T$ can be initialized with the variable x if x has type T . In particular, if x 's precise type U is a subtype of T , then x has type T by SUB, so we can still create a $\text{ref } x \ T$ (REF-I).
- Conversely, dereferencing a variable of a reference type $\text{Ref } T$ yields the type T (REF-E).
- Finally, if x is a reference of type $\text{Ref } T$, we are allowed to store a variable y into it if y has type T . To avoid the need to add a Unit type to the type system, we define an assignment $x := y$ to reduce to y , so the type of the assignment is T (ASGN).

11.4 SUBTYPING RULES

The subtyping rules of Mutable DOT include an added heap typing, and a subtyping rule for references. The rules are shown in Figure 11.4.

Subtyping between reference types is invariant: usually, $\text{Ref } T <: \text{Ref } U$ if and only if $T = U$. Invariance is required because reference types need to be (i) covariant for reading, or dereferencing, and (ii) contravariant for writing, or assignment.

However, in DOT, co- and contra-variance between types does not imply type equality: the calculus contains examples of types that are not equal, yet are equivalent with respect to subtyping. For example, for any types T and U , $T \wedge U <: U \wedge T <: T \wedge U$. Yet, $T \wedge U \neq U \wedge T$. Therefore, subtyping between reference types requires both covariance and contravariance:

$$\frac{\Gamma, \Sigma \vdash T <: U \quad \Gamma, \Sigma \vdash U <: T}{\Gamma, \Sigma \vdash \text{Ref } T <: \text{Ref } U} \quad (\text{REF-SUB})$$

$$\begin{array}{c}
\Gamma, \Sigma \vdash T <: \top \quad (\text{TOP}) \\
\Gamma, \Sigma \vdash \perp <: T \quad (\text{BOT}) \\
\Gamma, \Sigma \vdash T <: T \quad (\text{REFL}) \\
\frac{\Gamma, \Sigma \vdash S <: T \quad \Gamma, \Sigma \vdash T <: U}{\Gamma, \Sigma \vdash S <: U} \quad (\text{TRANS}) \\
\Gamma, \Sigma \vdash T \wedge U <: T \quad (\text{AND}_1\text{-<:}) \\
\Gamma, \Sigma \vdash T \wedge U <: U \quad (\text{AND}_2\text{-<:}) \\
\frac{\Gamma, \Sigma \vdash S <: T \quad \Gamma, \Sigma \vdash S <: U}{\Gamma, \Sigma \vdash S <: T \wedge U} \quad (\text{<:-AND}) \\
\frac{\Gamma, \Sigma \vdash x: \{A: S..T\}}{\Gamma, \Sigma \vdash S <: x.A} \quad (\text{<:-SEL}) \\
\frac{\Gamma, \Sigma \vdash x: \{A: S..T\}}{\Gamma, \Sigma \vdash x.A <: T} \quad (\text{SEL-<:})
\end{array}
\qquad
\begin{array}{c}
\frac{\Gamma, \Sigma \vdash S <: T \quad \Gamma, \Sigma \vdash S <: U}{\Gamma, \Sigma \vdash S <: T \wedge U} \quad (\text{<:-AND}) \\
\frac{\Gamma, \Sigma \vdash T <: U}{\Gamma, \Sigma \vdash \{a: T\} <: \{a: U\}} \quad (\text{FLD-<:-FLD}) \\
\frac{\Gamma, \Sigma \vdash S_2 <: S_1 \quad \Gamma, \Sigma \vdash T_1 <: T_2}{\Gamma, \Sigma \vdash \{A: S_1..T_1\} <: \{A: S_2..T_2\}} \quad (\text{TYP-<:-TYP}) \\
\frac{\Gamma, \Sigma \vdash S_2 <: S_1 \quad \Gamma, x: S_2, \Sigma \vdash T_1 <: T_2}{\Gamma, \Sigma \vdash \forall(x: S_1) T_1 <: \forall(x: S_2) T_2} \quad (\text{ALL-<:-ALL}) \\
\frac{\Gamma \vdash T <: U \quad \Gamma \vdash U <: T}{\Gamma \vdash \text{Ref } T <: \text{Ref } U} \quad (\text{REF-SUB})
\end{array}$$

Figure 11.4: Mutable DOT subtyping rules

TYPE SAFETY

In this section, we outline the soundness proof of Mutable DOT as an extension of the simple DOT soundness proof presented in Part I.

The type-safety proof of Mutable DOT involved adding the following definitions to the proof recipe (see Section 5.1):

1. a new case to the definition of *inert types*: any reference type $\text{Ref } T$ is inert;
2. *tight typing* extends the original tight typing definition with the new typing rules of Mutable DOT (exact versions of REF-I, REF-E, LOC, ASGN, and REF-SUB):

cf. *inert types* in Definition 2

cf. *tight typing* in Figure 5.1

$$\begin{array}{c}
 \frac{\Gamma, \Sigma \vdash_{\#} x : T}{\Gamma, \Sigma \vdash_{\#} \text{ref } x T : \text{Ref } T} \text{ (REF-I-}\#) \qquad \frac{\Gamma, \Sigma \vdash_{\#} x : \text{Ref } T}{\Gamma, \Sigma \vdash_{\#} !x : T} \text{ (REF-E-}\#) \\
 \\
 \frac{\Sigma(x) = T}{\Gamma, \Sigma \vdash_{\#} l : \text{Ref } T} \text{ (LOC-}\#) \qquad \frac{\Gamma, \Sigma \vdash_{\#} x : \text{Ref } T \quad \Gamma, \Sigma \vdash_{\#} y : T}{\Gamma, \Sigma \vdash_{\#} x := y : T} \text{ (ASGN-}\#) \\
 \\
 \frac{\Gamma, \Sigma \vdash_{\#} T <: U \quad \Gamma \vdash_{\#} U <: T}{\Gamma, \Sigma \vdash_{\#} \text{Ref } T <: \text{Ref } U} \text{ (REF-SUB-}\#)
 \end{array}$$

3. *invertible typing* extends the invertible rules with an additional case that inlines the REF-SUB subtyping rule:

cf. *invertible typing* in Figure 5.3

$$\frac{\Gamma, \Sigma \vdash_i x : \text{Ref } T \quad \Gamma \vdash_{\#} T <: U \quad \Gamma \vdash_{\#} U <: T}{\Gamma, \Sigma \vdash_i x : \text{Ref } U} \text{ (REF-i)}$$

4. *precise typing* extends the precise value-typing rules with a new case for locations:

cf. *precise typing* in Figure 5.2

$$\frac{\Sigma(x) = T}{\Gamma, \Sigma \vdash_{!} l : \text{Ref } T} \text{ (LOC-!)}$$

Since the typing relation depends on a heap typing, the *well-formedness* relation also needs to include Σ .

cf. *well-formedness* in Definition 16

DEFINITION 26 (Well-formed Environments). A store $\gamma = \overline{(x_i, v_i)}$ is *well-formed with respect to a type environment* $\Gamma = \overline{(x_i, T_i)}$ *and heap typing* Σ , written $\gamma : \Gamma, \Sigma$, if for each i , $\Gamma, \Sigma \vdash_{!} v_i : T_i$.

Additionally, we need to define well-formedness for heaps with respect to heap typings:

DEFINITION 27 (Well-Typed Heap). *A heap $\sigma = \overline{l_i \mapsto x_i}$ is well-typed with respect to an environment Γ and heap typing $\Sigma = \overline{l_i \mapsto T_i}$, written $\Gamma, \Sigma \vdash \sigma$, if for each i , $\Gamma, \Sigma \vdash x_i : T_i$.*

We can now present the central lemmas required to prove the Mutable DOT soundness theorems.

The helper lemmas for canonical forms (Lemma 11 (\forall to $\Gamma(x)$), Lemma 12 (\forall to λ), Lemma 13 (μ to $\Gamma(x)$), Lemma 14 (μ to ν), and Lemma 17 (Corresp. Types)) are left unchanged. However, we need two additional lemmas for variables and values that have reference types.

LEMMA 28 (Ref T to $\Gamma(x)$).

$$\frac{\text{inert } \Gamma \quad \Gamma, \Sigma \vdash x : \text{Ref } T}{\Gamma(x) = \text{Ref } U \quad \Gamma, \Sigma \vdash \text{Ref } U <: \text{Ref } T \quad \Gamma, \Sigma \vdash \text{Ref } T <: \text{Ref } U}$$

LEMMA 29 (Ref T to l).

$$\frac{\text{inert } \Gamma \quad \Gamma, \Sigma \vdash x : \text{Ref } T}{\Gamma(x) = \text{Ref } U \quad \Gamma, \Sigma \vdash \text{Ref } U <: \text{Ref } T \quad \Gamma, \Sigma \vdash \text{Ref } T <: \text{Ref } U}$$

With that we can prove canonical forms for mutable references.

LEMMA 30 (Canonical Forms for References). *Suppose that a store γ is well-formed with respect to an inert environment Γ and a heap typing Σ , that a heap σ is well-typed with respect to Γ and Σ , and that $\Gamma \vdash x : \text{Ref } T$. Then the store γ assigns a location l of type $\text{Ref } T$ to x , and l points to a variable y of type T in the heap.*

$$\frac{\gamma : \Gamma \quad \Gamma, \Sigma \vdash \sigma}{\text{inert } \Gamma \quad \Gamma, \Sigma \vdash x : \text{Ref } T} \quad \frac{\gamma(x) = l \quad \sigma(l) = y}{\Gamma, \Sigma \vdash l : \text{Ref } T} \quad \frac{\Gamma, \Sigma \vdash y : T}{(\text{Can. Forms for } l)}$$

Finally, we can present the progress and preservation theorems, as well as the ultimate soundness result.

$$\frac{\gamma : \Gamma \quad \Gamma, \Sigma \vdash \sigma}{\text{inert } \Gamma \quad \Gamma \vdash t : T} \quad \frac{\gamma \mid t \mapsto \gamma' \mid t'}{\Gamma \vdash t' : T} \quad \frac{\Gamma', \Sigma' \vdash \sigma'}{\gamma' : \Gamma' \quad \text{inert } \Gamma'} \quad (\text{PRESERVATION})$$

$$\frac{\gamma : \Gamma \quad \Gamma, \Sigma \vdash \sigma}{\text{inert } \Gamma \quad \Gamma \vdash t : T} \quad \frac{\gamma \mid t \mapsto \gamma' \mid t'}{(\text{PROGRESS})}$$

$$\frac{\emptyset, \Sigma \vdash t : T}{(\sigma \mid \gamma \mid t \mapsto \sigma' \mid \gamma' \mid t' \wedge t' \twoheadrightarrow) \vee t \Uparrow} \quad (\text{Mutable DOT Soundness})$$

THEOREM 31 (Mutable DOT Preservation). *Let Γ be an inert typing environment and Σ a heap typing such that $\gamma : \Gamma, \Sigma$ and $\Gamma, \Sigma \vdash \sigma$. If $\Gamma, \Sigma \vdash t : T$ and $\sigma \mid \gamma \mid t \mapsto \sigma' \mid \gamma' \mid t'$ for some store γ' and heap σ' then there exists an inert context Γ' and heap typing Σ' such that $\gamma' : \Gamma', \Sigma'$, $\Gamma', \Sigma' \vdash \sigma'$, and $\Gamma', \Sigma' \vdash t' : T$.*

THEOREM 32 (Mutable DOT Progress). *Let γ be a store, σ a heap, Γ an inert typing environment, and Σ a heap typing such that $\gamma : \Gamma, \Sigma$ and $\Gamma, \Sigma \vdash \sigma$. If $\Gamma, \Sigma \vdash t : T$ then there exists a term t' , store γ' , and heap σ' such that $\sigma \mid \gamma \mid t \mapsto \sigma' \mid \gamma' \mid t'$.*

THEOREM 33 (Mutable DOT Type Soundness). *If $\emptyset, \Sigma \vdash t : T$ then either t diverges ($t \Uparrow$) or t reduces to a normal form t' , i.e. $\sigma \mid \emptyset \mid t \mapsto \sigma' \mid \gamma' \mid t'$, $t' \twoheadrightarrow$, and $\Gamma, \Sigma' \vdash t' : T$ for some Γ and Σ' such that $\gamma' : \Gamma, \Sigma'$.*

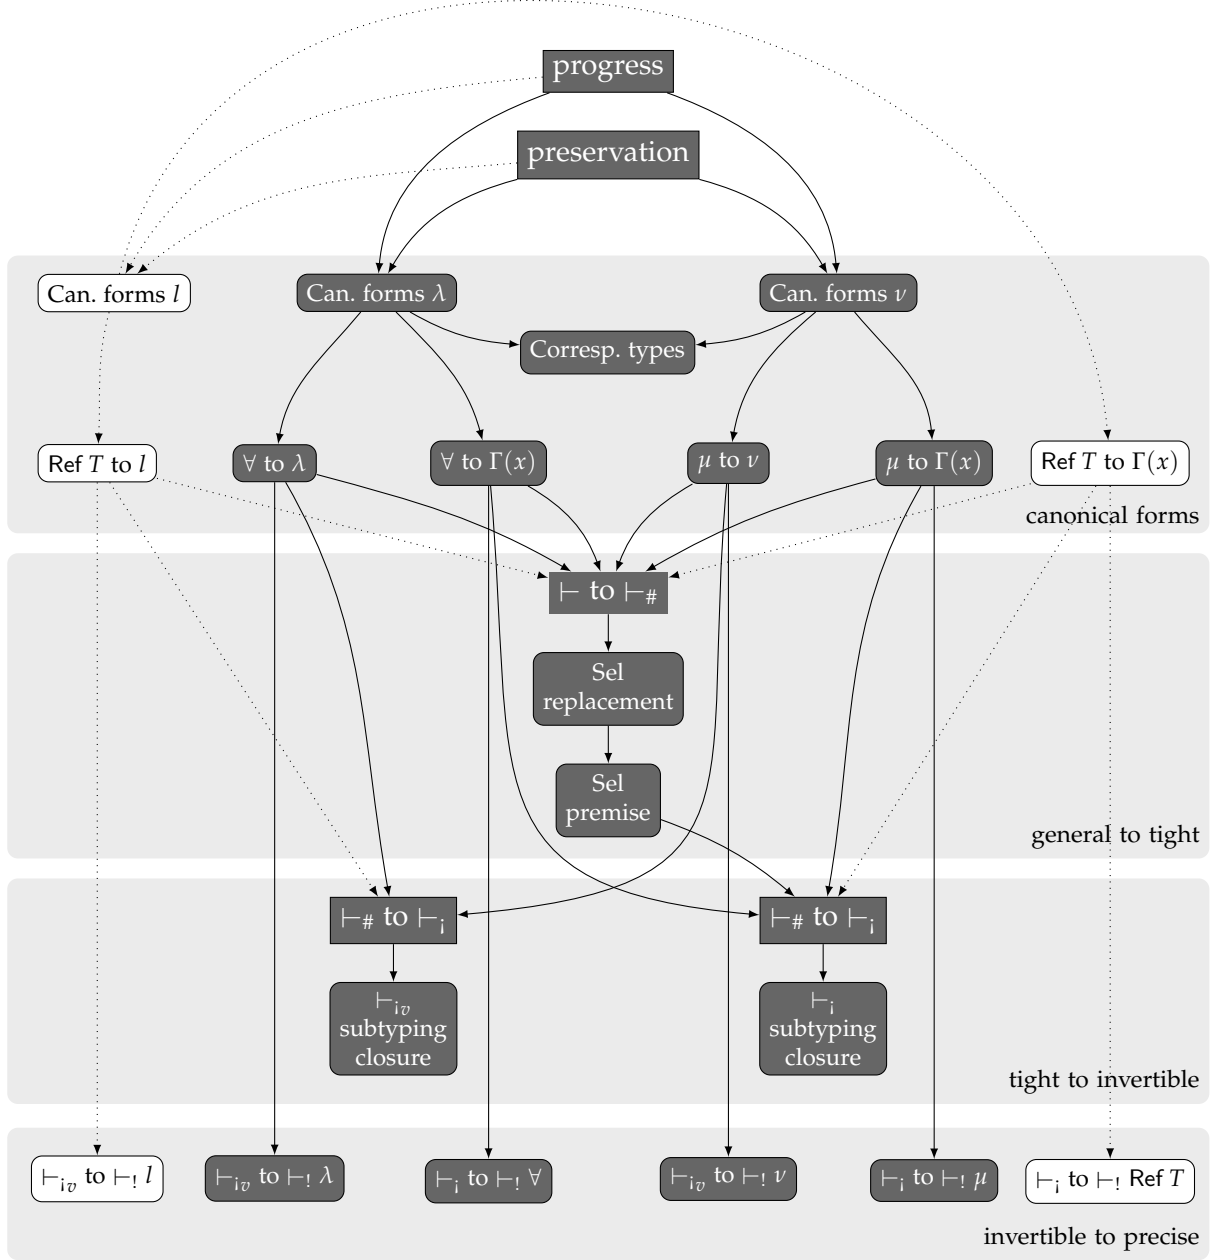


Figure 12.1: An instance of the dependency graph from Figure 5.4 showing the main lemmas in the Mutable DOT proof as an extension of the simple DOT proof (Part I). Gray nodes denote existing lemmas. White nodes denote Mutable DOT specific lemmas

The changes to the lemmas proof are shown using white nodes in the dependence graph in Figure 12.1.

The overall structure of the dependencies between the lemmas did not change. The new canonical forms lemmas followed the proof recipe that we have described in Chapter 5. In the proofs of some lemmas, we had additional new cases to prove, but the structure of the proof of each lemma did not change. In general, we found that the new lemmas and new cases in the existing lemmas were easy to prove.

DISCUSSION

In this section, we explain the design choices of Mutable DOT in more detail and discuss possible alternative implementations.

13.1 MOTIVATION FOR A HEAP OF VARIABLES

One unusual aspect of the design of Mutable DOT is that the heap contains variables rather than values. We experimented with alternative designs that contained values, and observed the following undesirable interactions with the existing design of DOT.

A key desirable property is that the heap should be well-typed with respect to a heap typing: $\forall l. \Gamma, \Sigma \vdash \sigma(l) : \Sigma(l)$.

Many of the DOT type assignment rules apply only to variables, and not to values. For example, the type $\{a : \top\}$ is not inhabited by any value, but a variable can have this type. This is because an object value has a recursive type, and the REC-E rule that opens a recursive type $\mu(x : \{a : \top\})$ into $\{a : \top\}$ applies only to variables, not to values. In particular, in the term

$$\text{let } x = \nu(y : \{a : \top\})\{a = t\} \text{ in ref } x \{a : \top\}$$

x has type $\{a : \top\}$ but $\nu(y : \{a : \top\})\{a = t\}$ does not, even though the let binding suggests that the variable and the value should be equal. If memory cells were to contain values, a cell of type $\{a : \top\}$ would not make sense, because no values have that type.

We could try to restrict reference types to always store recursive (or function) types. However, this would severely restrict the polymorphism of memory cells, because DOT does not support subtyping between recursive types (subtyping between recursive structural types is not supported by Scala either). In particular, it would be impossible to define a memory cell containing objects with a field a of type \top and possibly additional fields.

The above example let term demonstrates another problem: type preservation. The type system should admit the term $\text{ref } x \{a : \top\}$ because x has type $\{a : \top\}$. This term should reduce to a fresh location l of type $\text{Ref } \{a : \top\}$. But a heap that maps l to $\nu(y : \{a : \top\})\{a = t\}$ would not be well typed, because the value does not have type $\{a : \top\}$.

13.2 CORRECTNESS OF A HEAP OF VARIABLES

Putting variables instead of values in the heap raises a concern: when we write a variable into the heap, we expect that when we read it back, it will still be in scope, and it will still be bound to the same value. For example, in the following program fragment, the variable x gets saved in the heap inside the function f .

```

let  $f$   =   $\lambda(x: \top) \text{ref } x \ T$   in
let  $y$   =   $v$                     in
let  $r$   =   $f \ y$                   in
 $!r$ 

```

Will x go out of scope by the time we read it from the heap?

The reduction sequence for this program is shown in Figure 13.1. Notice that before the body $\text{ref } x \ T$ of the function is reduced, the parameter x is first substituted with the argument y , which does not go out of scope.

\emptyset		$f \mapsto \lambda(x: \top) \text{ref } x \ T, y \mapsto v$		let $r = f \ y$ in $!r$	\mapsto
\emptyset		$f \mapsto \lambda(x: \top) \text{ref } x \ T, y \mapsto v$		let $r = \text{ref } x \ T[y/x]$ in $!r$	\mapsto
\emptyset		$f \mapsto \lambda(x: \top) \text{ref } x \ T, y \mapsto v$		let $r = \text{ref } y \ T$ in $!r$	\mapsto
$l \mapsto y$		$f \mapsto \lambda(x: \top) \text{ref } x \ T, y \mapsto v$		let $r = l$ in $!r$	\mapsto
$l \mapsto y$		$f \mapsto \lambda(x: \top) \text{ref } x \ T, y \mapsto v, r \mapsto l$		$!r$	\mapsto
$l \mapsto y$		$f \mapsto \lambda(x: \top) \text{ref } x \ T, y \mapsto v, r \mapsto l$		y	

Figure 13.1: Reduction sequence for example program

More generally, from the store-based reduction semantics, it is immediately obvious that when a variable x is saved in the heap using $\text{ref } x \ T$ or $y := x$, the only variables that are in scope are those in the store. There are no function parameters in scope that could go out of scope when the function finishes.

Moreover, once a variable is in the store, it never goes out of scope, and the value that it is bound to never changes. This is because the only reduction rule that modifies the store is LET-VALUE, and it only adds a new variable binding, but does not affect any existing bindings.

Another natural question is whether a heap of variables limits the expressiveness of the calculus. Since a program contains only a finite number of variables, one might think that the size of the heap is restricted by that number. However, during execution, the reduction rule for function application performs capture-avoiding substitution using alpha renaming, which introduces fresh variables as necessary. Thus, the use of variables in the heap does not impose any restrictions on the number of objects that can be created.

$$\begin{array}{c}
 \sigma \mid \gamma \mid \text{let } x = v \text{ in } t \\
 \mapsto \\
 \sigma \mid \gamma, x \mapsto v \mid t \\
 \text{(LET-VALUE)}
 \end{array}$$

13.3 CREATING REFERENCES

The Mutable DOT reference creation term $\text{ref } x \ T$ requires both a type T and an initial variable x . The variable is needed so that a reference cell is always initialized, to avoid the need to add a null value to DOT. If desired, it is possible to model uninitialized memory cells in Mutable DOT by explicitly creating a sentinel null value.

Some other calculi with mutable references (e.g. Types and Programming Languages (Pierce, 2002)) do not require the type T to be given explicitly, but just adopt the precise type of x as the type for the new cell. Such a design does not fit well with subtyping in DOT. In particular, it would prevent the creation of a cell with some general type T initialized with a variable x of a more specific subtype of T .

More seriously, such a design (together with subtyping) would break type preservation. Suppose that $\Gamma, \Sigma \vdash y : S$ and $\Gamma, \Sigma \vdash S <: T$. Then we could arrive at the following reduction sequence:

$$\begin{array}{lcl}
 \emptyset & | & f \mapsto \lambda(x : T) \text{ref } x, y \mapsto v \quad | \quad f y \quad \mapsto \\
 \emptyset & | & f \mapsto \lambda(x : T) \text{ref } x, y \mapsto v \quad | \quad \text{ref } x [y/x] \quad \mapsto \\
 \emptyset & | & f \mapsto \lambda(x : T) \text{ref } x, y \mapsto v \quad | \quad \text{ref } y
 \end{array}$$

The term at the beginning of the reduction sequence has type $\text{Ref } T$, while the term at the end, $\text{ref } y$, has type $\text{Ref } S$. Preservation would require $\text{Ref } S$ to be a subtype of $\text{Ref } T$, but this is not the case in general since the only condition that this example imposes on S and T is that $\Gamma, \Sigma \vdash S <: T$.

RELATED WORK

This part of the thesis shows how to extend the DOT calculus with ML-style mutable references, to serve as a basis for further extensions that involve mutation. An earlier version of the Mutable DOT calculus with a soundness proof that extended the original proof by Amin, Grütter, et al., (2016) appeared in a previous publication (Rapoport and Lhoták, 2017) and a technical report (Rapoport and Lhoták, 2016).

Amin and Rompf, (2017) present mechanized soundness proofs using definitional interpreters for big-step DOT-like calculi. The paper and an earlier technical report (Rompf and Amin, 2016a) mention that mutable references can be added to this class of calculi and present a Coq formalization of System $F_{<}$ with mutable references. To our knowledge, the formalization does not include a soundness proof for $D_{<}$ or DOT with mutation. Additionally, the calculi discussed in the papers are based on a big-step semantics, whereas our work focuses on a small-step semantics for DOT.

Kabir and Lhoták, (2018) present κ DOT, an extension of DOT with constructors and mutable fields. Unlike DOT, κ DOT's objects support strict initialization of fields, which corresponds closer to Scala than DOT's field semantics. κ DOT's fields are mutable by default, and they are not explicitly typed as references, unlike in our version of Mutable DOT. Still, κ DOT can express read-only fields through setting a field's lower-bound type to \perp . The difference between Mutable DOT and κ DOT is that the latter focuses on modelling constructors and makes all fields mutable by default, whereas Mutable DOT presents the first and minimal addition to DOT with mutation, and supports both mutable fields and variables.

SUMMARY

DOT formalizes the essence of Scala, but it lacks mutation, which is an important feature of object-oriented languages. In this part of the thesis we showed how DOT can be extended to handle mutation in a type-safe way.

We have seen that adding a mutable heap to the semantics of DOT is not straightforward. The lack of subtyping between recursive types leads to situations where variables and values, even though they are bound together, have incompatible types. As a result, if DOT were extended with a conventional heap containing values, it would be impossible for a cell of a given type T to store values of different subtypes of T , thus significantly restricting the kinds of mutable code that could be expressed.

Our key idea was to enable support for mutation in DOT by using a heap that contains variables instead of values. We showed that by using a heap of variables, it is possible to extend DOT with mutable references in a type-safe way. This leads to a formalization of a language with path-dependent types and mutation, and also brings DOT one step closer to encoding the full Scala language.

Although designing Mutable DOT involved non-trivial reasoning about how to maintain preservation while allowing a mutable heap, *proving* type-safety of Mutable DOT involved a straightforward extension of the simple type-safety proof presented in Part I. The extension of DOT with mutation serves as evidence that the simple proof is indeed simple. In the next part, we present an extension of DOT that is more complex and involves significant changes to the proof. Nevertheless, as we will show in Part III, the basic organization and key ideas of the proof remain the same.

Part III

FULLY PATH-DEPENDENT TYPES

INTRODUCTION

Path-dependent types embody two universal principles of modular programming: abstraction and composition. Abstraction allows us to leave values or types in a program unspecified to keep it generic and reusable. For example, in Scala, we can define trees where the node type remains abstract:

$\underbrace{\text{path-dependent}}_{\text{composition}} \quad \underbrace{\text{type}}_{\text{abstraction}}$

```
trait Tree {
  type Node
  val root: Node
  def add(node: Node): Tree
}
```

If an object x has type `Tree`, then the path-dependent type `x.Node` denotes the type of abstract nodes.

Composition is the ability to build our program out of smaller components. For example, if we are interested in a specific kind of tree, say a red-black tree, then we can refine the abstract `Node` type to contain a `Color` type:

```
trait RedBlackTree extends Tree {
  type Node <: { type Color }
}
```

This exemplifies composition in at least two ways: by having `RedBlackTree` extend `Tree` we have *inherited* its members; and by nesting the refined definition of `Node` within `RedBlackTree` we have used *aggregation*. If an object r is a `RedBlackTree`, then the path-dependent type `r.root.Color` allows us to traverse the composition and access the `Color` type member.

As described in the previous chapters, the long struggle to formalize path-dependent types recently led to machine-verified soundness proofs for several variants of the [DOT](#) calculus. In spite of its apparent simplicity DOT is an expressive calculus that can encode a variety of language features, and the discovery of its soundness proof was a breakthrough for the Scala community.

However, a crucial limitation is that the existing DOT calculi restrict path-dependent types to depend only on variables, not on general paths. That is, they allow the type `x.Node` (path of length 1) but not a longer type such as `r.root.Color` (length 2). We need to lift this restriction in order to faithfully model Scala which does allow general path-dependent types. More importantly, this restriction must be lifted to fulfill the goal of *scalable component abstraction* (Odersky and Zenger, 2005b), in which modules of a program can be arbitrarily nested to form other, larger modules.

Scala: <pre> package dotty { package core { object types { class Type class TypeRef extends Type { val symb: core.symbols.Symbol } } object symbols { class Symbol { val tpe: core.types.Type } } } } </pre>	DOT: <pre> let dotty = v(dotty) { core = v(core) { types = v(types) { Type = ... TypeRef = Type ^ { symb: <u>core.symbols.Symbol</u> } } symbols = v(symbols) { Symbol = { tpe: <u>core.types.Type</u> } } } in ... </pre>
--	--

Figure 16.1: A simplified excerpt from the Dotty compiler in Scala. This code fragment cannot be expressed in DOT, as shown on the right

The final part of this thesis formalizes and proves sound a generalization of the DOT calculus with path-dependent types of arbitrary length. We call the new path-dependent calculus pDOT. Our Coq-verified proof is built on top of the simple proof presented in Part I.

At this point, two questions naturally arise. Are fully path-dependent types really necessary? That is, do they provide additional expressiveness, or are they just syntactic sugar over variable-dependent types? And if fully path-dependent types are in fact useful, what are the barriers to adding them to DOT?

WHY FULLY PATH-DEPENDENT TYPES ARE NECESSARY

The need for paths of arbitrary length is illustrated by the simplified excerpt from the implementation of the Scala 3 (“Dotty”) compiler in Figure 16.1. Type references (TypeRef) are Types that have an underlying class or trait definition (Symbol), while Symbols in the language also have a Type. Additionally, TypeRefs and Symbols are nested in different packages, `core.types` and `core.symbols`.

It is impossible to express the above type dependencies in DOT while maintaining the nested program structure, as shown on the right side of Figure 16.1 in DOT syntax (see Chapter 2 for a description of the DOT syntax and type system). To replicate nested Scala modules, DOT uses objects and fields. Unfortunately, we run into problems when typing the `symb` field because its desired path-dependent type `core.symbols.Symbol` has a path of length two.

We are then tempted to find a workaround. One option is to try to reference `Symbol` as a path-dependent type of length one: `symbols.Symbol` instead of `core.symbols.Symbol`. However, this will not do because `symbols`

is a field, and DOT requires that field accesses happen through the enclosing object (core). Another option is to move the definition of the Symbol type member to the place it is accessed from, to ensure that the path to the type member has length 1:

```
types = v(types) {
  Type = ...;
  Symbol = ...;
  TypeRef = Type  $\wedge$  { symb: types.Symbol }
}
```

However, such a transformation would require flattening the nested structure of the program whenever we need to use path-dependent types. This would limit encapsulation and our ability to organize a program according to its logical structure. Yet another approach is to assign the symbols object to a variable that is defined before the dotted object:

```
let symbols = v(symbols) {
  Symbol = { tpe: dotted.core.types.Type }
}
in let dotted = v(dotted) ...
```

This attempt fails as well, as the symbols object can no longer reference the dotted package. For the above example this means that a Symbol cannot have a Type.

This real-world pattern with multiple nested modules and intricate dependencies between them (sometimes even *recursive* dependencies, as in our example), leads to path-dependent types of length greater than one. Because path-dependent types are used in DOT to formalize features like parametric and family polymorphism (Ernst, 2001), covariant specialization (Bruce, Odersky, and Wadler, 1998), and wildcards, among others, a version of DOT with just variable-dependent types can only formalize these features in special cases. Thus, to unleash the full expressive power of DOT we need path-dependent types on paths of arbitrary length.

see Section 17.1 for a minimal example that illustrates this limitation

variable-dependent type
limited composition abstraction

WHY FULLY PATH-DEPENDENT TYPES ARE HARD

The restriction to types dependent on variables rather than paths is not merely cosmetic; it is fundamental. A key challenge in formalizing the DOT calculi is the *bad bounds* problem, discussed in Chapter 4: the occurrence of a type member in a program introduces new subtyping relationships, and these subtyping relationships could undermine type safety in the general case. To maintain type safety, the existing DOT calculi ensure that whenever a type $x.A$ is in scope, any code in the same scope will not execute until x has been assigned some concrete value; the value serves as evidence that type soundness has not been subverted. If we allow a type to depend on a path, rather than a

Section 17.2.1 shows why extending DOT with path can lead to bad bounds

variable, we must extend this property to paths: we must show that whenever a scope admits a given path, that path will always evaluate to some stable value. The challenge of ensuring that the paths of type-selections always evaluate to a value is to rule out the possibility that paths cyclically alias each other, while at the same time keeping the calculus expressive enough to allow recursion. By contrast, the DOT calculus automatically avoids the problem of type selections on non-terminating paths (i.e. paths whose evaluation does not terminate) because in DOT all paths are variables, and variables are considered normal form.

A second challenge of extending DOT with support for general paths is to track *path equality*. Consider the following program:

```
val t1 = new ConcreteTree
val t2 = new ConcreteTree
val t3 = t2
```

A subclass of `Tree` such as `ConcreteTree` (not shown) refines `Node` with a concrete type that implements some representation of nodes. We want the types `t1.Node` and `t2.Node` to be considered distinct even though `t1` and `t2` are both initialized to the same expression. That way, we can distinguish the nodes of different tree instances. On the other hand, notice that the reference `t3` is initialized to be an alias to the same tree instance as `t2`. We therefore want `t2.Node` and `t3.Node` to be considered the same type.

How can the type system tell the difference between `t1.Node` and `t2.Node`, so that the former is considered distinct from `t3.Node`, but the latter is considered the same? Scala uses *singleton types* for this purpose. In Scala, `t3` can be typed with the singleton type `t2.type` which guarantees that it is an alias for the same object as `t2`. The type system treats paths that are provably aliased (as evidenced by singleton types) as interchangeable, so it considers `t2.Node` and `t3.Node` as the same type. We add singleton types to pDOT for two reasons: first, we found singleton types useful for formalizing path-dependent types, and second, enabling singleton types brings DOT closer to Scala.

The contributions are as follows:

1. The *pDOT* calculus, a generalization of DOT with *path-dependent types of arbitrary length* that lifts DOT's type-selection-on-variables restriction.
2. The first extension of DOT with *singleton types*, a Scala feature that, in addition to tracking path equality, enables the method chaining pattern and hierarchical organization of components (Odersky and Zenger, 2005b).
3. A *Coq-mechanized type soundness proof* of pDOT that is based on the simple soundness proof presented in Part I. Our proof maintains the simple proof's modularity properties which makes it easy to extend pDOT with new features.

Chapter 18 provides
an intuition for
pDOT's main ideas;
Chapter 19 presents
the calculus in detail

Coq proof:
<https://git.io/dotpaths>
described in
Chapter 21

4. *Formalized examples* that illustrate the expressive power of pDOT: *see Chapter [20](#)*
the compiler example from this section that uses general path-dependent types, a method chaining example that uses singleton types, and a covariant list implementation.

CHALLENGES OF ADDING PATHS TO DOT

This chapter shows why the existing DOT calculus cannot encode path-dependent types on paths of arbitrary length and describes how naively extending DOT with support for full paths leads to bad bounds.

17.1 PATH LIMITATIONS IN DOT: A MINIMAL EXAMPLE

Consider the following example DOT object in which a type member B refers to a type member A that is nested inside the definition of a field c :

$$\begin{aligned} \text{let } x = v(x) \\ \quad \{c = v(_) \{A = x.B\}\} \wedge \\ \quad \{B = x.c.A\} \text{ in } \dots \end{aligned}$$

In the example, to reference the field c , we must first select the field's enclosing object x through its self variable. As a result, the path to A leads through $x.c$ which is a path of length two. Since DOT does not allow paths of length two, this definition of B cannot be expressed in DOT without flattening the program structure so that all fields and type members become global members of one top-level object.

Since DOT uses ANF, we may try to decompose the path of length two into dereferences of simple variables, but as we will see, this will fail. It does not work to first assign $x.c$ to a local variable y outside of the object x and then use the type $y.A$:

$$\begin{aligned} \text{let } y = x.c & \quad \text{in} \\ \text{let } x = v(x) & \\ \quad \{c = v(_) \{A = x.B\}\} \wedge & \\ \quad \{B = y.A\} & \quad \text{in } \dots \end{aligned}$$

This program is invalid because at the declaration site of y , x is not yet defined. We could try other ways of let-binding the inner objects to variables before defining the enclosing object, but all such attempts are doomed to failure. A sequence of let bindings imposes a total ordering on the objects and restricts an object to refer only to objects that are defined before it. In the presence of recursive references between the objects, as in this example, no valid ordering of the let bindings is possible.

17.2 CHALLENGES OF ADDING PATHS TO DOT

If restricting path-dependent types exclusively to variables limits the expressivity of DOT then why does the calculus impose such a constraint? Before we explain the soundness issue that makes it difficult to extend DOT with paths we must first review the problem of bad bounds, a key challenge that makes it difficult to ensure soundness of the DOT calculus.

As discussed in Chapter 4, Scala’s abstract type members make it possible to define custom subtyping relationships between types. This is a powerful but tricky feature. For example, given any types S and U , consider the function $\lambda(x: \{A: S..U\}) t$. In the body of the function, we can use $x.A$ as a placeholder for some type that is a supertype of S and a subtype of U . Some concrete type will be bound to $x.A$ when the function is eventually called with some specific argument. Due to transitivity of subtyping, the constraints on $x.A$ additionally introduce an assumption inside the function body that $S <: U$, because $S <: x.A <: U$ according to the type rules $<:-\text{SEL}_{\text{DOT}}$ and $\text{SEL-}<:\text{DOT}$:

*We add an explicit
DOT annotation to
DOT rules to
distinguish them
from pDOT rules*

$$\frac{\Gamma \vdash x: \{A: S..T\}}{\Gamma \vdash S <: x.A} (<:-\text{SEL}_{\text{DOT}}) \qquad \frac{\Gamma \vdash x: \{A: S..T\}}{\Gamma \vdash x.A <: T} (\text{SEL-}<:\text{DOT})$$

However, recall that S and U are arbitrary types, possibly with no existing subtyping relationship. The key to soundness is that although the function body is type-checked under the possibly unsound assumption $S <: U$, the body executes only when the function is called, and calling the function requires an argument that specifies a concrete type T to be bound to $x.A$. This argument type must satisfy the constraints $S <: T <: U$. Thus, the argument type embodies a form of *evidence* that the assumption $S <: U$ which is used to type-check the function body is actually valid.

More generally, given a term t of type $\{A: S..U\}$, we can rule out the possibility of bad bounds caused by the use of a dependent type $t.A$ if there exists some object with the same type $\{A: S..U\}$. This is because the object must bind the type member A to some concrete type T respecting the subtyping constraints $S <: T$ and $T <: U$, so the object is evidence that $S <: U$.

Existing DOT calculi ensure that whenever some variable x of type T is in scope in some term t , the term reduces only after x has already been assigned a value. The value assigned to x is evidence that T does not have bad bounds. To ensure that any code that uses the type $x.A$ executes only after x has been bound to a value of a compatible type, DOT employs a strict operational semantics. A variable x can be introduced by one of the three binding constructs: **let** $x = t$ **in** u , $\lambda(x: T) t$, or $\nu(x: T) d$. In the first case, x is in scope within u , and the reduction semantics requires that before u can execute, t must

first reduce to a value with the same type as x . In the second case, x is in scope within t , which cannot execute until an argument value is provided for the parameter x . In the third case, the object itself is bound to the self variable x . In summary, the semantics ensures that by the time that evaluation reaches a context with x in scope, x is bound to a value, and therefore x 's type does not introduce bad bounds.

17.2.1 Naive Path Extension Leads to Bad Bounds

When we extend the type system with types $p.A$ that depend on paths rather than variables, we must take similar precautions to control bad bounds. If a path p has type $\{A: S..U\}$ and some normal form n also has this type, then n must be an object that binds to type member A a type T such that $S <: T <: U$.

However, not all syntactic paths in DOT have this property. For example, in an object $v(x) \{a = t\}$, where t can be an arbitrary term, t could loop instead of reducing to a normal form of the same type. In that case, there is no guarantee that a value of the type exists, and it would be unsound to allow the path $x.a$ as the prefix of a path-dependent type $x.a.A$.

The following example, in which a function $x.b$ is typed as an object (a record with field c), demonstrates this unsoundness:

$$\begin{array}{ll} v(x : \{a : \{C : \forall(y : \top) \top.. \{c : \top\}\} \} \wedge \{b : \{c : \top\}\}) & \\ \{a = x.a\} & \wedge \{b = \lambda(y : \top) y\} \end{array}$$

Here, $x.b$ refers to a function $\lambda(y : \top) y$ of type $\forall(y : \top) \top$. If we allowed such a definition, the following would hold:

$$\forall(y : \top) \top <: x.a.C <: \{c : \top\}.$$

Then by subsumption, $x.b$, a function, has type $\{c : \top\}$ and therefore it must be an object. To avoid this unsoundness, we have to rule out the type selection $x.a.C$ on the non-terminating path $x.a$.

In general, if a path p has a field declaration type $\{a : T\}$, then the extended path $p.a$ has type T , but we do not know whether there exists a value of type T because $p.a$ has not yet reduced to a variable. Therefore, the type T could have bad bounds, and we should not allow the path $p.a$ to be used in a path-dependent type $p.a.A$.

The main difficulty we encountered in designing pDOT was to ensure that type selections occur only on terminating paths while ensuring that the calculus still permits non-terminating paths in general, since that is necessary to express recursive functions and maintain Turing completeness of the calculus.

This chapter outlines the main ideas that have shaped our definition of pDOT. To distinguish the DOT typing rules by Amin, Grütter, et al., (2016) presented in Chapter 2 from the pDOT typing rules we will postfix the DOT typing rules with “_{DOT}”. For example, VAR_{DOT} is a DOT typing rule, but VAR is a pDOT typing rule.

Chapter 19 presents
pDOT detail

18.1 PATHS INSTEAD OF VARIABLES

To support fully-path-dependent types, our calculus needs to support paths in all places where DOT permitted variables. Consider the following example:

let $x = \nu(y) \{a = \nu(z) \{B = U\}\}$
in $x.a$

In order to make use of the fact that $U <: x.a.B <: U$, we need a type rule that reasons about path-dependent types. In DOT, this is done through the $\text{SEL-}<:_{\text{DOT}}$ and $<:-\text{SEL}_{\text{DOT}}$ rules, as mentioned in Chapter 4. Since we need to select B on a path $x.a$ and not just on a variable x , we need to extend the rules (merged into one here for brevity) to support paths:

$$\frac{\Gamma \vdash x : \{A : S..T\}}{\Gamma \vdash S <: x.A <: T} <:-\text{SEL-}<:_{\text{DOT}} \quad \Rightarrow \quad \frac{\Gamma \vdash p : \{A : S..T\}}{\Gamma \vdash S <: p.A <: T} <:-\text{SEL-}<:$$

However, before we can use this rule we need to also generalize the recursion elimination rule $\text{Rec-E}_{\text{DOT}}$. In the above example, how do we obtain the typing $\Gamma \vdash x.a : \{B : U..U\}$? The only identifier of the inner object is $x.a$, a path. The type of the path is $\mu(z : \{B : U..U\})$. In order to use the type member B , it is necessary to specialize this recursive type, replacing the recursive self variable z with the path $x.a$. This is necessary because the type U might refer to the self variable z , which is not in scope outside the recursive type. Thus, in order to support path-dependent types, it is necessary to allow recursion elimination on objects identified by paths:

$$\frac{\Gamma \vdash x : \mu(y : T)}{\Gamma \vdash x : T[x/y]} \text{REC-E}_{\text{DOT}} \quad \Rightarrow \quad \frac{\Gamma \vdash p : \mu(y : T)}{\Gamma \vdash p : T[p/y]} \text{REC-E}$$

By similar reasoning, we need to generalize all DOT variable-typing rules to path-typing rules. As we show later, we also have to generalize DOT’s ANF syntax to use paths wherever DOT uses variables, so all

the DOT reduction rules that operate on variables are generalized to paths in pDOT.

18.2 PATHS AS IDENTIFIERS

A key design decision of pDOT is to let *paths represent object identity*. In DOT, object identity is represented by variables, which works out because variables are irreducible. In pDOT, *paths are irreducible*, because reducing paths would strip objects of their identity and break preservation.

18.2.1 Variables are Identifiers in DOT

In the DOT calculus by Amin, Grütter, et al., (2016), variables do not reduce to values for two reasons:

- *type safety*: making variables irreducible is necessary to maintain preservation, and
- *object identity*: to access the members of objects (which can recursively reference the object itself), objects need to have a name; reducing variables would strip objects of their identity.

If variables in DOT reduced to values, then in the previous example program,⁴ x would reduce to

$$\begin{array}{l} {}^4 \text{let } x = v(y) \\ \{a = v(z) \{B = U\}\} \\ \text{in } x.a \end{array} \quad v = v(y) \{a = v(z) \{B = U\}\}.$$

To maintain type preservation, for any type T such that $\Gamma \vdash x : T$, we also must be able to derive $\Gamma \vdash v : T$. Since

$$\Gamma \vdash x : \mu(y : \{a : \mu(z : \{B : U..U\})\}),$$

by recursion elimination $\text{Rec-E}_{\text{DOT}}$,

$$\Gamma \vdash x : \{a : \mu(z : \{B : U[x/y]..U[x/y]\})\}.$$

Does v also have that type? No!

$$\frac{\frac{\Gamma \vdash x : \mu(y : \{a : \mu(z : \{B : U..U\})\})}{\Gamma \vdash x : \{a : \mu(z : \{B : U[x/y]..U[x/y]\})\}} \text{Rec-E}_{\text{DOT}} \quad \gamma : \Gamma \quad \frac{\gamma(x) = v}{\gamma \mid x \mapsto \gamma \mid v} \text{HYPOTHETICAL REDUCTION}}{\Gamma \vdash v : \{a : \mu(z : \{B : U[x/y]..U[x/y]\})\}} \text{PRESERVATION}_{\text{DOT}}$$

The value v has only the *recursive* type

$$\mu(y : \{a : \mu(z : \{B : U..U\})\}).$$

Since v is no longer connected to any specific name, no recursion elimination is possible on its type. In particular, it does not make sense to give this value the type

$$\{a: \mu(z: \{B: U[x/y]..U[x/y]\})\}$$

because this type refers to x , but after the reduction, the value is no longer associated with this name.

The example illustrates that in DOT, variables represent the identity of objects. This is necessary in order to access an object's members: object members can reference the object itself, for which the object needs to have a name.

18.2.2 Paths are Identifiers in pDOT

In pDOT, *paths* represent the identity of objects and therefore they must be irreducible. Similarly to DOT, reducing paths would lead to unsoundness and strip nested objects of their identity. Making paths irreducible means that in pDOT, we cannot have an analog of DOT's field selection rule Proj_{DOT} .

Consider the field selection $x.a$ from the previous example. What is its type? By recursion elimination,

$$\Gamma \vdash x.a: \{B: U[x.a/z]..U[x.a/z]\}$$

If pDOT had a path-reduction rule PROJ analogous to DOT's PROJ_{DOT} , then $x.a$ would reduce to $v(z)\{B = U\}$. However, that value does not have the type $\{B: U[x.a/z]..U[x.a/z]\}$; it only has the recursive type $\mu(z: \{B: U..U\})$.

$$\frac{\frac{\Gamma \vdash x.a: \mu(z: \{B: U..U\})}{\Gamma \vdash x.a: \{B: U[x.a/z]..U[x.a/z]\}} \text{REC-E} \quad \gamma: \Gamma \quad \frac{\gamma(x) = v(y)\{a = v(z)\{B = U\}\}}{\gamma \mid x.a \mapsto \gamma \mid v(z)\{B = U\}} \text{HYPOTHETICAL PROJ}}{\Gamma \vdash v(z)\{B = U\}: \{B: U[x.a/z]..U[x.a/z]\}} \text{PRESERVATION}$$

The reduction step from $x.a$ to $v(z)\{B = U\}$ caused the object to lose its name. Since the non-recursive type of the term depends on the name, the loss of the name also caused the term to lose its non-recursive type. This reduction step violates type preservation and type soundness.

18.2.3 Well-Typed Paths Don't Go Wrong

If pDOT programs can return paths without reducing them to values, could these paths be nonsensical? The type system ensures that they cannot. In particular, we ensure that if a path p has a type then p either identifies some value, and looking up p in the runtime configuration

$$\frac{\gamma(x) = v(x: T) \dots \{a = t\} \dots}{\gamma \mid x.a \mapsto \gamma \mid t} \text{(PROJ}_{\text{DOT}}\text{)}$$

see Chapter 21 for
pDOT safety proof

terminates, or p is a path that cyclically aliases other paths. Additionally, the pDOT safety proof ensures that if a path has a function or object type, then it can be looked up to a value; if p can *only* be typed with a singleton type (or \top), then the lookup will loop.

18.3 PATH REPLACEMENT

We introduce a *path replacement* operation for types that contain paths which reference the same object. If a path q is assigned to a path p then q *aliases* p . In the tree example from Chapter 16, t_3 aliases t_2 , but t_1 does not alias t_2 , even though they identify syntactically equal objects.

```
val t1 =
  new ConcreteTree
val t2 =
  new ConcreteTree
val t3 = t2
```

Section 19.2 defines
the replacement
operation precisely

If q is an alias of p we want to ensure that we can use q in the same way as p . For example, any term that has a type $T \rightarrow p.A$ should also have the type $T \rightarrow q.A$, and vice versa. In pDOT, we achieve this by introducing a subtyping relationship between *equivalent types*: if p and q are aliases, and a type T can be obtained from type U by *replacing* instances of p in U with q then T and U are equivalent. For example, $T \rightarrow q.A$ can be obtained from $T \rightarrow p.A$ by replacing p with q , and these types are therefore equivalent.

18.4 SINGLETON TYPES

To keep track of path aliases in the type system we use *singleton types*.

Suppose that a pDOT program assigns the path q to p , and that a type T can be obtained from U by replacing an instance of p with q . How does the type system know that T and U are equivalent? We could try passing information about the whole program throughout the type checker. However, that would make reasoning about types depend on reasoning about values, which would make typechecking more complicated and less modular, as shown in Part I.

Instead, we ensure that the type system keeps track of path aliasing using singleton types, an existing Scala feature. A singleton type of a path p , denoted $p.\text{type}$, is a type that is inhabited *only* with the value that is represented by p . In the tree example from Chapter 16, to tell the type system that t_3 aliases t_2 , we ensure that t_3 has the singleton type $t_2.\text{type}$. This information is used to allow subtyping between aliased paths, and to allow such paths to be typed with the same types.

Section 19.2
describes how
singleton types track
path identity

In pDOT, singleton types are an essential feature that is necessary to encode fully path-dependent types. However, this makes pDOT also the first DOT formalization of Scala's singleton types. In Chapter 20, we show a pDOT encoding of an example that motivates this Scala feature.

18.5 DISTINGUISHING FIELDS AND METHODS

Scala distinguishes between fields (*vals*, immutable fields that are strictly evaluated at the time of object initialization) and methods (*defs*, which are re-evaluated at each invocation). By contrast, DOT unifies the two in the concept of a term member. Since the distinction affects which paths are legal in Scala, we must make some similar distinction in pDOT. Consider the following Scala program:

```
val x = new {
  val a: { type A } = ta
  def b: { type B } = tb
}
val y: x.a.A
val z: x.b.B
```

Scala allows path-dependent types only on *stable* paths (Documentation, 2018b). A *val* can be a part of a stable path but a *def* cannot. Therefore, the type selection *x.a.A* is allowed but *x.b.B* is not.

DOT unifies the two concepts in one:

$$\text{let } x = v(x) \{a = t_a\} \wedge \{b = t_b\} \quad \text{in } \dots$$

However, this translation differs from Scala in the order of evaluation. Scala's fields, unlike DOT's, are fully evaluated to values when the object is constructed. Therefore, a more accurate translation of this example would be as follows:

$$\begin{array}{ll} \text{let } a' = t_a & \text{in} \\ \text{let } x = v(x) \{a = a'\} \wedge \{b = \lambda(_).t_b\} & \text{in } \dots \end{array}$$

This translation highlights the fact that although Scala can initialize *x.a* to an arbitrary term, that term will be already reduced to a value before evaluation reaches a context that contains *x*. The reason is that the constructor for *x* will strictly evaluate all of *x*'s *val* fields when *x* is created.

To model the fact that Scala field initializers are fully evaluated when the object is constructed, we require field initializers in pDOT to be values or paths, rather than arbitrary terms. We use the name *stable term* for a value or path.

This raises the question of how to model a Scala method such as *b*. A method can still be represented by making the delayed evaluation of the body explicit: instead of initializing the field *b* with the method body itself, we delay the body inside a lambda abstraction. The lambda abstraction, a value, can be assigned to the field *b*. The body of the lambda abstraction can be an arbitrary term; it is not evaluated during object construction, but later when the method is called and the lambda is applied to some dummy argument.

18.6 PRECISE SELF TYPES

DOT allows powerful type abstractions, but it demands *objects* as proof that the type abstractions make sense. An object assigns actual types to its type members and thus provides concrete evidence for the declared subtyping relationships between abstract type members. To make the connection between the object value and the type system, DOT requires the self type in an object literal to precisely describe the concrete types assigned to type members, and we need to define similar requirements for self types in pDOT.

In the object

$$\nu(x: \{A: T..T\}) \{A = T\}$$

DOT requires the self-type to be $\{A: T..T\}$ rather than some wider type $\{A: S..U\}$. This is not merely a convenience, but it is essential for soundness. Without the requirement, DOT could create and type the object

$$\nu(x: \{A: \top..\perp\}) \{A = T\}$$

which introduces the subtyping relationship $\top <: \perp$ and thus makes every type a subtype of every other type. Although we can require the actual assigned type T to respect the bounds (i.e. $\top <: T <: \perp$), such a condition is not sufficient to prohibit this object. The assigned type T and the bounds (\top and \perp in this example) can in general depend on the self variable, and thus the condition makes sense only in a typing context that contains the self variable with its declared self type. But in such a context, we already have the assumption that $\top <: x.A <: \perp$, so it holds that $\top <: T$ (since $\top <: x.A <: \perp <: T$) and similarly $T <: \perp$.

In pDOT, a path-dependent type $p.A$ can refer to type members not only at the top level, but also deep inside the object. Accordingly, we need to extend the precise self type requirement to apply recursively within the object, as follows:

1. An object containing a type member definition $\{A = T\}$ must declare A with tight bounds, using $\{A: T..T\}$ in its self type.
2. An object containing a definition $\{a = \nu(x: T)d\}$ must declare a with the recursive type $\mu(x: T)$, using $\{a: \mu(x: T)\}$ in its self type.
3. An object containing a definition $\{a = \lambda(x: T)U\}$ must declare a with a function type, using $\{a: \forall(x: S) V\}$ in its self type.
4. An object containing a definition $\{a = p\}$ must declare a with the singleton type $p.type$, using $\{a: p.type\}$ in its self type.

The first requirement is the same as in DOT. The second and third requirements are needed for soundness of paths that select type members from deep within an object. The fourth requirement is needed to

prevent unsoundness in the case of cyclic references. For example, if we were to allow the object

$$v(x: \{a: \{A: \top.. \perp\}\}) \{a = x.a\}$$

we would again have $\top <: \perp$. The fourth requirement forces this object to be declared with a precise self type:

$$v(x: \{a: x.a.type\}) \{a = x.a\}$$

Now, $x.a$ no longer has the type $\{A: \top.. \perp\}$, so it no longer collapses the subtyping relation. The precise typing thus ensures that cyclic paths can be only typed with singleton types but not function or object types, and therefore we cannot have type or term selection on cyclic paths.

Although both DOT and pDOT require precision in the self type of an object, the object itself can be typed with a wider type once it is assigned to a variable. For example, in DOT, if we have

$$\text{let } x = v(x: \{A: T..T\}) \{A = T\} \text{ in } \dots$$

then x also has the wider type $\{A: \perp.. \top\}$. Similarly, in pDOT, if we have

$$\text{let } x = v(x: \{a: \mu(y: \{b: \forall(z: T) U\}) \wedge \{c: x.a.b.type\}\})d \text{ in } \dots$$

then x also has all of the following types:

$$\begin{aligned} \Gamma &\vdash x: \{a: \{b: \forall(z: T) U\}\} \\ \Gamma &\vdash x: \{a: \mu(y: \{b: \top\})\} \\ \Gamma &\vdash x: \{c: x.a.b.type\} \\ \Gamma &\vdash x: \{c: \forall(z: T) U\} \end{aligned}$$

In fact, the typings for this object in pDOT are more expressive than in DOT. Because DOT does not open types nested inside of field declarations, DOT cannot assign the first two types to x . In Section 19.2, we show one simple type rule that generalizes pDOT to open and abstract types of term members nested deeply inside an object. In Chapter 20, we encode several examples from previous DOT papers in pDOT and show that the real-world compiler example from Chapter 16 that uses types depending on long paths can be encoded in pDOT as well.

In summary, both DOT and pDOT require the self type in an object literal to precisely describe the values in the literal, but this does not limit the ability to ascribe a more abstract type to the paths that identify the object.

The pDOT calculus generalizes DOT by allowing paths wherever DOT allows variables (except in places where variables are used as binders, such as x in $\lambda(x: T) t$).

19.1 SYNTAX

Figure 19.1 shows the abstract syntax of pDOT. The signature construct in pDOT is a path, defined to be a variable followed by zero or more field selections (e.g. $x.a.b.c$). pDOT uses paths wherever DOT uses variables. In particular, field selections $x.a$ and function application xy are done on paths: $p.a$ and pq . Most importantly, pDOT also generalizes DOT's types by allowing path-dependent types $p.A$ on paths rather than just on variables. Additionally, as described in Section 18.4, the pDOT calculus formalizes Scala's singleton types. A singleton type $p.\text{type}$ is inhabited with only one value: the value that is assigned to the path p . A singleton type thus indicates that a term designates the same object as the path p . Just as a path-dependent type $p.A$ depends on the value of p , a singleton type $q.\text{type}$ depends on the value of q . Singleton types are therefore a second form of dependent types in the calculus.

In pDOT, paths and values are considered *stable terms*. As motivated in Section 18.5, to distinguish between fields and methods, term members of object definitions can only be assigned stable terms, while still allowing the same (and even more expressive) type abstractions as in DOT. We use the meta-variable s to denote stable terms.

19.2 PDOT TYPING RULES

The typing and subtyping rules of pDOT are shown in Figures 19.2 and 19.3.

19.2.1 From Variables to Paths

The first thing to notice in the pDOT typing and subtyping rules is that all variable-specific rules, except VAR, are generalized to paths, as motivated in Section 18.1. The key rules that make DOT and pDOT interesting are the type-selection rules <-:SEL and SEL-<: . These rules enable us to make use of the type member in a path-dependent type. When a path p has type $\{A: S..U\}$, the rules introduce the path-dependent type $p.A$ into the subtyping relation by declaring the

```

 $p, q, r$    path
 $s$          stable term

 $p, q, r :=$ 
   $x$ 
   $p.a$ 
 $t, u :=$ 
   $s$ 
   $pq$ 
  let  $x = t$  in  $u$ 
 $s :=$ 
   $p$ 
   $v$ 
 $v :=$ 
   $v(x: T)d$ 
   $\lambda(x: T) t$ 
 $d :=$ 
   $\{a = s\}$ 
   $\{A = T\}$ 
   $d \wedge d'$ 
 $S, T, U :=$ 
   $\top$ 
   $\perp$ 
   $\{a: T\}$ 
   $\{A: S..T\}$ 
   $p.A$ 
   $p.\text{type}$ 
   $S \wedge T$ 
   $\mu(x: T)$ 
   $\forall(x: S) T$ 

```

Figure 19.1: Abstract syntax of pDOT (cf. DOT syntax in Figure 2.1)

Term typing
 $\Gamma \vdash t : T$

$$\begin{array}{c}
\frac{\Gamma(x) = T}{\Gamma \vdash x : T} \quad (\text{VAR}) \qquad \frac{\Gamma \vdash p : q.\text{type} \quad \Gamma \vdash q : T}{\Gamma \vdash p : T} \quad (\text{SNGL-TRANS}) \\
\\
\frac{\Gamma, x : T \vdash t : U \quad x \notin \text{fv}(T)}{\Gamma \vdash \lambda(x : T) t : \forall(x : T) U} \quad (\text{ALL-I}) \qquad \frac{\Gamma \vdash p : q.\text{type} \quad \Gamma \vdash q.a}{\Gamma \vdash p.a : q.a.\text{type}} \quad (\text{SNGL-E}) \\
\\
\frac{\Gamma \vdash p : \forall(z : S) T \quad \Gamma \vdash q : S}{\Gamma \vdash p q : T [q/z]} \quad (\text{ALL-E}) \\
\\
\frac{x; \Gamma, x : T \vdash d : T}{\Gamma \vdash \nu(x : T) d : \mu(x : T)} \quad (\{\}-I) \qquad \frac{\Gamma \vdash p : T [p/x]}{\Gamma \vdash p : \mu(x : T)} \quad (\text{REC-I}) \\
\\
\frac{\Gamma \vdash p : \{a : T\}}{\Gamma \vdash p.a : T} \quad (\text{FLD-E}) \qquad \frac{\Gamma \vdash p : \mu(x : T)}{\Gamma \vdash p : T [p/x]} \quad (\text{REC-E}) \\
\\
\frac{\Gamma \vdash p.a : T}{\Gamma \vdash p : \{a : T\}} \quad (\text{FLD-I}) \qquad \frac{\Gamma \vdash p : T \quad \Gamma \vdash p : U}{\Gamma \vdash p : T \wedge U} \quad (\&-I) \\
\\
\frac{\Gamma \vdash t : T \quad \Gamma, x : T \vdash u : U \quad x \notin \text{fv}(U)}{\Gamma \vdash \text{let } x = t \text{ in } u : U} \quad (\text{LET}) \qquad \frac{\Gamma \vdash t : T \quad \Gamma \vdash T <: U}{\Gamma \vdash t : U} \quad (\text{SUB})
\end{array}$$

Definition typing
 $p; \Gamma \vdash d : T$

$$\begin{array}{c}
p; \Gamma \vdash \{A = T\} : \{A : T..T\} \quad (\text{DEF-TYP}) \qquad \frac{\Gamma \vdash q}{p; \Gamma \vdash \{a = q\} : \{a : q.\text{type}\}} \quad (\text{DEF-PATH}) \\
\\
\frac{p; \Gamma \vdash \lambda(x : T) t : \forall(x : U) V}{p; \Gamma \vdash \{a = \lambda(x : T) t\} : \{a : \forall(x : U) V\}} \quad (\text{DEF-ALL}) \\
\\
\frac{p.a; \Gamma \vdash d [p.a/y] : T [p.a/y] \quad \text{tight } T}{p; \Gamma \vdash \{a = \nu(y : T) d\} : \{a : \mu(y : T)\}} \quad (\text{DEF-NEW}) \qquad \frac{p; \Gamma \vdash d_1 : T_1 \quad p; \Gamma \vdash d_2 : T_2 \quad \text{dom}(d_1), \text{dom}(d_2) \text{ disjoint}}{p; \Gamma \vdash d_1 \wedge d_2 : T_1 \wedge T_2} \quad (\text{ANDDEF-I})
\end{array}$$

Typeable paths
 $\Gamma \vdash p$

$$\frac{\Gamma \vdash p : T}{\Gamma \vdash p} \quad (\text{WF})$$

Tight bounds
 $\text{tight } T$

$$\text{tight } T = \begin{cases} U = V & \text{if } T = \{A : U..V\} \\ \text{tight } U & \text{if } T = \mu(x : U) \text{ or } T = \{a : U\} \\ \text{tight } U \text{ and tight } V & \text{if } T = U \wedge V \\ \text{true} & \text{otherwise} \end{cases}$$

Figure 19.2: pDOT typing rules (cf. DOT typing in Figure 2.3)

$$\begin{array}{c}
\begin{array}{c}
\Gamma \vdash T <: \top \quad (\text{TOP}) \\
\Gamma \vdash \perp <: T \quad (\text{BOT}) \\
\Gamma \vdash T <: T \quad (\text{REFL}) \\
\frac{\Gamma \vdash S <: T \quad \Gamma \vdash T <: U}{\Gamma \vdash S <: U} \quad (\text{TRANS}) \\
\Gamma \vdash T \wedge U <: T \quad (\text{AND}_1-<:) \\
\Gamma \vdash T \wedge U <: U \quad (\text{AND}_2-<:) \\
\frac{\Gamma \vdash S <: T \quad \Gamma \vdash S <: U}{\Gamma \vdash S <: T \wedge U} \quad (<:-\text{AND}) \\
\frac{\Gamma \vdash T <: U}{\Gamma \vdash \{a: T\} <: \{a: U\}} \quad (\text{FLD-<:-FLD})
\end{array}
&
\begin{array}{c}
\frac{\Gamma \vdash S_2 <: S_1 \quad \Gamma \vdash T_1 <: T_2}{\Gamma \vdash \{A: S_1..T_1\} <: \{A: S_2..T_2\}} \quad (\text{Typ-<:-Typ}) \\
\frac{\Gamma \vdash \bar{p}: \{A: S..T\}}{\Gamma \vdash S <: \bar{p}.A} \quad (<:-\text{SEL}) \\
\frac{\Gamma \vdash p: q.\text{type} \quad \Gamma \vdash q}{\Gamma \vdash T <: T[q/p]} \quad (\text{SNGL}_{pq}-<:) \\
\frac{\Gamma \vdash p: q.\text{type} \quad \Gamma \vdash q}{\Gamma \vdash T <: T[p/q]} \quad (\text{SNGL}_{qp}-<:) \\
\frac{\Gamma \vdash \bar{p}: \{A: S..T\}}{\Gamma \vdash \bar{p}.A <: T} \quad (\text{SEL-<:}) \\
\frac{\Gamma \vdash S_2 <: S_1 \quad \Gamma, x: S_2 \vdash T_1 <: T_2}{\Gamma \vdash \forall(x: S_1) T_1 <: \forall(x: S_2) T_2} \quad (\text{ALL-<:-ALL})
\end{array}
\end{array}$$

subtyping constraints $S <: p.A$ and $p.A <: U$. Thanks to these two rules, pDOT supports fully-path-dependent types.

19.2.2 Object Typing

Similarly to the DOT calculus, the $\{\}$ -I rule gives an object $v(x: T)d$ with declared type T which may depend on the self variable x the recursive type $\mu(x: T)$. The rule also checks that the definitions d of the object actually do have type T under the assumption that the self variable has this type. The object's definitions d are checked by the Definition typing rules. As discussed in Section 18.6, the rules assign a precise self type for objects, ensuring that paths are declared with singleton types, functions with function types, and objects with object types. For objects, the tight T condition ensures that all type members that can be reached by traversing T 's fields have equal bounds, while still allowing arbitrary bounds in function types.⁵

A difference with DOT is that pDOT's definition-typing judgment keeps track of the path that represents an object's identity. When we typecheck an outermost object that is not nested in any other object, we use the $\{\}$ -I rule. The rule introduces x as the identity for the object and registers this fact in the definition-typing judgment. To typecheck an object that is assigned to a field a of another object p we use the DEF-NEW rule. This rule typechecks the object's body assuming the object identity $p.a$ and replaces the self-variable of the object with that path. The definition-typing judgment keeps track of the path to the

Figure 19.3: pDOT subtyping rules (cf. DOT subtyping in Figure 2.3)

⁵ see Section 21.1 for details

definition's enclosing object starting from the root of the program. This way the type system knows what identities to assign to nested objects. For example, when typechecking the object assigned to $x.a$ in the expression

let $x = v(x) \{a = v(y) \{b = y.b\}\}$ **in** ...

we need to replace y with the path $x.a$:

$$\frac{\frac{\Gamma, x: \{a: \mu(y: \{b: y.b.type\})\} \vdash x.a}{x.a; \Gamma, x: \{a: \mu(y: \{b: y.b.type\})\} \vdash \{b = x.a.b\} : \{b: x.a.b.type\}} \text{DEF-PATH} \quad \text{tight } \{b: y.b.type\}}{x; \Gamma, x: \{a: \mu(y: \{b: y.b.type\})\} \vdash \{a = v(y) \{b = y.b\}\} : \{a: \mu(y: \{b: y.b.type\})\}} \text{DEF-NEW} \quad \{\}-I \quad \Gamma \vdash v(x) \{a = v(y) \{b = y.b\}\} : \mu(x: \{a: \mu(y: \{b: y.b.type\})\})$$

An alternative design of the DEF-NEW rule can be to introduce a fresh variable y into the context (similarly to the $\{\}-I$ rule). However, we would have to assign y the type $x.a.type$ to register the fact that these two paths identify the same object. We decided to simplify the rule by immediately replacing the nested object's self variable with the outer path to avoid the indirection of an additional singleton type.

19.2.3 Path Alias Typing

In pDOT, singleton-type related typing and subtyping rules are responsible for the handling of aliased paths and equivalent types.

SINGLETON TYPE CREATION How does a path p obtain a singleton type? A singleton type indicates that in the initial program, a prefix of p (which could be all of p) is assigned a path q . For example, in the program

let $x = v(x: \{a: x.type\} \wedge \{b: S\}) \{a = x\} \wedge \{b = s\}$ **in** ...

the path $x.a$ should have the type $x.type$ because $x.a$ is assigned the path x . The singleton type for $x.a$ can be obtained as follows. Suppose that in the typing context of the **let** body, x is mapped to the type of its object, $\mu(x: \{a: x.type\} \wedge \{b: S\})$. Through applying recursion elimination (REC-E), field selection (FLD-E), and finally subsumption (SUB) with the intersection subtyping rule $\text{AND}_1-<:$, we will obtain that $\Gamma \vdash x.a: x.type$.

In the above example, $x.a$ aliases x , so anything that we can do with x we should be able to do with $x.a$. Since x has a field b and we can create a path $x.b$, we want to also be able to create a path $x.a.b$. Moreover, we want to treat $x.a.b$ as an alias for $x.b$. This is done through the SNGL-E rule: it says that if p aliases q , and $q.a$ has a type (denoted with $\Gamma \vdash q.a$), then $p.a$ aliases $q.a$. This rule allows us to conclude that $\Gamma \vdash x.a.b: x.b.type$.

$$\frac{\Gamma \vdash q.a}{\Gamma \vdash p: q.type} \quad \Gamma \vdash p.a: q.a.type \quad (\text{SNGL-E})$$

SINGLETON TYPE PROPAGATION In the above example we established that the path $x.a.b$ is an alias for $x.b$. Therefore, we want to be able to type $x.a.b$ with any type with which we can type $x.b$. The SNGL-TRANS rule allows us to do just that: if p is an alias for q , then we can type p with any type with which we can type q . Using that rule, we can establish that $\Gamma \vdash x.a.b : S$ because $\Gamma \vdash x.b : S$.

$$\frac{\Gamma \vdash q : T \quad \Gamma \vdash p : q.\text{type}}{\Gamma \vdash p : T} \text{ (SNGL-TRANS)}$$

EQUIVALENT TYPES As described in Section 18.3, we call two types equivalent if they are equal up to path aliases. We need to ensure that equivalent types are equivalent by subtyping, i.e. that they are subtypes of each other. For example, suppose that $\Gamma \vdash p : q.\text{type}$, and the path r refers to an object $v(x) \{a = p\} \wedge \{b = p\}$. Then we want to be able to type r with all of the following types:

$$\begin{array}{ll} \Gamma \vdash r : \{a : p.\text{type}\} \wedge \{b : p.\text{type}\} & \Gamma \vdash r : \{a : p.\text{type}\} \wedge \{b : q.\text{type}\} \\ \Gamma \vdash r : \{a : q.\text{type}\} \wedge \{b : q.\text{type}\} & \Gamma \vdash r : \{a : q.\text{type}\} \wedge \{b : p.\text{type}\} \end{array}$$

The pDOT subtyping rules $\text{SNGL}_{pq} < :$ and $\text{SNGL}_{qp} < :$ allow us to assign these types to r by establishing subtyping between equivalent types. Specifically, if we know that $\Gamma \vdash p : q.\text{type}$ then the rules allow us to replace any occurrence of p in a type T with q , and vice versa, while maintaining subtyping relationships in both directions.

$$\frac{\Gamma \vdash p : q.\text{type} \quad \Gamma \vdash q}{\Gamma \vdash T < : T[q/p]} \text{ (SNGL}_{pq} < :)$$

$$\frac{\Gamma \vdash p : q.\text{type} \quad \Gamma \vdash q}{\Gamma \vdash T < : T[p/q]} \text{ (SNGL}_{qp} < :)$$

We express that two types are equivalent using the *replacement operation*. The operation is similar to the substitution operation, except that we replace paths with paths instead of variables with terms, and we replace only one path at a time rather than all of its occurrences. The statement $T[q/p] = U$ denotes that the type T contains one or more paths that start with p , e.g. $p.\overline{b_1}, \dots, p.\overline{b_n}$, and that exactly one of these occurrences $p.\overline{b_i}$ is replaced with $q.\overline{b_i}$, yielding the type U . Note that it is not specified exactly in which occurrence of the above paths the prefix p is replaced with q . The precise definition of the replacement operation is presented in Figure 19.4.

Given the path r from the above example, we can choose whether to replace the first or second occurrence of p with q ; for example, we can derive

$$\frac{\frac{\dots}{\Gamma \vdash r : \{a : p.\text{type}\} \wedge \{b : p.\text{type}\}} \text{ REC-E} \quad \frac{\Gamma \vdash p : q.\text{type} \quad \frac{\dots}{\{a : p.\text{type}\} \wedge \{b : p.\text{type}\} [q/p] = \{a : p.\text{type}\} \wedge \{b : q.\text{type}\}} \text{ REPL-AND}_2}{\Gamma \vdash \{a : p.\text{type}\} \wedge \{b : p.\text{type}\} < : \{a : p.\text{type}\} \wedge \{b : q.\text{type}\}} \text{ SNGL}_{pq} < : \quad \text{SUB}$$

To replace several occurrences of a path with another, we repeatedly apply $\text{SNGL}_{pq} < :$ or $\text{SNGL}_{qp} < :$.

$$\begin{array}{c}
\frac{}{p.\bar{b}.A [q/p] = q.\bar{b}.A} \text{ (REPL-PATH)} \\
\frac{}{p.\bar{b}.\text{type} [q/p] = q.\bar{b}.\text{type}} \text{ (REPL-SNGL)} \\
\frac{T_1 [q/p] = T_2}{(T_1 \wedge U) [q/p] = T_2 \wedge U} \text{ (REPL-AND}_1\text{)} \\
\frac{T_1 [q/p] = T_2}{(U \wedge T_1) [q/p] = U \wedge T_2} \text{ (REPL-AND}_2\text{)} \\
\frac{T [q/p] = U}{\mu(x: T) [q/p] = \mu(x: U)} \text{ (REPL-REC)} \\
\frac{T_1 [q/p] = T_2}{(\forall(x: T_1) U) [q/p] = \forall(x: T_2) U} \text{ (REPL-ALL}_1\text{)} \\
\frac{T_1 [q/p] = T_2}{(\forall(x: U) T_1) [q/p] = \forall(x: U) T_2} \text{ (REPL-ALL}_2\text{)} \\
\frac{T_1 [q/p] = T_2}{\{a: T_1\} [q/p] = \{a: T_2\}} \text{ (REPL-FLD)} \\
\frac{T_1 [q/p] = T_2}{\{A: T_1..U\} [q/p] = \{A: T_2..U\}} \text{ (REPL-TYP}_1\text{)} \\
\frac{T_1 [q/p] = T_2}{\{A: U..T_1\} [q/p] = \{A: U..T_2\}} \text{ (REPL-TYP}_2\text{)}
\end{array}$$

Figure 19.4: Replacement of a path p in a type by q

19.2.4 Abstracting Over Field Types

Finally, we describe one of the most interesting pDOT rules which adds significant expressivity to pDOT.

Consider a function

$$f = \lambda(x: \{a: T\}) \dots$$

and a path p that refers to the object

$$\nu(x: \{a: q.\text{type}\}) \{a = q\}$$

where $\Gamma \vdash q: T$. Since

$$\Gamma \vdash p: \mu(x: \{a: q.\text{type}\})$$

by REC-E, we have

$$\Gamma \vdash p: \{a: q.\text{type}\}$$

we assume for simplicity that q does not start with x

Therefore, since $\Gamma \vdash q: T$, we would like to be able to pass p into the function f which expects an argument of type $\{a: T\}$. Unfortunately, the typing rules so far do not allow us to do that because although q has type T , $q.\text{type}$ is *not* a subtype of T , and therefore $\{a: q.\text{type}\}$ is not a subtype of $\{a: T\}$.

The type rule FLD-I allows us to bypass that limitation. If a path p has a record type $\{a: T\}$ (and therefore $\Gamma \vdash p.a: T$), then the rule lets

$$\frac{\Gamma \vdash p.a: T}{\Gamma \vdash p: \{a: T\}} \text{ (FLD-I)}$$

us type p with any type $\{a: U\}$ as long as $p.a$ can be typed with U . More generally, using this typing rule, we can derive

$$\frac{\Gamma \vdash p.a_1.a_2.\dots.a_n: U}{\Gamma \vdash p: \{a_1: \{a_2: \dots \{a_n: U\}\}\}}$$

For the above example, we can prove that $\Gamma \vdash p: \{a: T\}$ and pass it into f as follows:

$$\frac{\frac{\frac{\Gamma \vdash p: \{a: q.\text{type}\}}{\Gamma \vdash p.a: q.\text{type}} \text{FLD-E} \quad \Gamma \vdash q: T}{\Gamma \vdash p.a: T} \text{SNGL-TRANS}}{\Gamma \vdash p: \{a: T\}} \text{FLD-I}$$

The FLD-I rule allows us to eliminate recursion on types that are nested inside fields, which is not possible in DOT. If a DOT function f expects a parameter of type $\{a: \mu(x: T)\}$, then in DOT, we cannot pass a variable y of type $\{a: \mu(x: T \wedge U)\}$ or a variable z of type $\{a: T[z.a/x]\}$ into f because there is no subtyping between recursive types, and there is no subtyping relationship between $\mu(x: T)$ and $T[z.a/x]$ (and the latter type might not exist in the first place due to the lack of fully-path-dependent types). All of the above is possible in pDOT because both $y.a$ and $z.a$ can be typed with $\mu(x: T)$, which allows us to use the FLD-I rule and type y and z as $\{a: \mu(x: T)\}$.

19.3 REDUCTION SEMANTICS

The operational semantics of pDOT is presented in Figure 19.5. pDOT's reduction rules mirror the DOT rules with three distinctions:

- *paths everywhere*: wherever DOT uses (as opposed to defines) variables, pDOT uses paths;
- *no PROJ_{DOT}*: there is no reduction rule for field projection because in pDOT, paths are normal form (as motivated in Section 18.2.2);
- *path lookup*: pDOT uses the reflexive, transitive closure of the *path lookup* operation \rightsquigarrow that generalizes variable lookup in stores to paths.

The path lookup operation is presented in Figure 19.6. This operation allows us to look up a value that is nested deeply inside an object. If a path is a variable the lookup operation is a straightforward variable lookup (LOOKUP-STEP-VAR). If in a store γ , a path p is assigned an object $v(x) \{a = s\}$ then $\gamma \vdash p.a \rightsquigarrow s[p/x]$ because the self variable x in s gets replaced with p (LOOKUP-STEP-VAL). If p is equal to another path q then $\gamma \vdash p.a \rightsquigarrow q.a$ (LOOKUP-STEP-PATH).

Finally, we want to be able to follow a sequence of paths in a store: for example, if $\gamma \vdash p \rightsquigarrow q$ and $\gamma \vdash q \rightsquigarrow v$, we want to conclude that

$$\begin{array}{c}
\frac{\gamma \vdash p \rightsquigarrow^* \lambda(z:T) t}{\gamma \mid \underline{pq} \mapsto \gamma \mid t[q/z]} \quad (\text{APPLY}) \\
\\
\gamma \mid \text{let } x = \underline{p} \text{ in } t \mapsto \gamma \mid t[p/x] \quad (\text{LET-PATH}) \\
\\
\gamma \mid \text{let } x = v \text{ in } t \mapsto \gamma, x \mapsto v \mid t \quad (\text{LET-VALUE}) \\
\\
\frac{\gamma \mid t \mapsto \gamma' \mid t'}{\gamma \mid \text{let } x = t \text{ in } u \mapsto \gamma' \mid \text{let } x = t' \text{ in } u} \quad (\text{CTX})
\end{array}$$

Figure 19.5: Operational semantics of pDOT

$$\begin{array}{c}
\frac{\gamma(x) = v}{\gamma \vdash x \rightsquigarrow v} \quad (\text{LOOKUP-STEP-VAR}) \\
\\
\frac{\gamma \vdash p \rightsquigarrow v(x:T) \dots \{a=s\} \dots}{\gamma \vdash p.a \rightsquigarrow s[p/x]} \quad (\text{LOOKUP-STEP-VAL}) \\
\\
\frac{\gamma \vdash p \rightsquigarrow q}{\gamma \vdash p.a \rightsquigarrow q.a} \quad (\text{LOOKUP-STEP-PATH}) \\
\\
\gamma \vdash s \rightsquigarrow^* s \quad (\text{LOOKUP-REFL}) \\
\\
\frac{\gamma \vdash s_1 \rightsquigarrow s_2 \quad \gamma \vdash s_2 \rightsquigarrow^* s_3}{\gamma \vdash s_1 \rightsquigarrow^* s_3} \quad (\text{LOOKUP-TRANS})
\end{array}$$

Figure 19.6: Value-environment path lookup

looking up p yields v . This is done through the reflexive, transitive closure \rightsquigarrow^* of the \rightsquigarrow relation (LOOKUP-REFL and LOOKUP-TRANS).

For example, looking up $x.a.c$ in the environment

$$\gamma = (y, v(y')\{b = v(y'')\{c = \lambda(z: \top) z\}\}), \\ (x, v(x)\{a = y.b\})$$

yields $\lambda(z: \top) z$:

$$\begin{array}{ll} \gamma(x) = & v(x)\{a = y.b\} & \gamma(y) = & v(y')\{b = v(y'')\{c = \lambda(z: \top) z\}\} \\ \gamma \vdash x \rightsquigarrow & v(x)\{a = y.b\} & \gamma \vdash y \rightsquigarrow & v(y')\{b = v(y'')\{c = \lambda(z: \top) z\}\} \\ \gamma \vdash x. & a \rightsquigarrow y.b & \gamma \vdash y. & b \rightsquigarrow v(y'')\{c = \lambda(z: \top) z\} \\ \gamma \vdash x. & a.c \rightsquigarrow y.b.c & \gamma \vdash y. & b. \quad c \rightsquigarrow \lambda(z: \top) z \end{array}$$

$$\gamma \vdash x.a.c \rightsquigarrow^* \lambda(z: \top) z$$

The reduction rule that uses the lookup operation is the function application rule APPLY: to apply p to q we must be able to look up p in the store and obtain a function. Since pDOT permits cycles in paths, does this mean that the lookup operation for this type rule might not terminate? Fortunately, pDOT's type safety ensures that this will not happen. As shown in Section 21.4, if $\Gamma \vdash p: \forall(T: U)$ then lookup of p eventually terminates and results in a function value. Therefore, a function application $p q$ always makes progress.

EXAMPLES

In this chapter, we present three pDOT program examples that illustrate different features of the calculus. All of the programs were formalized and typechecked in Coq.

To make the examples easier to read, we simplify the notation for objects $v(x: U)d$ by removing type annotations where they can be easily inferred, yielding a new notation $v(x \Rightarrow d')$ where d' are the definitions d modified as follows:

- a type definition $\{A = T\}$ can be only typed with $\{A: T..T\}$, so we will skip type declarations;
- in a definition $\{a = p\}$, the field a is assigned a path and can be only typed with a singleton type; we will therefore skip the type $\{a: p.type\}$;
- in a definition $\{a = v(x: T)d\}$, a is assigned an object that must be typed with $\mu(x: T)$; since we can infer T by looking at the object definition, we will skip the typing $\{a: \mu(x: T)\}$;
- we inline the type of abstractions into the field definition (e.g. $\{a: \forall(T: U) = \lambda(x: T) t\}$).

For readability we will also remove the curly braces around object definitions and replace the \wedge delimiters with semicolons. As an example of our abbreviations, the object

$$v(x: \{A: T..T\} \wedge \{a: p.type\} \wedge \{b: \mu(y: U)\} \wedge \{c: \forall(z: S) V\} \\ \{A = T\} \wedge \{a = p\} \wedge \{b = v(y: U)d\} \wedge \{c = \lambda(z: S') t\})$$

will be encoded as

$$v(x \Rightarrow A = T; a = p; b = v(y \Rightarrow d'); c: \forall(z: S) V = \lambda(z: S') t)$$

Note that translating the nested object $v(y: U)d$ yielded a new notation $v(y \Rightarrow d')$ where d' stands for the new encoding of d .

20.1 CLASS ENCODINGS

Fully path-dependent types allow pDOT to define encodings for Scala's module system and classes, as we will see in the examples below.

In Scala, declaring a class $A(args)$ automatically defines both a type A for the class and a constructor for A with parameters $args$. We will

encode such a Scala class in pDOT as a type member A and a method newA that returns an object of type A :

$$\begin{array}{ll} \nu(p \Rightarrow & \text{package } p \{ \\ & A = \mu(\mathbf{this}: \{\text{foo}: \forall(_) U\}); & \text{class } A(x: U) \{ \\ & \text{newA}: \forall(x: U) p.A & \quad \text{def } \text{foo}: U = x \\ & = \nu(\mathbf{this}) \{\text{foo} = \lambda(_).x\} & \} \\ & & \} \end{array}$$

To encode subtyping we use type intersection. For example, we can define a class B that extends A as follows:

$$\begin{array}{ll} \nu(p \Rightarrow & \text{package } p \{ \\ & B = p.A \wedge \{C: \perp.. \top\}; & \text{class } B(x:U) \text{ extends } A(x:U) \{ \\ & \text{newB}: \forall(x: U) p.B & \\ & = \nu(\mathbf{this}) \{\text{foo} = \lambda(_).x; & \quad \text{type } C \\ & \quad C = \dots\} & \} \\ & & \} \end{array}$$

20.2 LISTS

As an example to illustrate that pDOT supports the type abstractions of DOT we formalize the covariant-list library by Amin, Grütter, et al., (2016) in pDOT, presented in Figure 20.1. The encoding defines `List` as a data type with an element type A and methods `head` and `tail`. The library contains `nil` and `cons` fields for creating lists. To soundly formalize the list example, we encode `head` and `tail` as *methods* (defs) as opposed to *vals* by wrapping them in lambda abstractions, as discussed in Section 18.5. This encoding also corresponds to the Scala standard library where `head` and `tail` are *defs* and not *vals*, and hence one cannot perform a type selection on them.

By contrast, the list example by Amin, Grütter, et al., (2016) encodes `head` and `tail` as fields without wrapping their results in functions. For a DOT that supports paths, such an encoding is unsound because it violates the property that paths to objects with type members are acyclic. In particular, since no methods should be invoked on `nil`, its `head` and `tail` methods are defined as non-terminating loops, and `nil`'s element type is instantiated to \perp . If we allowed `nil.head` to have type \perp then since $\perp <: \{A: \top.. \perp\}$, we could derive $\top <: \text{nil.head.A} <: \perp$.

```

ν(sci ⇒ List = μ(self: {A: ⊥..⊤} ∧
    {head: ∀(⊥) self.A} ∧
    {tail: ∀(⊥) (sci.List ∧ {A: ⊥..self.A})});
nil: ∀(x: {A: ⊥..⊤}) sci.List ∧ {A: ⊥..⊥}
    = λ(x: {A: ⊥..⊤})
        let result = ν(self ⇒ A = ⊥;
            head: ∀(y: ⊤) self.A = λ(y: ⊤) self.head y;
            tail: ∀(y: ⊤) (sci.List ∧ self.A) = λ(y: ⊤) self.tail y)
        in result;
cons: ∀(x: {A: ⊥..⊤}) ∀(hd: x.A) ∀(tl: sci.List ∧ {A: ⊥..x.A}) sci.List ∧ {A: ⊥..x.A}
    = λ(x: {A: ⊥..⊤}) λ(hd: x.A) λ(tl: sci.List ∧ {A: ⊥..x.A})
        let result = ν(self ⇒ A = x.A;
            head: ∀(⊥) self.A = λ⊥. hd
            tail: ∀(⊥) (sci.List ∧ self.A) = λ⊥. tl)
        in result)

```

Figure 20.1: A **co-variant list** library in pDOT

20.3 MUTUALLY RECURSIVE MODULES

The second example, presented in Figure 20.2, illustrates pDOT's ability to use path-dependent types of arbitrary length. It formalizes the compiler example from Chapter 16 in which the nested classes `Type` and `Symbol` recursively reference each other.

```

ν(dc ⇒ types = ν(types ⇒ Type = μ(this: {symb: dc.symbols.Symbol});
    newType: ∀(s: dc.symbols.Symbol) types.Type
    = λ(s: dc.symbols.Symbol)
        let result' = ν(this ⇒ symb = s) in result');
symbols = ν(symbols ⇒ Symbol = μ(this: {tpe: dc.types.Type});
    newSymbol: ∀(t: dc.types.Type) symbols.Symbol
    = λ(t: dc.types.Type)
        let result' = ν(this ⇒ tpe = t) in result'))

```

Figure 20.2: Mutually recursive types in a compiler package: [fully-path-dependent types](#)

20.4 CHAINING METHODS WITH SINGLETON TYPES

The last example focuses on pDOT's ability to use singleton types as they are motivated by Odersky and Zenger, (2005b). An example from that paper introduces a class C with an `incr` method that increments a mutable integer field x and returns the object itself (`this`). A class D extends C and defines an analogous `decr` method. The example shows how we can invoke a chain of `incr` and `decr` methods on an object of type D using singleton types: if $C.\text{incr}$ returned an object of type C this would be impossible since C does not have a `decr` member, so the method's return type is `this.type`, a singleton type.

Our formalization of the example is displayed in Figure 20.3. Since pDOT does not support mutation, our example excludes the mutation side effect of the original example which is there to make the example more practical.

```

let pkg =  $\nu(p \Rightarrow C = \mu(\text{this: } \{\text{incr: this.type}\});$ 
            $D = \mu(\text{this: } p.C \wedge \{\text{decr: this.type}\});$ 
           newD:  $\forall(\_) p.D = \lambda \_.$ 
               let result =  $\nu(\text{this} \Rightarrow \text{incr} = \text{this}; \text{decr} = \text{this})$ 
               in result)
in let d = pkg.newD _
in d.incr.decr

```

Figure 20.3: Chaining method calls using singleton types

TYPE SAFETY

We implemented the type-safety proof of pDOT in Coq as an extension of the simple DOT soundness proof by Rapoport, Kabir, et al., (2017). Compared to the 2,051 LOC, 124 lemmas and theorems, and 65 inductive or function definitions in the simple DOT proof, the pDOT Coq formalization consists of 7,343 LOC, 429 lemmas and theorems, and 115 inductive or function definitions. This section presents an overview of the key challenges and insights of proving pDOT sound.

The challenges of adapting the DOT soundness proof to pDOT can be classified into three main themes:

- adapting the notion of inert types to pDOT,
- adapting the stratification of typing rules to pDOT, and
- adapting the canonical forms lemma to changes in the operational semantics in pDOT.

21.1 INERT TYPES IN PDOT

The purpose of inertness is to prevent the introduction of a possibly undesirable subtyping relationship between arbitrary types $S <: U$ arising from the existence of a type member that has those types as bounds. If a variable x has type $\{A: S..U\}$, then $S <: x.A$ and $x.A <: U$, so by transitivity, $S <: U$.

As presented in Section 5.2, a DOT type is *inert* if it is a function type or a recursive type $\mu(x: T)$ where T is a *record type*. A record type is either a type-member declaration with equal bounds $\{A: U..U\}$ or an arbitrary field declaration $\{a: S\}$. In DOT, this is sufficient to rule out the introduction of new subtyping relationships.

In pDOT, a new subtyping relationship $S <: U$ arises when a *path* p , rather than only a variable x , has a type member $\{A: S..U\}$. Therefore, the inertness condition needs to enforce equal bounds on type members not only at the top level of an object, but recursively in any objects nested deeply within the outermost object. Therefore, a field declaration $\{a: T\}$ is inert only if the field type T is also inert. Moreover, since pDOT adds singleton types to DOT, the definition of a record type is also extended to allow a field to have a singleton type.

DEFINITION 34 (Record Types in pDOT). *A record type is an intersection of types each of which is either a field declaration $\{a: T\}$ where T is inert or a singleton type, or a tight type declaration $\{A: U..U\}$.*

$$\begin{array}{c}
 \frac{\text{inert } T}{\text{record } \{a: T\}} \\
 \text{record } \{a: p.\text{type}\} \\
 \text{record } \{A: U..U\} \\
 \frac{\text{record } T \quad \text{record } U}{\text{record } T \wedge U} \\
 \text{inert } \forall(x: T) U \\
 \frac{\text{record } T}{\text{inert } \mu(x: T)} \\
 \text{(cf. Definition 2)}
 \end{array}$$

DEFINITION 35 (Inert Types in pDOT). A type U is inert if it is either a function type or a recursive type $\mu(x: T)$ where T is a record type.

Both the DOT and pDOT preservation lemmas must ensure that reduction preserves inertness of typing contexts.

ASIDE: WHY WE NEED THE `tight` JUDGMENT This section motivates the need for the `tight` judgment that occurs in the **DEF-NEW** definition-typing rule (Figure 19.2). The reader may skip this section as it explains the design of the type rules and not the pDOT proof.

In designing pDOT and its safety proof, we must ensure that reduction preserves inertness: that is, when $\gamma \mid t \mapsto \gamma' \mid t'$ there is an inert typing environment Γ' that is well-formed with respect to γ' and in which t' has the required type (i.e. we need to be able to prove the analogue of Lemma 20 (Value Typing)). In order to guarantee this, we need to ensure that *all well-typed values can be typed with an inert type*. As we will show now, if we omit the `tight` judgment from the object-definition typing rule **DEF-NEW**, there will be well-typed values that do not have a *precise* inert type. The values might still have an inert type under general typing but relying on such reasoning would further complicate the soundness proof.

$$\frac{\text{tight } T \quad p.a; \Gamma \vdash d[p.a/y]: T[p.a/y]}{p; \Gamma \vdash \{a = v(y: T)d\} : \{a: \mu(y: T)\}} \text{ (DEF-NEW)}$$

The object-definition typing rule **DEF-NEW** requires that the type T of an object $v(x: T)d$ be `tight`, i.e. that all type declarations of the type (except those nested inside function types) have equal bounds. Together with the pDOT definition-typing rules, this ensures that T has the form of a concatenation of record types, and, due to `tight`, that each type declaration $\{A: T..U\}$ of that concatenation has equal bounds ($T = U$).

Given that every record that declares a type member must be typed with the **DEF-TYP** rule, which requires `tight` bounds anyway, why do we need this additional tightness restriction?

If we do not require the tightness restriction in object-definition typing then there will exist well-typed objects $v(x: T)d$ whose precise recursive types $\mu(x: T)$ are not inert, which would complicate the soundness proof. Consider the following value:

$$\begin{aligned} &v(x: \{a: \mu(y: \{A: y.\text{type}..x.a.\text{type}\})\}) \\ &\quad \{a = v(y: \{A: y.\text{type}..x.a.\text{type}\})\} \\ &\quad \{A = y.\text{type}\} \end{aligned}$$

It cannot be typed with an inert type using the **{-I}** rule. Is this value well-typed? If we remove the tightness condition, then it is:

$$\frac{\frac{x.a; \Gamma \vdash \{A = x.a.\text{type}\} : \{A: x.a.\text{type}..x.a.\text{type}\}}{\text{DEF-TYP}} \quad \frac{x; \Gamma \vdash \{a = v(y: T_{Ay}^{x.a}) \{A = y.\text{type}\}\} : \{a: \mu(y: T_{Ay}^{x.a})\}}{\text{DEF-NEW}}}{\vdash v(x) \{a = v(y: T_{Ay}^{x.a}) \{A = y.\text{type}\}\} : \mu(x: \{a: \mu(y: T_{Ay}^{x.a})\})} \text{ {-I}}$$

where we use the shorthand

$$T_{Ay}^{x.a} = \{A: y.\text{type}..x.a.\text{type}\}$$

The reason that this works is that the substitution of the self-variable y with $x.a$ in DEF-NEW made the type $T_{Ay}^{x.a}$ have equal bounds, yielding $\{A: x.a.\text{type}..x.a.\text{type}\}$, and the DEF-TYP typing went through.

To avoid this problem we must either prevent the above value from being well-typed or change the definition of inertness to include the above type. Since the latter would significantly complicate the inertness definition which is supposed to be a simple syntactic property, we chose the former: a check in the definition typing rules that object types have tight bounds.

21.2 PROOF RECIPE FOR PDOT

The DOT proof presented in Part I employs the *proof recipe*, a stratification of the typing rules into multiple typing relations that rule out cycles in a typing derivation, but are provably as expressive as the general typing relation under the condition of an inert typing context. Recall that besides the general typing relation, the proof uses three intermediate relations:

- *tight* typing neutralizes the $<:-\text{SEL}$ and $\text{SEL}<:-$ rules that could introduce bad bounds,
- *invertible* typing contains introduction rules that create more complex types out of smaller ones, and
- *precise* typing contains elimination rules that decompose a type into its constituents.

see Section 5.3 for
tight typing in DOT

see Section 5.4 for
invertible typing in
DOT

see Figure 5.2 for
precise typing in
DOT

The language features that pDOT adds to DOT create new ways to introduce cycles in a typing derivation. The stratification of the typing rules needs to be extended to eliminate these new kinds of cycles.

21.2.1 Overview of Extended Proof Recipe

The notion of aliased paths is inherently symmetric: if p and q are aliases for the same object, then any term with type $p.A$ also has type $q.A$ and vice versa. This is complicated further because the paths p and q can occur deeply inside some complex type, and whether a term has such a type should be independent of whether the type is expressed in terms of p or q . Another complicating factor is that a prefix of a path is also a path, which may itself be aliased. For example, if p is an alias of q and $q.a$ is an alias of r , then by transitivity, $p.a$ should also be an alias of r .

The pDOT proof eliminates cycles due to aliased paths by breaking the symmetry of path aliasing. When p and q are aliased, either

$\Gamma \vdash p : q.\text{type}$ or $\Gamma \vdash q : p.\text{type}$. The typing rules carefully distinguish these two cases, so that for every pair p, q of aliased paths introduced by a typing declaration, we know whether the aliasing was introduced by the declaration of p or of q . A key lemma then proves that if we have any sequence of aliasing relationships

$$p_0 \sim p_1 \sim \dots \sim p_n$$

where for each i , either

$$\Gamma \vdash p_i : p_{i+1}.\text{type} \quad \text{or} \quad \Gamma \vdash p_{i+1} : p_i.\text{type}$$

we can reorder the replacements so that the ones of the first type all come first and the ones of the second type all come afterwards. More precisely, we can always find some “middle” path q such that

$$\Gamma \vdash p_0 : q.\text{type} \quad \text{and} \quad \Gamma \vdash p_n : q.\text{type}$$

(in degenerate cases, the middle path q might actually be p_0 or p_n). Therefore, we further stratify the proof recipe into two typing judgments the first of which accounts for the $\text{SNGL}_{pq}^-<$ rule, and the second for the $\text{SNGL}_{qp}^-<$ rule. This eliminates cycles in a typing derivation due to aliased paths, but the replacement reordering lemma ensures that it preserves expressiveness.

Another new kind of cycle is introduced by the field elimination rule FLD-E and the field introduction rule FLD-I that is newly added in pDOT . This cycle can be resolved in the same way as other cycles in DOT , by stratifying these rules in two different typing relations.

The final stratification of the pDOT typing rules requires seven typing relations rather than the four required in the soundness proof for DOT . General and tight typing serve the same purpose as in the DOT proof, but pDOT requires three elimination and two introduction typing relations.

Table 21.1 shows which typing rules of pDOT are handled by each of the auxiliary typing relations which are introduced in the next sections.

Relation	Type rules	Inlined subtyping rules ($\text{SUB} + \dots$)
Precise	$\vdash!$ VAR, FLD-E , REC-E	$\text{AND}_1^-<$, $\text{AND}_2^-<$
	$\vdash!!$ SNGL-E	
	$\vdash!!!$ SNGL-TRANS	
Invertible	\vdash_i $\&\text{-I}$	$\text{SNGL}_{pq}^-<$, $<:-\text{AND}$, TOP , $\text{ALL}^-<$, ALL , $\text{FLD}^-<$, FLD , $\text{TYP}^-<$, TYP
	\vdash_{ii} $\&\text{-I}$, FLD-I , REC-I	$\text{SNGL}_{qp}^-<$, $<:-\text{AND}$, $<:-\text{SEL}$
Tight	$\vdash_\#$ all	all, tight versions of SEL and SNGL rules
General	\vdash all	all

Table 21.1: Auxiliary typing relations that make up the proof recipe of pDOT

$\frac{\Gamma, x: T \vdash t: U \quad x \notin \text{fv}(T)}{\Gamma \vdash! \lambda(x: T) t: \forall(x: T) U} \text{ (ALL-I}_! \text{)}$	$\frac{x; \Gamma, x: T \vdash d: T}{\Gamma \vdash! \nu(x: T) d: \mu(x: T)} \text{ (}\{\!\!\}\text{-I}_! \text{)}$	<i>Precise typing for values</i> $\Gamma \vdash! v: T$
$\frac{\Gamma(x) = T}{\Gamma \vdash! x: T} \text{ (VAR}_! \text{)}$	$\frac{\Gamma \vdash! p: T \wedge U}{\Gamma \vdash! p: U} \text{ (AND}_2\text{-E}_! \text{)}$	<i>Precise-I typing</i> $\Gamma \vdash! p: T$
$\frac{\Gamma \vdash! p: \mu(z: T)}{\Gamma \vdash! p: T[p/z]} \text{ (REC-E}_! \text{)}$	$\frac{\Gamma \vdash! p: \{a: T\}}{\Gamma \vdash! p.a: T} \text{ (FLD-E}_! \text{)}$	
$\frac{\Gamma \vdash! p: T \wedge U}{\Gamma \vdash! p: T} \text{ (AND}_1\text{-E}_! \text{)}$		
$\frac{\Gamma \vdash! p: T}{\Gamma \vdash!! p: T} \text{ (PATH}_{!!} \text{)}$	$\frac{\Gamma \vdash!! p: q.\text{type} \quad \Gamma \vdash!! q.a}{\Gamma \vdash!! p.a: q.a.\text{type}} \text{ (SNGL-E}_{!!} \text{)}$	<i>Precise-II typing</i> $\Gamma \vdash!! p: T$
$\frac{\Gamma \vdash!! p: T}{\Gamma \vdash!!! p: T} \text{ (PATH}_{!!!} \text{)}$	$\frac{\Gamma \vdash!! p: q.\text{type} \quad \Gamma \vdash!!! q: U}{\Gamma \vdash!!! p: U} \text{ (SNGL-TRANS}_{!!!} \text{)}$	<i>Precise-III typing</i> $\Gamma \vdash!!! p: T$

21.2.2 Typing Judgments for pDOT's Proof Recipe

This section introduces the auxiliary typing relations that build up pDOT's proof recipe.

21.2.2.1 Three Levels of Elimination (Precise Typing) Rules

The precise typing rules (responsible for type elimination, such as REC-E) are presented in Figure 21.1. The rules are divided into three stages that closely correspond to store lookup.

The *precise-I* type represents a path's exact type as assigned by the environment (modulo possible recursion and intersection elimination). If a path is a variable x , precise-I typing is the same as DOT's precise typing. Additionally, precise-I typing can perform field selection (FLD-E_!): if x 's precise-I type is the recursive type

$$\mu(x: \cdots \wedge \{a_1: \mu(y: \cdots \wedge \{a_n: T_n\} \wedge \dots)\} \wedge \dots)$$

Figure 21.1: Precise typing in pDOT

store lookup is
defined in
Figure 19.6

see Figure 5.2 for
precise DOT typing

then the precise-I type of $x.a_1 \dots a_n$ returns the type T_n that is nested deeply in x 's type. Precise-I typing mimics the LOOKUP-STEP-VAR and LOOKUP-STEP-VAL path-lookup rules.

The purpose of *precise-II* typing is to do field selection on paths that have singleton types. If the typing environment assigns a singleton type to a path p then precise-II typing enables field selections on p . In particular, if p has $q.type$ and $q.a$ is typeable then under precise-II typing $p.a$ has $q.a.type$ (SNGL-E_{II}). This corresponds to the LOOKUP-STEP-PATH path-lookup rule.

Precise-III typing transitively follows a path's aliases through the typing environment. If a path p has a precise-II type $q.type$ (i.e. p aliases q) then precise-III typing allows p to be typechecked with any of q 's precise-III types. Precise-III typing is similar to the transitive closure of value lookup, realized by the LOOKUP-TRANS lookup rule.

Obtaining a path's precise-III type is the ultimate goal of the proof recipe. If p has type T in precise-III typing, we know that either the environment directly assigns p the type T or that p is assigned a singleton type q , and by recursively following path aliases starting with q we eventually arrive at T . More precisely, if $\Gamma \vdash_{!!!} p : T$, then one of the following is true:

1. $\Gamma \vdash_I p : T$, i.e. T is the most precise type that Γ assigns to p (modulo possible recursion and intersection elimination), or
2.
 - $p = p'.\bar{b}$ (i.e. p' is a prefix of p) and
 - $\Gamma \vdash_I p' : q.type$ (i.e. $\Gamma \vdash_{!!} p'.\bar{b} : q.\bar{b}.type$), and either
 - $T = q.\bar{b}.type$, or
 - $\Gamma \vdash_{!!!} q.\bar{b} : T$.

Precise typing is the only place in the proof recipe that encodes the SNGL-TRANS, SNGL-E, and FLD-E typing rules, along with the explicit elimination rules such as REC-E.

21.2.2.2 Two Levels of Introduction (Invertible Typing) Rules

Just like in DOT, type-introduction rules take place in *invertible* typing. However, pDOT's invertible typing is further split into two typing judgments. The reason is that the path-replacement subtyping rules $SNGL_{pq} <:$ and $SNGL_{qp} <:$ introduce new possibilities for cyclic typing derivations. For example, if $\Gamma \vdash p : q.type$ then we want to rule out an infinite derivation that changes back and forth between applying the $SNGL_{pq} <:$ and $SNGL_{qp} <:$ rules:

$$\begin{array}{c}
 \dots \quad \frac{\Gamma \vdash p : q.type}{\Gamma \vdash T <: T[q/p]} \text{SNGL}_{pq} <: \quad \frac{\Gamma \vdash p : q.type}{\Gamma \vdash T[q/p] <: T} \text{SNGL}_{qp} <: \\
 \frac{\Gamma \vdash r : T \quad \Gamma \vdash T <: T[q/p]}{\Gamma \vdash r : T[q/p]} \text{SUB} \quad \frac{\Gamma \vdash T[q/p] <: T}{\Gamma \vdash r : T} \text{SUB} \\
 \hline
 \Gamma \vdash r : T
 \end{array}$$

This derivation is infinite because it keeps alternating between typing r with T and $T [q/p]$.

In order to avoid such infinite derivations we restrict the direction in which path replacements can be performed. Specifically, we stratify invertible typing into two stages: one stage can only apply the $\text{SNGL}_{pq}\text{-<}$: replacement, the other can only apply the $\text{SNGL}_{qp}\text{-<}$: replacement. Invertible-I typing, denoted $\Gamma \vdash_i p : T$, contains introduction rules and inlined versions of the $\text{SNGL}_{pq}\text{-<}$: subtyping rule and is presented in Figure 21.2. Invertible-II typing, denoted $\Gamma \vdash_{ii} p : T$, contains inlined versions of the $\text{SNGL}_{qp}\text{-<}$: subtyping rule and is shown in Figure 21.3.

21.2.2.3 Tight Typing Rules

In DOT, new subtyping relationships are introduced through the $\text{SEL-<:}_{\text{DOT}}$ and $\text{<:-SEL}_{\text{DOT}}$ subtyping rules that allow us to use variable-dependent types. The ability to use such dependent types is restricted in tight typing, which is equivalent to general typing but does not admit the introduction of new subtyping relationships: the tight versions of the above rules, $\text{SEL-<:-\#}_{\text{DOT}}$ and $\text{<:-SEL-\#}_{\text{DOT}}$, require tight bounds and precise typing in their premises.

pDOT's tight typing, shown in Figure 21.4, has to generalize DOT's tight typing to paths for which it uses the third level of precise typing introduced in Section 21.2.2.1.

Additionally to path-dependent types, pDOT has a second form of dependent types: singleton types. The subtyping rules related to singleton types, $\text{SNGL}_{pq}\text{-<}$: and $\text{SNGL}_{qp}\text{-<}$., are handled in the same way as <:-SEL and SEL-<: . We define the tight versions of these rules, $\text{SNGL}_{pq}\text{-<:-\#}$ and $\text{SNGL}_{qp}\text{-<:-\#}$, with restricted precise-typing premises.

$$\frac{\Gamma \vdash! x : \{A : T..T\}}{\Gamma \vdash_{\#} x.A <: T} \quad (\text{SEL-<:-\#}_{\text{DOT}})$$

$$\frac{\Gamma \vdash_{!!!} p : \{A : T..T\}}{\Gamma \vdash_{\#} p.A <: T} \quad (\text{SEL-<:-\#})$$

$$\frac{\Gamma \vdash_{!!!} p : q.\text{type}}{\Gamma \vdash_{\#} T <: T [q/p]} \quad (\text{SNGL}_{pq}\text{-<:-\#})$$

21.2.3 Proof Recipe Lemmas

Given a typing $\Gamma \vdash p : T$ in an inert context Γ , the proof recipe works as follows:

1. prove that p 's *general* type T is equal to its *tight* type, i.e.

$$\Gamma \vdash_{\#} p : T$$

2. prove that p 's *tight* type is equal to its *invertible-ii* type T , i.e.

$$\Gamma \vdash_{ii} p : T$$

3. establish a relationship between p 's *invertible-II* type T and its *invertible-I* type T' , i.e.

$$\Gamma \vdash_i p : T'$$

*Invertible-I path typ-
ing*
 $\Gamma \vdash_i p : T$

$$\begin{array}{c}
\frac{\Gamma \vdash_{!!!} p : T}{\Gamma \vdash_i p : T} \quad (\text{PATH-i}) \\
\\
\frac{\Gamma \vdash_i p : \{a : S\} \quad \Gamma \vdash_{\#} S <: T}{\Gamma \vdash_i p : \{a : T\}} \quad (\text{FLD-}<:-\text{i}) \\
\\
\frac{\Gamma \vdash_i p : \{A : T..U\} \quad \Gamma \vdash_{\#} T' <: T \quad \Gamma \vdash_{\#} U <: U'}{\Gamma \vdash_i p : \{A : T'..U'\}} \quad (\text{Typ-i}) \\
\\
\frac{\Gamma \vdash_i p : \forall(x : S) T \quad \Gamma \vdash_{\#} S' <: S \quad \Gamma, x : S' \vdash T <: T'}{\Gamma \vdash_i p : \forall(x : S') T'} \quad (\text{ALL-i}) \\
\\
\frac{\Gamma \vdash_i p : S \quad \Gamma \vdash_i p : T}{\Gamma \vdash_i p : S \wedge T} \quad (\text{AND-i}) \\
\\
\frac{\Gamma \vdash_i r : \mu(x : T) \quad \Gamma \vdash_{!} p : q.\text{type} \quad \Gamma \vdash_{!!} q}{\Gamma \vdash_i r : \mu(x : T) [q/p]} \quad (\text{SNGL-REC-i}) \\
\\
\frac{\Gamma \vdash_i r : r'.A \quad \Gamma \vdash_{!} p : q.\text{type} \quad \Gamma \vdash_{!!} q}{\Gamma \vdash_i r : r'.A [q/p]} \quad (\text{SNGL-SEL-i}) \\
\\
\frac{\Gamma \vdash_i r : r'.\text{type} \quad \Gamma \vdash_{!} p : q.\text{type} \quad \Gamma \vdash_{!!} q}{\Gamma \vdash_i r : r'.\text{type} [q/p]} \quad (\text{SNGL-SNGL-i})
\end{array}$$

*Invertible-I value
typing*
 $\Gamma \vdash_i v : T$

$$\begin{array}{c}
\frac{\Gamma \vdash_{!} v : T}{\Gamma \vdash_i v : T} \quad (\text{PATH-v-i}) \\
\\
\frac{\Gamma \vdash_i v : \forall(x : S) T \quad \Gamma \vdash_{\#} S' <: S \quad \Gamma, x : S' \vdash T <: T'}{\Gamma \vdash_i v : \forall(x : S') T'} \quad (\text{ALL-v-i}) \\
\\
\frac{\Gamma \vdash_i v : S \quad \Gamma \vdash_i v : T}{\Gamma \vdash_i v : S \wedge T} \quad (\text{AND-v-i}) \\
\\
\frac{\Gamma \vdash_i v : \mu(x : T) \quad \Gamma \vdash_{!} p : q.\text{type} \quad \Gamma \vdash_{!!} q}{\Gamma \vdash_i v : \mu(x : T) [q/p]} \quad (\text{REC-SNGL-v-i}) \\
\\
\frac{\Gamma \vdash_i v : T}{\Gamma \vdash_i v : \top} \quad (\text{TOP-v-i})
\end{array}$$

Figure 21.2:
Invertible-I typing
in pDOT

$\frac{\Gamma \vdash_i p : T}{\Gamma \vdash_{ii} p : T} \quad (\text{PATH-ii})$	$\frac{\Gamma \vdash_{ii} r : \mu(x : T) \quad \Gamma \vdash_{!} p : q.\text{type} \quad \Gamma \vdash_{!!} q}{\Gamma \vdash_{ii} r : \mu(x : T) [p/q]} \quad (\text{SNGL-REC-ii})$	<i>Invertible-II path typing</i> $\Gamma \vdash_{ii} p : T$
$\frac{\Gamma \vdash_{ii} p : S \quad \Gamma \vdash_{ii} p : T}{\Gamma \vdash_{ii} p : S \wedge T} \quad (\text{AND-ii})$		
$\frac{\Gamma \vdash_{ii} p : T [p/x]}{\Gamma \vdash_{ii} p : \mu(x : T)} \quad (\text{REC-I-ii})$	$\frac{\Gamma \vdash_{ii} r : r'.A \quad \Gamma \vdash_{!} p : q.\text{type} \quad \Gamma \vdash_{!!} q}{\Gamma \vdash_{ii} r : r'.A [p/q]} \quad (\text{SNGL-SEL-ii})$	
$\frac{\Gamma \vdash_{ii} p.a : T}{\Gamma \vdash_{ii} p : \{a : T\}} \quad (\text{FLD-I-ii})$		
$\frac{\Gamma \vdash_{ii} p : T \quad \Gamma \vdash_{!} q : \{A : T..T\}}{\Gamma \vdash_{ii} p : q.A} \quad (\text{TYP-SEL-ii})$	$\frac{\Gamma \vdash_{ii} r : r'.\text{type} \quad \Gamma \vdash_{!} p : q.\text{type} \quad \Gamma \vdash_{!!} q}{\Gamma \vdash_{ii} r : r'.\text{type} [p/q]} \quad (\text{SNGL-SNGL-ii})$	
$\frac{\Gamma \vdash_i v : T}{\Gamma \vdash_{ii} v : T} \quad (\text{PATH-ii-v})$	$\frac{\Gamma \vdash_{ii} v : \mu(x : T) \quad \Gamma \vdash_{!} p : q.\text{type} \quad \Gamma \vdash_{!!} q}{\Gamma \vdash_{ii} v : \mu(x : T) [p/q]} \quad (\text{SNGL-REC-ii-v})$	<i>Invertible-II value typing</i> $\Gamma \vdash_{ii} v : T$
$\frac{\Gamma \vdash_{ii} v : S \quad \Gamma \vdash_{ii} p : T}{\Gamma \vdash_{ii} v : S \wedge T} \quad (\text{AND-ii-v})$	$\frac{\Gamma \vdash_{ii} v : T [p/x]}{\Gamma \vdash_{ii} v : \mu(x : T)} \quad (\text{REC-I-ii-v})$	
$\frac{\Gamma \vdash_{ii} v : T \quad \Gamma \vdash_{!} q : \{A : T..T\}}{\Gamma \vdash_{ii} v : q.A} \quad (\text{TYP-SEL-ii-v})$	$\frac{\Gamma \vdash_{ii} v : r'.A \quad \Gamma \vdash_{!} p : q.\text{type} \quad \Gamma \vdash_{!!} q}{\Gamma \vdash_{ii} v : r'.A [p/q]} \quad (\text{SNGL-SEL-ii-v})$	

Figure 21.3:
Invertible-II typing
in pDOT

Tight term typing
 $\Gamma \vdash_{\#} t : T$

$$\begin{array}{c}
\frac{\Gamma(x) = T}{\Gamma \vdash_{\#} x : T} \quad (\text{VAR}_{\#}) \qquad \frac{\Gamma \vdash_{\#} p : q.\text{type} \quad \Gamma \vdash q : T}{\Gamma \vdash_{\#} p : T} \quad (\text{SNGL-TRANS}_{\#}) \\
\\
\frac{\Gamma, x : T \vdash t : U \quad x \notin \text{fv}(T)}{\Gamma \vdash_{\#} \lambda(x : T) t : \forall(x : T) U} \quad (\text{ALL-I}_{\#}) \qquad \frac{\Gamma \vdash_{\#} p : q.\text{type} \quad \Gamma \vdash_{\#} q.a}{\Gamma \vdash_{\#} p.a : q.a.\text{type}} \quad (\text{SNGL-E}_{\#}) \\
\\
\frac{\Gamma \vdash_{\#} p : \forall(z : S) T \quad \Gamma \vdash_{\#} q : S}{\Gamma \vdash_{\#} p q : T[q/z]} \quad (\text{ALL-E}_{\#}) \qquad \frac{\Gamma \vdash_{\#} p : T[p/x]}{\Gamma \vdash_{\#} p : \mu(x : T)} \quad (\text{REC-I}_{\#}) \\
\\
\frac{x; \Gamma, x : T \vdash d : T}{\Gamma \vdash_{\#} \nu(x : T) d : \mu(x : T)} \quad (\text{I-I}_{\#}) \qquad \frac{\Gamma \vdash_{\#} p : \mu(x : T)}{\Gamma \vdash_{\#} p : T[p/x]} \quad (\text{REC-E}_{\#}) \\
\\
\frac{\Gamma \vdash_{\#} p : \{a : T\}}{\Gamma \vdash_{\#} p.a : T} \quad (\text{FLD-E}_{\#}) \qquad \frac{\Gamma \vdash_{\#} p.a : T}{\Gamma \vdash_{\#} p : \{a : T\}} \quad (\text{FLD-I}_{\#}) \\
\\
\frac{\Gamma \vdash_{\#} p : T \quad \Gamma \vdash_{\#} p : U}{\Gamma \vdash_{\#} p : T \wedge U} \quad (\&\text{-I}_{\#}) \\
\\
\frac{\Gamma \vdash_{\#} t : T \quad \Gamma, x : T \vdash u : U \quad x \notin \text{fv}(U)}{\Gamma \vdash_{\#} \text{let } x = t \text{ in } u : U} \quad (\text{LET}_{\#}) \qquad \frac{\Gamma \vdash_{\#} t : T \quad \Gamma \vdash T <: U}{\Gamma \vdash_{\#} t : U} \quad (\text{SUB}_{\#})
\end{array}$$

Tight subtyping
 $\Gamma \vdash_{\#} T <: U$

$$\begin{array}{c}
\Gamma \vdash_{\#} T <: \top \quad (\text{TOP}_{\#}) \qquad \frac{\Gamma \vdash_{\#} p : q.\text{type} \quad \Gamma \vdash_{\#} q}{\Gamma \vdash_{\#} T <: T[q/p]} \quad (\text{SNGL}_{pq} \text{-<:}_{\#}) \\
\\
\Gamma \vdash_{\#} \perp <: T \quad (\text{BOT}_{\#}) \\
\\
\Gamma \vdash_{\#} T <: T \quad (\text{REFL}_{\#}) \qquad \frac{\Gamma \vdash_{\#} p : q.\text{type} \quad \Gamma \vdash_{\#} q}{\Gamma \vdash_{\#} T <: T[p/q]} \quad (\text{SNGL}_{qp} \text{-<:}_{\#}) \\
\\
\frac{\Gamma \vdash_{\#} S <: T \quad \Gamma \vdash_{\#} T <: U}{\Gamma \vdash_{\#} S <: U} \quad (\text{TRANS}_{\#}) \\
\\
\Gamma \vdash_{\#} T \wedge U <: T \quad (\text{AND}_1 \text{-<:}_{\#}) \qquad \frac{\Gamma \vdash_{\#} T <: U}{\Gamma \vdash_{\#} \{a : T\} <: \{a : U\}} \quad (\text{FLD-<:}_{\#}) \\
\\
\Gamma \vdash_{\#} T \wedge U <: U \quad (\text{AND}_2 \text{-<:}_{\#}) \\
\\
\frac{\Gamma \vdash_{\#} S <: T \quad \Gamma \vdash_{\#} S <: U}{\Gamma \vdash_{\#} S <: T \wedge U} \quad (\text{<:-AND}_{\#}) \qquad \frac{\Gamma \vdash_{\#} S_2 <: S_1 \quad \Gamma \vdash_{\#} T_1 <: T_2}{\Gamma \vdash_{\#} \{A : S_1..T_1\} <: \{A : S_2..T_2\}} \quad (\text{TYP-<:-TYP}_{\#}) \\
\\
\frac{\Gamma \vdash_{\#} p : \{A : S..S\}}{\Gamma \vdash_{\#} S <: p.A} \quad (\text{<:-SEL}_{\#}) \\
\\
\frac{\Gamma \vdash_{\#} p : \{A : S..S\}}{\Gamma \vdash_{\#} p.A <: S} \quad (\text{SEL-<:}_{\#}) \\
\\
\frac{\Gamma \vdash_{\#} S_2 <: S_1 \quad \Gamma, x : S_2 \vdash T_1 <: T_2}{\Gamma \vdash_{\#} \forall(x : S_1) T_1 <: \forall(x : S_2) T_2} \quad (\text{ALL-<:-ALL}_{\#})
\end{array}$$

Figure 21.4: Tight typing for pDOT

4. establish a relationship between p 's *invertible-I* type T' and its *precise-III* type T'' , i. e.

$$\Gamma \vdash_{\text{III}} p : T''$$

This type is the type that is directly assigned to p (or its aliases) by Γ , modulo recursion and intersection elimination.

The lemmas that establish the proof recipe have to be proved in reverse order and this is how we present them below. Note that except for tight typing, each of the proof-recipe relations comes in two versions: for paths and values. We present only the lemmas for paths; the formulations for values are similar and also simpler since there are fewer rules that apply to values than to paths.

21.2.3.1 From Invertible-I to Precise-III Typing

The lemmas in this section are specialized to function, record, and singleton types. The proof-recipe lemmas for values are specialized to function and object types but we omit them here because the function-type-related lemmas are the same as for paths, and the reasoning about object types is similar to the reasoning for singleton types.

FUNCTION TYPES To convert a function type from invertible-I typing into precise-III typing, we use the following lemma.

LEMMA 36 (\vdash_i to $\vdash_{\text{III}} \forall$). *If $\Gamma \vdash_i p : \forall(x : T) U$ and Γ is inert then there exist types T' and U' such that $\Gamma \vdash_{\#} T <: T'$, $\Gamma, x : T \vdash_{\#} U' <: U$, and $\Gamma \vdash_{\text{III}} p : \forall(x : T') U'$.*

For the proof recipe, we do not need to further convert p 's type into precise-II and -I typings. However, these typing relations are needed to prove the lemmas of the proof recipe.

RECORD TYPES The following lemma converts from an invertible-I type declaration to a precise-III type:

LEMMA 37 (\vdash_i to $\vdash_{\text{III}} \{A\}$). *If $\Gamma \vdash_i p : \{A : S..U\}$ and Γ is inert then there exists a type T such that $\Gamma \vdash_{\text{III}} p : \{A : T..T\}$, $\Gamma \vdash_{\#} S <: T$ and $\Gamma \vdash_{\#} T <: U$.*

SINGLETON TYPES If a path has a precise-III singleton type T_{III} then invertible-I typing can introduce pq replacements to that type, yielding a type T_i . Subsequent invertible-II typing can introduce qp replacements to T_i , yielding a type T_{II} . To reason about the exact relationship between T_{III} , T_i , and T_{II} , we distinguish between pq - and qp -replacements, defined as follows.

DEFINITION 38. *If $\Gamma \vdash_i p : q.\text{type}$ then replacing an occurrence of the path p with q in a type is a pq -replacement, and replacing an occurrence of q with p in a type is a qp -replacement.*

$$\frac{\text{inert } \Gamma \quad \Gamma \vdash_i p : \forall(x : T) U}{\Gamma \vdash_{\text{III}} p : \forall(x : T') U'} \quad \frac{\Gamma \vdash_{\#} T <: T' \quad \Gamma, x : T \vdash_{\#} U' <: U}{\Gamma \vdash_{\text{III}} p : \forall(x : T') U'} \quad (\vdash_i \text{ TO } \vdash_{\text{III}} \forall)$$

$$\frac{\text{inert } \Gamma \quad \Gamma \vdash_i p : \{A : S..U\}}{\Gamma \vdash_{\text{III}} p : \{A : T..T\}} \quad \frac{\Gamma \vdash_{\#} S <: T <: U}{\Gamma \vdash_{\text{III}} p : \{A : T..T\}} \quad (\vdash_i \text{ TO } \vdash_{\text{III}} \{A\})$$

$$\frac{\Gamma \vdash_{\text{II}} q \quad \Gamma \vdash_i p : q.\text{type}}{\Gamma \vdash T \rightsquigarrow T[q/p]}$$

We denote the fact that U is the result of a pq replacement in T as

$$\Gamma \vdash T \rightsquigarrow U$$

Additionally, if $\Gamma \vdash T \rightsquigarrow U$ and $\Gamma \vdash S \rightsquigarrow U$ we will write

$$\Gamma \vdash T \rightsquigarrow U \leftarrow S$$

Finally, we denote the reflexive, transitive closure of \rightsquigarrow as \rightsquigarrow^* .

The following lemma establishes the relationship between a path's invertible-I singleton type and its precise-III type.

$$\frac{\text{inert } \Gamma \quad \Gamma \vdash_i p : q.\text{type}}{\Gamma \vdash_{!!!} p : q'.\text{type} \quad \Gamma \vdash q'.\text{type} \rightsquigarrow q.\text{type} \quad (\vdash_i \text{ TO } \vdash_{!!!} \forall)}$$

LEMMA 39 (\vdash_i to $\vdash_{!!!} _.\text{type}$). *If $\Gamma \vdash_i p : q.\text{type}$ and Γ is inert then there exists a path q' such that $\Gamma \vdash_{!!!} p : q'.\text{type}$ and $\Gamma \vdash q'.\text{type} \rightsquigarrow q.\text{type}$.*

21.2.3.2 From Invertible-II to Invertible-I Typing

This section presents the conversions from invertible-II to invertible-I typing for paths of function, record, and singleton types.

FUNCTION TYPES Invertible-II typing does not have rules for function types. Therefore, if a path has a function type in invertible-II typing it must have had the same type in invertible-I typing:

$$\frac{\text{inert } \Gamma \quad \Gamma \vdash_{!!} p : \forall(x : T) U}{\Gamma \vdash_i p : \forall(x : T) U \quad (\vdash_{!!} \text{ TO } \vdash_i \forall)}$$

LEMMA 40 ($\vdash_{!!}$ to $\vdash_i \forall$). *If $\Gamma \vdash_{!!} p : \forall(x : T) U$ and Γ is inert then $\Gamma \vdash_i p : \forall(x : T) U$.*

RECORD TYPES Similarly, invertible-II typing does not affect the invertible-I type of a path that is typed with a type declaration:

$$\frac{\text{inert } \Gamma \quad \Gamma \vdash_{!!} p : \{A : S..U\}}{\Gamma \vdash_i p : \{A : S..U\} \quad (\vdash_{!!} \text{ TO } \vdash_i \{A\})}$$

LEMMA 41 ($\vdash_{!!}$ to $\vdash_i \{A\}$). *If $\Gamma \vdash_{!!} p : \{A : S..U\}$ and Γ is inert then $\Gamma \vdash_i p : \{A : S..U\}$.*

SINGLETON TYPES Invertible-II typing may introduce qp replacements to singleton types. Therefore, if

$$\Gamma \vdash_{!!} p : q.\text{type}$$

then p must have had a singleton type $q'.\text{type}$ in invertible-I typing:

$$\Gamma \vdash_i p : q'.\text{type}$$

such that $q.\text{type}$ could be obtained from $q'.\text{type}$ through a sequence of qp replacements. Since a qp replacement in a singleton type is the inverse of a pq replacement⁶ we can say that $q'.\text{type}$ was obtained from $q.\text{type}$ through a series of pq replacements:

$$\Gamma \vdash q.\text{type} \rightsquigarrow^* q'.\text{type}$$

⁶ note that $q.\bar{b}.\text{type} [p/q] [q/p] = q.\bar{b}.\text{type}$

LEMMA 42 (\vdash_{ii} to \vdash_i $_.$ type). *If Γ is an inert context and $\Gamma \vdash_{\text{ii}} p : q.\text{type}$ then there exists a path q' such that $\Gamma \vdash_i p : q'.\text{type}$ and $\Gamma \vdash q.\text{type} \rightsquigarrow^* q'.\text{type}$.*

21.2.3.3 From Tight to Invertible-II Typing

Given the tight type of a path, we can convert it into an invertible-II type (cf. Theorem 10 ($\vdash_{\#}$ to \vdash_{ii})):

LEMMA 43 ($\vdash_{\#}$ to \vdash_{ii}). *If $\Gamma \vdash_{\#} p : T$ and Γ is inert then $\Gamma \vdash_{\text{ii}} p : T$.*

Just as in DOT, to prove this lemma we need to first show that invertible-II typing is closed under tight subtyping:

LEMMA 44 (Invertible-II $<$: Closure). *If Γ is inert, $\Gamma \vdash_{\text{ii}} p : T$, and $\Gamma \vdash_{\#} T <: U$ then $\Gamma \vdash_{\text{ii}} p : U$.*

The proof of the last lemma is by induction on the tight-subtyping derivation. The $\text{SNGL}_{pq} <:-\#$ and $\text{SNGL}_{qp} <:-\#$ induction cases require the proofs of the following replacement closure lemmas:

LEMMA 45 (Invertible-II qp -Replacement Closure). *If Γ is inert, $\Gamma \vdash_{\text{ii}} r : T$, $\Gamma \vdash_{\text{iii}} p : q.\text{type}$, q has a precise-II type, and T contains a path q , then $\Gamma \vdash_{\text{ii}} r : T[p/q]$.*

Since invertible-II typing explicitly inlines the $\text{SNGL}_{qp} <:-\#$ typing rule which performs qp -replacement, the proof of this lemma is straightforward. However, the proof of the analogous lemma for the reverse replacement direction (pq) is more challenging. Before presenting it we first need to prove that invertible-I typing is closed under pq replacement since it explicitly inlines the $\text{SNGL}_{pq} <:-\#$ rule:

LEMMA 46 (Invertible-I pq -Replacement Closure). *If Γ is inert, $\Gamma \vdash_i r : T$, $\Gamma \vdash_{\text{iii}} p : q.\text{type}$, q has a precise-II type and T contains a path p , then $\Gamma \vdash_i r : T[q/p]$.*

Using this lemma it is possible to prove the pq -replacement closure for invertible-II typing:

LEMMA 47 (Invertible-II pq -Replacement Closure). *If Γ is inert, $\Gamma \vdash_{\text{ii}} r : T$, $\Gamma \vdash_{\text{iii}} p : q.\text{type}$, q has a precise-II type, and T contains a path p , then $\Gamma \vdash_{\text{ii}} r : T[q/p]$.*

The proof is by induction on the invertible-II typing derivation. To illustrate a point of difficulty in the proof consider, for instance, the SNGL-SEL-ii induction case where a path r 's type T is derived from a qp -replacement. We want to prove that if we apply a pq -replacement (possibly of other paths p' and q') to T , yielding a type T' , then $\Gamma \vdash_{\text{ii}} r : T'$. However, there is no invertible-II type rule that allows pq replacements, so we need to show that the same type T' could have been derived if we had applied the pq -replacement *before* the qp -replacement that yielded T . We omit the details of this proof here and refer the interested reader instead to the Coq proof.

$$\frac{\text{inert } \Gamma \quad \Gamma \vdash_{\text{ii}} p : q.\text{type}}{\Gamma \vdash_i p : q'.\text{type}} \quad \Gamma \vdash q.\text{type} \rightsquigarrow^* q'.\text{type} \quad (\vdash_{\text{ii}} \text{ TO } \vdash_i _.\text{type})$$

$$\frac{\text{inert } \Gamma \quad \Gamma \vdash_{\#} p : T}{\Gamma \vdash_{\text{ii}} p : T} \quad (\vdash_{\#} \text{ TO } \vdash_{\text{ii}})$$

$$\frac{\text{inert } \Gamma \quad \Gamma \vdash_{\text{ii}} p : T \quad \Gamma \vdash_{\#} T <: U}{\Gamma \vdash_{\text{ii}} p : U} \quad (\vdash_{\text{ii}} <: \text{ CLOSURE})$$

$$\frac{\text{inert } \Gamma \quad \Gamma \vdash_{\text{iii}} q \quad \Gamma \vdash_{\text{iii}} p : q.\text{type} \quad \Gamma \vdash_{\text{ii}} r : T}{\Gamma \vdash_{\text{ii}} r : T[p/q]} \quad (\vdash_{\text{ii}} qp \text{ CLOSURE})$$

$$\frac{\text{inert } \Gamma \quad \Gamma \vdash_{\text{iii}} q \quad \Gamma \vdash_{\text{iii}} p : q.\text{type} \quad \Gamma \vdash_i r : T}{\Gamma \vdash_i r : T[q/p]} \quad (\vdash_i pq \text{ CLOSURE})$$

$$\frac{\text{inert } \Gamma \quad \Gamma \vdash_{\text{iii}} r \quad \Gamma \vdash_{\text{iii}} p : q.\text{type} \quad \Gamma \vdash_{\text{ii}} r : T}{\Gamma \vdash_{\text{ii}} r : T[q/p]} \quad (\vdash_{\text{ii}} pq \text{ CLOSURE})$$

$$\frac{\Gamma \vdash_{\text{iii}} q \quad \Gamma \vdash_{\text{ii}} r : r'.A \quad \Gamma \vdash_i p : q.\text{type}}{\Gamma \vdash_{\text{ii}} r : r'.A[p/q]} \quad (\text{SNGL-SEL-ii})$$

21.2.3.4 From General to Tight Typing

The translation from general into tight typing is the start of the proof recipe. Recall that to prove that DOT's general typing implies tight typing in Theorem 6 (\vdash to $\vdash_{\#}$), we needed to be able to “replace” the restricted precise-typing premise of $\text{SEL-}<:_{\text{DOT}}$ with a more general tight-typing premise, for which we proved Lemma 7 ($\text{SEL-}<:_{\#}$ Replacement) and Lemma 8 ($\text{SEL-}<:_{\#}$ Premise).

In pDOT, we have to do the same:

$$\frac{\text{inert } \Gamma \quad \Gamma \vdash_{\#} p: \{A: S..U\}}{\Gamma \vdash_{\#} p: \{A: T..T\}} \quad (\text{SEL-}<:_{\#} \text{ PREMISE})$$

LEMMA 48 ($\text{SEL-}<:_{\#}$ Premise). *If Γ is an inert context, then if $\Gamma \vdash_{\#} p: \{A: S..U\}$, then there exists a type T such that $\Gamma \vdash_{\#} p: \{A: T..T\}$, $\Gamma \vdash_{\#} S <: T$, and $\Gamma \vdash_{\#} T <: U$.*

Proof. The proof of this lemma uses the proof recipe lemmas described in the previous sections. It first converts p 's tight-typing type $\{A: S..U\}$ to the same invertible-II type using Lemma 43 and to the same invertible-I type using Lemma 41. It then uses Lemma 37 to arrive at a precise-III type $\{A: T..T\}$ for p . \square

$$\frac{\text{inert } \Gamma \quad \Gamma \vdash_{\#} p: \{A: S..U\}}{\Gamma \vdash_{\#} S <: p.A <: U} \quad (\text{SEL-}<:_{\#} \text{ REPLACEMENT})$$

LEMMA 49 ($\text{SEL-}<:_{\#}$ Replacement). *If Γ is an inert context, then if $\Gamma \vdash_{\#} p: \{A: S..U\}$, then $\Gamma \vdash_{\#} S <: p.A$ and $\Gamma \vdash_{\#} p.A <: U$.*

Proof. Directly follows from Lemma 48. \square

Similarly to strengthening the precise premises of the SEL tight-subtyping rules, we have to strengthen the precise-typing premises to tight-typing premises for the $\text{SNGI}_{pq}<:$ and $\text{SNGI}_{qp}<:$ rules:

$$\frac{\text{inert } \Gamma \quad \Gamma \vdash_{\#} q \quad \Gamma \vdash_{\#} p: q.\text{type}}{\Gamma \vdash_{\#} p: q'.\text{type}} \quad (\text{SNGI-}<:_{\#} \text{ PREMISE})$$

LEMMA 50 ($\text{SNGI-}<:_{\#}$ Premise). *If Γ is an inert context, then if $\Gamma \vdash_{\#} p: q.\text{type}$ where $\Gamma \vdash_{\#} q$, then there exists a path q' such that*

- $\Gamma \vdash_{\#} q'$,
- $\Gamma \vdash_{\#} p: q'.\text{type}$, and
- $\Gamma \vdash_{\#} q'.\text{type} <: q.\text{type}$.

Proof. By Lemma 43 ($\vdash_{\#}$ to \vdash_{ii}) p 's tight type is the same as its invertible-II type, i.e.

$$\Gamma \vdash_{\text{ii}} p: q.\text{type}$$

By Lemma 42 (\vdash_{ii} to $\vdash_i _.\text{type}$) and Lemma 39 (\vdash_i to $\vdash_{\text{iii}} _.\text{type}$),

$$\Gamma \vdash_{\text{iii}} p: q'.\text{type}$$

where

$$\Gamma \vdash q.\text{type} \rightsquigarrow^* q''.\text{type} \leftarrow^* q'.\text{type}$$

which allows us to prove that $\Gamma \vdash_{\#} q'.\text{type} <: q''.\text{type} <: q.\text{type}$. Note that the actual proof needs to use stronger versions of Lemma 42 and Lemma 39 in order to prove that q' is typeable; however, we leave these details out to just display the main idea of the proof. \square

$$\begin{array}{c}
\vdash \emptyset \Diamond \quad (\text{TP-EMPTY}) \\
\\
\frac{\vdash \Gamma \Diamond \quad x \notin \text{fv}(\Gamma) \quad \forall \bar{b}, q, \frac{\Gamma, x: T \vdash! x.\bar{b}: q.\text{type}}{\exists U, \Gamma, x: T \vdash!! q: U}}{\vdash \Gamma, x: T \Diamond} \quad (\text{TP})
\end{array}$$

The proof of the next lemma immediately follows from Lemma 50.

LEMMA 51 (SNGL-<:# Replacement). *If Γ is an inert context, then if $\Gamma \vdash_{\#} p: q.\text{type}$ where $\Gamma \vdash_{\#} q$, then for any type T that contains a path p , $\Gamma \vdash_{\#} T <: T[q/p]$, and for any type U that contains a path q , $\Gamma \vdash_{\#} U <: U[p/q]$.*

21.3 TYPED-PATHS ENVIRONMENTS

A pDOT type can depend on paths rather than just variables, and the type makes sense only if the paths within it make sense, i. e. if all paths have a type. To ensure that all paths in a type are typeable, we introduce the notion of *typed-paths environments*. The typed-path property requires that any path that appears in a type should itself also be typeable in the same typing context. Without this property, it would be possible for the typing rules to derive types for ill-formed paths, and there could be paths that have types but do not resolve to any value during program execution. We will need the typed-path property when we formulate the canonical-forms, progress, and preservation lemmas.

The precise definition of typed-path environments, denoted $\vdash \Gamma \Diamond$, is given in Figure 21.5. The WT rule defines how to extend a typed-path environment Γ with a type T to maintain the typed-paths property. Specifically, it considers all the paths $(x.\bar{b})$ that are introduced by T and, if a path has a singleton type $q.\text{type}$, requires q to be precisely typeable in an environment.

How does the rule retrieve all paths that are introduced by T ? For that we just need to consider all precise-I typings of paths that start with x (recall that precise-I typing gives us information about the exact type, modulo recursion and intersection elimination, that an environment assigns to a path). For example, if $T = \mu(y: \{a: y.a.a.\text{type}\})$, there are the following precise-I typings that T introduces:

$$\Gamma, x: T \vdash! x: \mu(y: \{a: y.a.a.\text{type}\}) \quad (21.1)$$

$$\Gamma, x: T \vdash! x: \{a: x.a.a.\text{type}\} \quad (21.2)$$

$$\Gamma, x: T \vdash! x.a: x.a.a.\text{type} \quad (21.3)$$

The only typing applicable to our definition is (21.3).

Finally, how does the rule ensure that the paths in the above singleton types are typeable? For that it is not sufficient to require precise-I

Figure 21.5:
Typed-paths
environments

$$\begin{array}{c}
\text{inert } \Gamma \quad \Gamma \vdash_{\#} q \\
\hline
\Gamma \vdash_{\#} p: q.\text{type} \\
\hline
\Gamma \vdash_{\#} T <: T[q/p] \\
\hline
\Gamma \vdash_{\#} U <: U[p/q] \\
(\text{SNGL-<:# REPLACEMENT})
\end{array}$$

see Figure 21.1 for
precise typing

typing: precise-I typing only allows us to do field selection on paths that have recursive or record types, whereas we also need to do field selection on paths that have singleton types. In the above example, the path $x.a.a$ has a precise-II type but no precise-I type.

Note that we also want our type system to allow paths that have been accessed through precise-III typing, which transitively follows path aliases in the environment. However, it can be easily shown that every precise-III typeable path also has a precise-II typeable path, which is what we require in the above definition of typed-paths environments.

21.4 CANONICAL FORMS IN PDOT

As described in Section 5.5, the DOT proof depends on canonical forms lemmas that state that if a variable has a function type, then it resolves to a corresponding function at execution time, and if it has a recursive object type, then it resolves to a corresponding object. The change from DOT to pDOT involves several changes to Canonical Forms.

Two changes follow directly from pDOT's operational semantics. The DOT canonical forms lemmas apply to variables. Since the pDOT APPLY reduction rule applies to paths rather than variables, the canonical forms lemma is needed for paths. Since paths are normal forms in pDOT and there is no PROJ reduction rule for them, on the surface, pDOT needs a canonical forms lemma only for function types but not for object types. However, to reason about a path with a function type, we need to reason about the prefixes of the path, which have an object type. Therefore, the induction hypothesis in the canonical forms lemma for function types must still include canonical forms for object types. Moreover, since pDOT adds singleton types to the type system, the induction hypothesis needs to account for them as well.

A more subtle but important change is that lookup of a path in an execution environment is a recursive operation in pDOT, and therefore its termination cannot be taken for granted. An infinite loop in path lookup would be a hidden violation of progress for function application, since the APPLY reduction rule steps only once path lookup has finished finding a value for the path. Therefore, the canonical forms lemma proves that if a path has a function type, then lookup of that path does terminate, and the value with which it terminates is a function of the required type. The intuitive argument for termination requires connecting the execution environment with the typing environment: if direct lookup of path p yields another path q , then the context assigns p the singleton type $q.type$. But in order for p to have a function type, there cannot be a cycle of paths in the typing context (because a cycle would limit p to have only singleton types),

and therefore there cannot be a cycle in the execution environment. The statement of the canonical forms lemma is:

LEMMA 52 (Canonical Forms for Functions). *Let γ be a store and Γ be an inert, typed-paths environment such that $\gamma: \Gamma$. If $\Gamma \vdash p: \forall(x: T) U$ then there exists a type T' and a term t such that*

1. $\gamma \vdash p \rightsquigarrow^* \lambda(x: T') t$,
2. $\Gamma \vdash T <: T'$, and
3. $\Gamma, x: T \vdash t: U$.

This simple statement hides an intricate induction hypothesis and a long, tedious proof, since it needs to reason precisely about function, object, and singleton types and across all seven typing relations in the stratification of typing. We describe the key lemmas that are necessary for this proof below.

$$\frac{\text{inert } \Gamma \quad \gamma: \Gamma \quad \vdash \Gamma \Diamond \quad \Gamma \vdash p: \forall(x: T) U}{\gamma \vdash p \rightsquigarrow^* \lambda(x: T') t \quad \Gamma \vdash T <: T' \quad \Gamma, x: T \vdash t: U} \text{(CANONICAL FORMS } \forall \text{)}$$

21.4.1 Canonical Forms Proof

One difficulty in proving Canonical Forms for pDOT is to establish a correspondence between a typing environment and a store. In DOT, if a store γ is well-formed with respect to a typing environment Γ ($\gamma: \Gamma$), then it immediately follows that any variable that has a type in Γ is mapped to a value in γ . We would like the same property for paths: if $\Gamma \vdash p$ then we would like there to exist a *value v of the same type* such that $\gamma \vdash p \rightsquigarrow^* v$. However, there might not be such a value: it is possible that p cyclically references other paths. Fortunately, since we only need a canonical-forms lemma for functions, we can prove a more restricted form of the corresponding-types lemma for function types. We split this lemma into two. The first one states that if a path p has a function type then p looks up to a value in the store.

LEMMA 53 (Corresponding Values for Functions). *Suppose that a store γ is well-formed with respect to an inert, typed-paths environment Γ , and that Γ assigns type $\forall(x: T) U$ to a path p , i.e. $\Gamma \vdash_{!!!} p: \forall(x: T) U$. Then there exists a value v that is assigned to p by the store γ , i.e. $\gamma \vdash p \rightsquigarrow^* v$.*

$$\frac{\text{inert } \Gamma \quad \gamma: \Gamma \quad \vdash \Gamma \Diamond \quad \Gamma \vdash_{!!!} p: \forall(x: T) U}{\gamma \vdash p \rightsquigarrow^* v} \text{(CORRESP. VALUES } \forall \text{)}$$

The purpose of the second lemma is to show that if a path p has a function type and looks up to a value v in the store (as established by the previous lemma) then v has the same type as p . To prove that we state a more general lemma. It says that store lookup preserves function types:

LEMMA 54 (Corresponding Types for Functions). *Suppose that a store γ is well-formed with respect to an inert, typed-paths environment Γ , and that Γ assigns type $\forall(x: T) U$ to a stable term s (recall that a stable term is a path or a value), i.e. $\Gamma \vdash s: \forall(x: T) U$, and s looks up to s' in the store, i.e. $\gamma \vdash s \rightsquigarrow^* s'$. Then s' has the same type as s , i.e. $\Gamma \vdash s': \forall(x: T) U$.*

$$\frac{\text{inert } \Gamma \quad \gamma: \Gamma \quad \vdash \Gamma \Diamond \quad \gamma \vdash s \rightsquigarrow^* s' \quad \Gamma \vdash s: \forall(x: T) U}{\Gamma \vdash s': \forall(x: T) U} \text{(CORRESP. TYPES } \forall \text{)}$$

We present the main lemmas necessary to prove the above two.

The following lemma states that if a path p has a precise-III function type T then p either has type T under precise-II typing, or p aliases another path q that has T under precise-II typing.

LEMMA 55. *If Γ is an inert environment and $\Gamma \vdash_{!!!} p : \forall(x : T) U$ then either*

- $\Gamma \vdash_{!} p : \forall(x : T) U$, or
- *there exists a path q such that $\Gamma \vdash_{!!!} p : q.\text{type}$ and $\Gamma \vdash_{!} q : \forall(x : T) U$.*

Proof. By induction on the precise-III typing derivation of p . \square

This simple lemma states that opening a recursive type in the typing environment does not affect the typing of a term.

LEMMA 56. *If $\Gamma, x : \mu(x : U) \vdash t : T$ then $\Gamma, x : U \vdash t : T$.*

The following lemma establishes a correspondence between the precise-I type of a path and the term that it looks up to in the store.

LEMMA 57 (Lookup Preservation I). *Let γ be a store that is well-formed with respect to an inert, typed-paths environment Γ , and for some path p , stable term s , and type U ,*

- $\gamma \vdash p \rightsquigarrow s$
- $\Gamma \vdash_{!} p : U$.

Then one of the following is true:

1. *s is a function value of type U ,*
2. *there exist types T, S, W and definitions d such that*
 - *s is an object $v(S : d)$,*
 - $\Gamma \vdash_{!} p : \mu(y : T)$,
 - $p; \Gamma \vdash d[p/y] : S[p/y]$, and
 - $\Gamma \vdash S \rightsquigarrow^* W^* \leftarrow T$ for some type W , or
3. *there exist paths q, r , and r' such that*
 - $t = q$,
 - $U = r.\text{type}$, and
 - $\Gamma \vdash q.\text{type} \rightsquigarrow^* r'.\text{type}^* \leftarrow r.\text{type}$

see Definition 16 for the definition of $\gamma : \Gamma$

Proof. The proof is by induction on the derivation of $\gamma : \Gamma$. The empty-environment case cannot happen since a path cannot be typed in an

empty environment. In the inductive case, suppose that $p = x_p.\bar{b}$; then we have

$$\text{inert } (\Gamma, x: T) \quad (21.4)$$

$$\gamma: \Gamma \quad (21.5)$$

$$\vdash \Gamma, x: T \quad (21.6)$$

$$\Gamma \vdash v: T \quad (21.7)$$

$$\Gamma, x: T \vdash! x_p.\bar{b}: U \quad (21.8)$$

$$\gamma, x \mapsto v \vdash x_p.\bar{b} \rightsquigarrow t \quad (21.9)$$

We consider two cases depending on whether x is equal to x_p .

Case 1: $x = x_p$.

We proceed by induction on the lookup-derivation (21.9).

see Figure 19.6 for
the lookup definition

Case 1.i: LOOKUP-STEP-VAR. In this case, $p = x, s = v$. If v is a function then it is easy to show that $T = U$, and we have proved that s is a function value of type U as required.

If v is an object, then since by (21.4), T is inert, T must be a function, object, or singleton type. Applying the the proof recipe to v 's typing (21.7) in each case, we can rule out the function and singleton-type cases and obtain that

$$T = \mu(y: T') \quad (21.10)$$

$$\Gamma \vdash! v(y: S)d: \mu(y: S) \quad (21.11)$$

$$\Gamma \vdash S \rightsquigarrow^* U' \leftarrow T' \quad (21.12)$$

$$\Gamma, x: \mu(y: T') \vdash! x: \mu(y: T') \quad (21.13)$$

From here the result can be obtained by inverting (21.11) and using Narrowing (see Lemma 15) and Lemma 56.

Case 1.ii: LOOKUP-STEP-VAL. We have $p = x.\bar{b}.a$ and

$$\gamma, x \mapsto v \vdash x.\bar{b} \rightsquigarrow v(y: T')d \quad (21.14)$$

$$\Gamma, x: T \vdash! x.\bar{b}.a: U \quad (21.15)$$

where $d = \dots \wedge \{a = s\} \wedge \dots$. To apply the induction hypothesis to (21.14) we need to know the precise-I type of $x.\bar{b}$. It is easy to show that (21.15) implies that for some type V ,

$$\Gamma, x: T \vdash! x.\bar{b}: \mu(y: V) \quad (21.16)$$

The only applicable case in the result of applying the induction hypothesis to (21.14) is (2) because $x.\bar{b}$'s type must be recursive. This yields

$$x.\bar{b}; \Gamma, x: T \vdash d[x.\bar{b}/y]: S[x.\bar{b}/y] \quad (21.17)$$

$$\Gamma \vdash S \rightsquigarrow^* W \leftarrow V \quad (21.18)$$

Since the definitions d contain a record $\{a = s\}$ we can infer that

$$V = \dots \wedge \{a: V'\} \wedge \dots \quad (21.19)$$

$$\Gamma, x: T \vdash s: V' \quad (21.20)$$

Using (21.19) and (21.18) we can then show that

$$W = \dots \wedge \{a: W'\} \wedge \dots \quad \text{and} \quad \Gamma \vdash V' \rightsquigarrow^* W' \quad (21.21)$$

$$S = \dots \wedge \{a: S'\} \wedge \dots \quad \text{and} \quad \Gamma \vdash S' \rightsquigarrow^* W' \quad (21.22)$$

The remainder of the proof is based on a case analysis on the shapes of s and V' and is straightforward.

Case 1.iii: LOOKUP-STEP-PATH. In this case, $p = x.\bar{b}.a$, $s = q.a$, and we have

$$\gamma, x \mapsto v \vdash x.\bar{b} \rightsquigarrow q \quad (21.23)$$

$$\Gamma, x: T \vdash! x.\bar{b}.a: U \quad (21.24)$$

To apply the induction hypothesis to (21.23) we proceed as in the LOOKUP-STEP-VAL case, obtaining

$$\Gamma, x: T \vdash! x.\bar{b}: \mu(y: V) \quad (21.25)$$

The only applicable case in the result of applying the induction hypothesis to (21.23) is (3) since there, s is a path. However, in that case $x.\bar{b}$'s precise-I type is a singleton type which is incompatible with its recursive type according to (21.25), a contradiction.

Case 2: $x \neq x_p$. We have

$$\Gamma, x: T \vdash! x_p.\bar{b}: U \quad (21.26)$$

$$\gamma, x \mapsto v \vdash x_p.\bar{b} \rightsquigarrow s \quad (21.27)$$

Since the receiver x_p of the path $x_p.\bar{b}$ is not equal to x , x_p must be contained in Γ , which means that the whole path $x_p.\bar{b}$ must be typeable in Γ :

$$\Gamma \vdash! x_p.\bar{b}: U \quad (21.28)$$

$$(21.29)$$

By similar reasoning we can look up the path in γ :

$$\gamma \vdash x_p.\bar{b} \rightsquigarrow s \quad (21.30)$$

This allows us to apply the induction hypothesis and obtain the result we need to prove.

□

Lemma 57 establishes a connection between a path's precise-I type and the term it looks up to in the store. The following lemma does the same for a path's precise-II type.

LEMMA 58 (Lookup Preservation II). *Let γ be a store that is well-formed with respect to an inert, typed-paths environment Γ , and for some path p , stable term s , and type T ,*

- $\gamma \vdash p \rightsquigarrow s$ and
- $\Gamma \vdash_{!!} p : T$.

Then one of the following is true:

1. *there exists a type S and term u such that*
 - *s is a function $\lambda(y : S) u$,*
 - *$\Gamma \vdash s : T$, and*
 - *$\Gamma \vdash_{!} p : T$, or*
2. *there exist types S and W , definitions d , and a path p such that*
 - *t is an object $v(y : S)d$,*
 - *$\Gamma \vdash_{!} p : \mu(z : U)$,*
 - *$\Gamma \vdash_{!} p : T$,*
 - *$p; \Gamma \vdash d[p/y] : S[p/y]$, and*
 - *$\Gamma \vdash S \rightsquigarrow^* W^* \leftarrow U$, or*
3. *there exist paths q , r , and r' such that*
 - *$s = q$,*
 - *$T = r.\text{type}$, and*
 - *$\Gamma \vdash r.\text{type} \rightsquigarrow^* r'.\text{type}^* \leftarrow q.\text{type}$.*

Note that a path's precise-I typing does not have to be unique since precise-I typing can apply type-elimination rules

Proof. The proof is by induction on the precise-II derivation of p .

Case 1: PATH_{!!}. In this case, we have a precise-I derivation for p , and the result follows from Lemma 57.

Case 2: SNGL-E_{!!}. We have $T = q.a.\text{type}$ and

$$\Gamma \vdash_{!!} p : q.\text{type} \tag{21.31}$$

$$\Gamma \vdash_{!!} q.a : U \tag{21.32}$$

$$\gamma \vdash p.a \rightsquigarrow s \tag{21.33}$$

We proceed by induction on the lookup derivation (21.33). The variable-case is immediately ruled since our path ends with a field selection.

Case 2.i: LOOKUP-STEP-VAL. This case can be ruled out by applying the outer induction hypothesis.

see Figure 21.1 for the definition of precise-II typing

Case 2.ii: LOOKUP-STEP-PATH. We have

$$\gamma \vdash p \rightsquigarrow q' \quad (21.34)$$

We apply the outer induction hypothesis, which yields only one possible case (3) where for some path r' ,

$$\Gamma \vdash q.\text{type} \rightsquigarrow^* r'.\text{type} \leftarrow q'.\text{type} \quad (21.35)$$

We now have $s = q'.a$ and $T = q.a.\text{type}$. Equations (21.31) and (21.32) imply that $\Gamma \vdash_{!!} p : q.a.\text{type}$, and from (21.35) it follows that

$$\Gamma \vdash q.a.\text{type} \rightsquigarrow^* r'.a.\text{type} \leftarrow q'.a.\text{type} \quad (21.36)$$

which is what we needed to show. \square

Next, we prove a lemma similar to the above two lemmas (57 and 58) that establishes a correspondence between a path's precise-III typing and the term it looks up in the store. Note that we only prove this lemma for the case where a path's type is inert. The case where a path has a precise-III singleton type is proved separately.

LEMMA 59 (Lookup Preservation III-Inert). *Let γ be a store that is well-formed with respect to an inert, typed-paths environment Γ , and for some path p , stable term s , and type T ,*

- $\gamma \vdash p \rightsquigarrow s$ and
- $\Gamma \vdash_{!!!} p : T$ where T is inert.

Then one of the following is true:

- T is a function type $\forall(x : S) U$, and $\Gamma \vdash s : \forall(x : S) U$, or
- T is a recursive type $\mu(x : U)$, and there exist types S and W , definitions d , and a path p such that
 - t is an object $v(y : S)d$,
 - $p; \Gamma \vdash d[p/y] : S[p/y]$, and
 - $\Gamma \vdash S \rightsquigarrow^* W \leftarrow U$, or
- t is a path q , and $\Gamma \vdash_{!!!} q : T$.

see Figure 21.1 for
the definition of
precise-III typing

Proof. The proof is by induction on the precise-III typing of p . In each case, we get a precise-II typing for p to which we apply Lemma 58 and consider the three possible cases that the lemma yields. The proof for each case involves low-level reasoning about properties of precise-I, -II, and -III typing. We omit the details of the proof here and refer the reader to the Coq version instead. \square

The following lemma states that if a path $x.\bar{b}$ looks up to a path $x.\bar{c}$ with the same receiver x , then the path's precise-II type is $x.\bar{c}.\text{type}$.

LEMMA 60. *Let γ be a store that is well-formed with respect to an inert, typed-paths environment Γ , and*

- $\gamma \vdash x.\bar{b} \rightsquigarrow x.\bar{c}.\text{type}$
- $\Gamma \vdash_{!!} x.\bar{b} : T$

Then $T = x.\bar{c}.\text{type}$.

Proof. The proof is by induction on $\gamma : \Gamma$. In the inductive case with the extended environment $\Gamma, y : U$, we distinguish between whether x is equal to y . If it is not, we simply apply the induction hypothesis. If $x = y$ then we finish the proof using Lemma 58. \square

The next lemma establishes that if a path has a precise-I type then it can be looked up in the store.

LEMMA 61. *Let γ be a store that is well-formed with respect to an inert, typed-paths environment Γ , and p has a precise-I type. Then there exists a stable term s to which p can be looked up in the store, i.e. $\gamma \vdash p \rightsquigarrow s$.*

Proof. Similarly to the proof of Lemma 57, we perform two inductions. The outer induction is on $\gamma : \Gamma$. In the inductive case where we have an extended environment $\Gamma, x : T$, we do a case analysis on whether p 's receiver x_p is equal to x . If $x = x_p$ we induct on the precise-I typing of p . The interesting case is the field-selection case FLD-E! where we have

$$\Gamma, x : T \vdash_{!} x.\bar{b} : \mu(y : U) \quad (21.37)$$

$$\Gamma, x : T \vdash_{!} x.\bar{b} : \{a : U'\} \quad (21.38)$$

and we need to prove that $x.\bar{b}.a$ looks up to a term. By induction,

$$\gamma, x \mapsto v \vdash x.\bar{b} \rightsquigarrow s' \quad (21.39)$$

Applying Lemma 57 to (21.38) and (21.39) yields only one possible case where for some types S, W , and definitions d ,

$$s' = v(z : S)d \quad (21.40)$$

$$x.\bar{b}; \Gamma, x : T \vdash d[x.\bar{b}/z] : S[x.\bar{b}/z] \quad (21.41)$$

$$\Gamma \vdash S \rightsquigarrow^* W \leftarrow^* U \quad (21.42)$$

From (21.37) and (21.38) we can infer that for some type U' ,

$$U = \dots \wedge \{a : U'\} \wedge \dots \quad (21.43)$$

because if a path has a recursive type T and a record type U , it means that U was obtained by performing recursion elimination followed by

intersection elimination on T . This allows us to conclude from (21.42) that for some types S' and W' ,

$$S = \cdots \wedge \{a: S'\} \wedge \cdots \quad (21.44)$$

$$W = \cdots \wedge \{a: W'\} \wedge \cdots \quad (21.45)$$

$$\Gamma \vdash S' \rightsquigarrow^* W'^* \leftarrow U' \quad (21.46)$$

because pq - and qp -replacements preserve a type's shape. Finally, we can show using (21.44) and (21.41) that the definitions d must match the type S , and therefore

$$d = \cdots \wedge \{a = s''\} \wedge \cdots \quad (21.47)$$

for some stable term s'' . Finally, this allows us to conclude that

$$\gamma \vdash x.\bar{b}.a \rightsquigarrow s''$$

□

The following two lemmas state the same for precise-II and -III typing: if a path has a precise-II and -III type then it can be looked up in the store.

LEMMA 62. *Let γ be a store that is well-formed with respect to an inert, typed-paths environment Γ , and p has a precise-II type. Then there exists a stable term s to which p can be looked up in the store, i.e. $\gamma \vdash p \rightsquigarrow s$.*

Proof. The proof is by induction on p 's precise-II derivation. The $\text{PATH}_{!!}$ case is proved using Lemma 61. In the $\text{SNGL-E}_{!!}$ case, we use Lemma 58 and simple properties of precise typing. □

LEMMA 63. *Let γ be a store that is well-formed with respect to an inert, typed-paths environment Γ , and p has a precise-III type. Then there exists a stable term s to which p can be looked up in the store, i.e. $\gamma \vdash p \rightsquigarrow s$.*

Proof. The proof is by induction on p 's precise-III typing; in both cases it is easily solved using Lemma 62. □

If a path shares the same receiver with its precise-II singleton type $q.\text{type}$ (i.e. $\Gamma \vdash_{!!} x.\bar{b}: x.\bar{c}.\text{type}$) then the path looks up to q in the store.

LEMMA 64. *Let γ be a store that is well-formed with respect to an inert, typed-paths environment Γ , and $\Gamma \vdash_{!!} x.\bar{b}: x.\bar{c}.\text{type}$. Then $\gamma \vdash x.\bar{b} \rightsquigarrow x.\bar{c}$.*

Proof. By Lemma 62, there exists a stable term s such that

$$\gamma \vdash x.\bar{b} \rightsquigarrow s$$

The result can be then proved using Lemma 58. □

If a path shares the same receiver with its precise-III singleton type $q.\text{type}$ (i.e. $\Gamma \vdash_{!!!} x.\bar{b}: x.\bar{c}.\text{type}$) then the path looks up to q in the store in a finite number of steps.

LEMMA 65. Let γ be a store that is well-formed with respect to an inert, typed-paths environment Γ , and $\Gamma \vdash_{!!!} x.\bar{b} : x.\bar{c}.\text{type}$. Then $\gamma \vdash x.\bar{b} \rightsquigarrow^* x.\bar{c}$.

Proof. We proceed by induction on the precise-III derivation of p . The PATH_{!!!} case follows from Lemma 64. In the SNGL-E_{!!!} case, we have for some path q

$$\Gamma \vdash_{!!} x.\bar{b} : q.\text{type} \quad (21.48)$$

$$\Gamma \vdash_{!!!} q : x.\bar{c}.\text{type} \quad (21.49)$$

Suppose that $q = x_q.\bar{b}_q$. We distinguish between two cases depending on whether x is equal to x_q .

Case 1: $x = x_q$. By the induction hypothesis applied to (21.49),

$$\gamma \vdash x.\bar{b}_q \rightsquigarrow^* x.\bar{c} \quad (21.50)$$

By Lemma 64 applied to (21.48),

$$\gamma \vdash x.\bar{b} \rightsquigarrow x.\bar{b}_q \quad (21.51)$$

The result follows from (21.50) and (21.51).

Case 2: $x \neq x_q$. If $x.\bar{b}$'s precise-II type is $x_q.\bar{b}_q.\text{type}$ where $x \neq x_q$ it means that the variable x_q must occur before x in the environment (otherwise the environment would not be a typed-paths environment since we would not be able to have a type for $x_q.\bar{b}_q$ at the moment when x has this path as its singleton type). However, according to (21.49), $x_q.\bar{b}_q$'s precise type is $x.\bar{c}.\text{type}$, which means that x must be defined before x_q by analogous reasoning. Since each variable occurs only once in the domain of an inert environment, we arrive at a contradiction. \square

The following lemma shows that if $x.\bar{b}$ has precise-III type $y.\bar{c}$, and $x \neq y$, then there must exist a type $x.\bar{b}'.\text{type}$ that is the “last” precise-III type of $x.\bar{b}$ that has x as a receiver.

LEMMA 66. Let γ be a store that is well-formed with respect to an inert, typed-paths environment Γ , and

$$\Gamma \vdash_{!!!} x.\bar{b} : y.\bar{c}.\text{type}$$

where $x \neq y$. Then there exist paths $x.\bar{b}'$ and $z.\bar{c}'$ such that $x \neq z$ and

$$- \Gamma \vdash_{!!!} x.\bar{b} : x.\bar{b}'.\text{type} \text{ or } x.\bar{b} = x.\bar{b}'$$

$$- \Gamma \vdash_{!!} x.\bar{b}' : z.\bar{c}'.\text{type}$$

$$- \Gamma \vdash_{!!!} z.\bar{c}' : y.\bar{c}.\text{type} \text{ or } z.\bar{c}' = y.\bar{c}$$

Proof. The proof is by induction on the precise-III typing of $x.\bar{b}$. The interesting case is SNGL-TRANS_{!!!}. We do a case analysis based on whether y is equal to z . If it is the proof follows immediately. If $y \neq z$ we have

$$\Gamma \vdash_{!!!} z.\bar{c}': y.\bar{c}.\text{type} \quad (21.52)$$

$$\Gamma \vdash_{!!} x.\bar{b}: z.\bar{c}'.\text{type} \quad (21.53)$$

By the induction hypothesis applied to (21.52), there exist paths $z.\bar{c}_1$ and $z_2.\bar{c}_2$ such that $z \neq z_2$ and

$$\Gamma \vdash_{!!!} z.\bar{c}': z.\bar{c}_1.\text{type} \quad \text{or} \quad z.\bar{c}' = z.\bar{c}_1 \quad (21.54)$$

$$\Gamma \vdash_{!!} z.\bar{c}_1: z_2.\bar{c}_2.\text{type} \quad (21.55)$$

$$\Gamma \vdash_{!!!} z_2.\bar{c}_2: y.\bar{c}.\text{type} \quad \text{or} \quad z_2.\bar{c}_2 = y.\bar{c} \quad (21.56)$$

The proof can be then finished by a case analysis on whether x is equal to z . \square

In a typed-paths environment, there is a precise-typing relation between equivalent singleton paths

LEMMA 67. *If Γ is an inert, typed-paths environment and $\Gamma \vdash p.\text{type} \rightsquigarrow^* q.\text{type}$ where p is a precise-III typeable path then either $p = q$ or $\Gamma \vdash_{!!!} p: q.\text{type}$.*

Proof. The proof is by induction on the replacement closure relation. We omit the details here. \square

We want to be able to prove that if a path p has a precise-III singleton type $r.\text{type}$ then p can be looked up to r . However, we can only prove that if we know that p and r do not participate in an aliasing cycle. To establish that, we use the fact that r 's type is a function type which guarantees that eventually the lookup chain results in a value. The following lemma says exactly this: if a path p 's precise-III type is $r.\text{type}$ and r has a precise-I function type then p can be looked up to r in the store in a finite number of steps.

LEMMA 68. *Let γ be a store that is well-formed with respect to an inert, typed-paths environment Γ ,*

$$- \Gamma \vdash_{!!!} p: r.\text{type}, \text{ and}$$

$$- \Gamma \vdash_{!} r: \forall(x: S) T.$$

Then $\gamma \vdash p \rightsquigarrow^ r$.*

Proof. As we have done before, we start with an induction on $\gamma: \Gamma$. In the inductive case, we have an extended environment $\Gamma, x: T$ and a path $p = x_p.\bar{b}$. We distinguish between whether x is equal to x_p . The

interesting case is when $x = x_p$ (the induction hypothesis takes care of the other case). We have

$$\Gamma, x: T \vdash_{\vdash} r: \forall(y: S) V \quad (21.57)$$

$$\Gamma, x: T \vdash_{\vdash} x.\bar{b}: r.\text{type} \quad (21.58)$$

and we need to prove that $\gamma, x \mapsto v \vdash x.\bar{b} \rightsquigarrow^* r$.

By Lemma 63, there exists a stable term s such that

$$\gamma, x \mapsto v \vdash x.\bar{b} \rightsquigarrow s \quad (21.59)$$

By Lemma 59 applied to (21.58) and (21.59), there exist paths r' and r'' such that $s = r'$ and

$$\Gamma, x: T \vdash r.\text{type} \rightsquigarrow^* r''.\text{type} \leftarrow r'.\text{type} \quad (21.60)$$

Furthermore, because r has a precise-I function type by (21.57), r and all its prefixes cannot have a precise-I or -II singleton type (we do not present the proof here). This means that the only way how the $\Gamma, x: T \vdash r.\text{type} \rightsquigarrow^* r''.\text{type}$ could have been obtained in (21.60) is if $r = r''$, i. e.

$$\Gamma, x: T \vdash r'.\text{type} \rightsquigarrow^* r.\text{type} \quad (21.61)$$

Suppose now that $r = x_r.\bar{c}$. We proceed by case analysis on whether x is equal to x_r .

Case 1: $x = x_r$. Follows from Lemma 65.

Case 2: $x \neq x_r$. By Lemma 66 applied to (21.58), there exist paths $x.\bar{b}'$ and $z.\bar{c}'$ such that

$$x \neq z \quad (21.62)$$

$$\Gamma, x: T \vdash_{\vdash} x.\bar{b}: x.\bar{b}'.\text{type} \quad (21.63)$$

$$\Gamma, x: T \vdash_{\vdash} x.\bar{b}': z.\bar{c}'.\text{type} \quad (21.64)$$

$$\Gamma, x: T \vdash_{\vdash} z.\bar{c}': x_r.\bar{c}.\text{type} \quad (21.65)$$

By Lemma 62 applied to (21.64), there exists a stable term s' such that

$$\gamma, x \mapsto v \vdash x.\bar{b}' \rightsquigarrow s' \quad (21.66)$$

Applying Lemma 58 to (21.64) and (21.66), after eliminating the impossible cases we obtain that there exist paths q and q' such that $s' = q$ and

$$\Gamma, x: T \vdash z.\bar{c}'.\text{type} \rightsquigarrow^* q'.\text{type} \leftarrow q.\text{type} \quad (21.67)$$

we leave out the cases of Lemma 66 when $x.\bar{b} = x.\bar{b}'$ or $z.\bar{c}' = x_r.\bar{c}$ here because they are simpler

By Lemma 65 applied to (21.63),

$$\gamma, x \mapsto v \vdash x.\bar{b} \rightsquigarrow^* x.\bar{b}' \quad (21.68)$$

By (21.68) and (21.66) we now know that

$$\gamma, x \mapsto v \vdash x.\bar{b} \rightsquigarrow^* q \quad (21.69)$$

and we need to prove that $\gamma, x \mapsto v \vdash x.\bar{b} \rightsquigarrow^* x_r.\bar{c}$. Therefore, it remains to show that $\gamma, x \mapsto v \vdash q \rightsquigarrow^* x_r.\bar{c}$. To do that we apply Lemma 67 to the three qp -replacement closures in (21.61) and (21.69), each of which yields a disjunction of equalities or precise-III typing relationships between the involved paths. The proof is based on a case analysis of the resulting eight cases and involves low-level reasoning about the properties of precise typing. The resulting cases are either ruled out due to contradictions or by applying the induction hypothesis.

□

$$\frac{\text{inert } \Gamma \quad \gamma: \Gamma \vdash \Gamma \Diamond \quad \Gamma \vdash_{!!!} p: \forall(x: T) U}{\gamma \vdash p \rightsquigarrow^* v} \quad (\text{CORRESP. VALUES } \forall)$$

We can now sketch the proof of Lemma 53 (Corresp. Values \forall).

Proof of Lemma 53 (Corresp. Values \forall). By Lemma 55 one of the following is true:

Case 1: $\Gamma \vdash_{!} p: \forall(x: T) U$. By Lemma 61, there exists a term s such that $\gamma \vdash p \rightsquigarrow s$. Using Lemma 57 we can infer that either

- s is a function, in which case we are done, or
- p has a recursive or singleton type under precise-I typing which is incompatible with p 's precise-I function type, a contradiction.

Case 2: There exists a path q such that $\Gamma \vdash_{!!!} p: q.\text{type}$ and $\Gamma \vdash_{!} q: \forall(x: T) U$. By Lemma 68,

$$\gamma \vdash p \rightsquigarrow^* q,$$

and by Lemma 61,

$$\gamma \vdash q \rightsquigarrow s$$

for some stable term s . According to Lemma 57, s must be a function $\lambda(x: T) u$. Therefore, $\gamma \vdash p \rightsquigarrow^* \lambda(x: T) u$.

□

$$\frac{\text{inert } \Gamma \quad \gamma: \Gamma \vdash \Gamma \Diamond \quad \gamma \vdash s \rightsquigarrow^* s' \quad \Gamma \vdash s: \forall(x: T) U}{\Gamma \vdash s': \forall(x: T) U} \quad (\text{CORRESP. TYPES } \forall)$$

Finally, we outline the proof of Lemma 54.

Proof of Lemma 54 (Corresp. Types \forall). The proof is by induction on the reflexive, transitive closure of store lookup; the reflexive case is vacu-

ously true. In the inductive case, we have

$$\gamma \vdash p \rightsquigarrow s_1 \quad (21.70)$$

$$\gamma \vdash s_1 \rightsquigarrow^* s' \quad (21.71)$$

$$\Gamma \vdash p : \forall(x : T) U \quad (21.72)$$

We would like to apply the induction hypothesis to (21.71), which would conclude the proof; however, for that we need to show that s' has the same type as p .

Applying the proof recipe to (21.72), we obtain that

$$\Gamma \vdash_{!!!} p : \forall(x : T') U'$$

where $\Gamma \vdash_{\#} T <: T'$ and $\Gamma, x : T \vdash_{\#} U' <: U$. By Lemma 59 applied to (21.70) and (21.72), s_1 must have the same type as p 's precise type:

$$\Gamma \vdash s_1 : \forall(x : T') U'$$

Since $\Gamma \vdash \forall(x : T') U' <: \forall(x : T) U$, by subsumption,

$$\Gamma \vdash s_1 : \forall(x : T) U$$

which is what we needed to show. \square

21.5 VALUE TYPING

Recall from Section 5.6 that the preservation lemma depends on Lemma 20 (Value Typing) according to which any well-typed value has an inert precise type. In pDOT, the lemma needs to additionally state that the value's precise type maintains the typed-paths property of an environment:

LEMMA 69 (Value Typing). *If Γ is an **inert, typed-paths** environment and $\Gamma \vdash v : T$, then there exists an inert type T' such that*

- $\Gamma \vdash_{!} v : T'$,
- $\Gamma \vdash T' <: T$, and
- $\vdash \Gamma, x : T' \Diamond$ where $x \notin \text{dom}(\Gamma)$

$$\frac{\Gamma \vdash v : T \quad \text{inert } \Gamma \vdash \Gamma \Diamond}{\Gamma \vdash_{!} v : T' \quad \text{inert } T' \quad \vdash \Gamma, x : T' \Diamond \quad \Gamma \vdash T' <: T} \text{ (VALUE TYPING)}$$

The proof of pDOT's version of value typing is more complicated than in DOT. We first introduce a *type-lookup* relation and a few auxiliary lemmas.

To prove Value Typing for pDOT we need to be able to look inside of deeply nested field declarations of types. For that we introduce a *type-lookup* relation on a path p and two types T, U , denoted $T \xrightarrow{p} U$, which allows us to follow the path p inside of the recursive type T to

$$\emptyset \xRightarrow{x} \emptyset \quad (\text{TL-EMPTY})$$

$$\frac{T \xRightarrow{x.\bar{b}} \mu(y: \dots \wedge \{a: S\} \wedge \dots)}{T \xRightarrow{x.\bar{b}.a} S[x.\bar{b}/y]} \quad (\text{TL})$$

Figure 21.6: The type lookup relation

yield the type U . For example, $\mu(x: \{a: \mu(y: \{b: \top\})\}) \xRightarrow{x.a.b} \top$. The definition of type lookup is presented in Figure 21.6.

The following lemma connects precise-I typing of a path $x.b_1 \dots .b_n$ with type lookup in x 's environment type $\Gamma(x)$. Specifically, it gives one access to the environment type that is assigned to b_1 and to b_n .

LEMMA 70. *Let $\Gamma, x: \mu(x: T)$ be an inert environment, and*

$$\Gamma, x: \mu(x: T) \vdash_! x.\bar{b}: W$$

where $\bar{b} = b_1 \dots .b_n$ and $n > 0$. Then there exists a type V such that

- $T = \dots \wedge \{b_1: V\} \wedge \dots$ and
- $\mu(x: T) \xRightarrow{x.b_1 \dots .b_{n-1}} \mu(z: \dots \wedge \{b_n: W\} \wedge \dots)$

Proof. By induction on precise-I typing. □

Suppose that the path p corresponds to an object that contains a record d of type $\{b: V\}$, and that $V = \dots \wedge \{a: q.\text{type}\} \wedge \dots$. Then q is well-typed.

LEMMA 71. *Suppose that for an inert type S ,*

- $p; \Gamma \vdash d: \{b: V\}$,
- $S \xRightarrow{p} \mu(y: T)$,
- $T[p/y] = \dots \wedge \{b: V\} \wedge \dots$,
- $S \xRightarrow{p.b.\bar{b}} \mu(z: U)$, and
- $U[p.b.\bar{b}/z] = \dots \wedge \{a: q.\text{type}\} \wedge \dots$

Then q is typeable in Γ .

Proof. The proof is done by a case analysis on d 's typing derivation. It requires reasoning about properties of type lookup which we omit here. □

The same generalizes to the case where a path p refers to an object with multiple declarations d whose record type contains $\{b: V\}$.

LEMMA 72. *Suppose that for an inert type S ,*

- $p; \Gamma \vdash d: \dots \wedge \{b: V\} \wedge \dots$,

- $S \xRightarrow{p} \mu(y: T)$,
- $T[p/y] = \dots \wedge \{b: V\} \wedge \dots$,
- $S \xRightarrow{p.b.\bar{b}} \mu(z: U)$, and
- $U[p.b.\bar{b}/z] = \dots \wedge \{a: q.\text{type}\} \wedge \dots$

Then q is typeable in Γ .

Proof. The proof is by induction on the definition-typing derivation of d and uses Lemma 71. \square

We can now outline the proof of Lemma 69 (Value Typing).

Proof of Lemma 69 (Value Typing). The proof is by induction on the typing of v . The interesting case is the object typing rule $\{\}$ -I (see Figure 19.2). We have $v = v(x: U)d$, $T = \mu(x: U)$, and

$$x; \Gamma, x: U \vdash d: U \quad (21.73)$$

We need to prove that there exists an inert type T' that is a subtype of $\mu(x: U)$, such that $\Gamma \vdash! v(x: U)d: T'$ and such that $\vdash \Gamma, x: T' \Diamond$. We choose $\mu(x: U)$ as this type. The challenge is to show that $\vdash \Gamma, x: \mu(x: U) \Diamond$. That is, we must show that if a path whose receiver is x has a singleton type $q.\text{type}$ then q must be precise-II typeable in the extended environment $\Gamma, x: T$ (see rule W_T in Figure 21.5). To prove that, suppose that

$$\Gamma, x: \mu(x: U) \vdash! x.\bar{b}: q.\text{type} \quad (21.74)$$

We first prove that q has a *general* type in the extended environment. We must show that there exists a type S such that $\Gamma, x: \mu(x: U) \vdash!! q: S$.

Suppose that $\bar{b} = b_1 \dots b_n$. If $n = 0$ then from (21.74) we can infer that $U = q.\text{type}$. However, then $\mu(x: U)$ is not inert, a contradiction.

Therefore, $n > 0$, and by Lemma 70 applied to (21.74), we have

$$\mu(x: U) \xRightarrow{x.b_1 \dots b_{n-1}} \mu(z: \dots \wedge \{b_n: q.\text{type}\} \wedge \dots) \quad (21.75)$$

Then by Lemma 72 applied to (21.74) and (21.75) (we choose $\mu(x: U)$ as S and x as p , and we use the fact that $\mu(x: U) \xRightarrow{x} \mu(x: U)$) we obtain that q has a general type. It is then easy to show using the proof recipe that if a path has a general type it also has a precise-II type, which concludes the proof. \square

$$\begin{array}{c} \Gamma \vdash v: T \\ \text{inert } \Gamma \quad \vdash \Gamma \Diamond \\ \hline \Gamma \vdash! v: T' \\ \text{inert } T' \\ \vdash \Gamma, x: T' \Diamond \\ \Gamma \vdash T' <: T \\ \text{(VALUE TYPING)} \end{array}$$

21.6 TYPE SOUNDNESS FOR PDOT

To formulate the soundness theorems we update the definition of normal forms for pDOT:

$p \rightarrow \quad v \rightarrow$

DEFINITION 73 (Normal Form). A term t is in normal form, denoted $t \rightarrow$, if t is either a path or a value.

The two central lemmas of the soundness proof are Progress and Preservation:

$$\frac{\gamma: \Gamma \text{ inert } \Gamma \vdash \Gamma \Diamond \quad \Gamma \vdash t: T}{\gamma \mid t \mapsto \gamma' \mid t'}$$

$\vee t \rightarrow$
(PROGRESS)

LEMMA 74 (Progress). Let γ be a store and Γ a typing environment. If all of the following hold:

- $\gamma: \Gamma$,
- Γ is inert,
- Γ is a typed-paths environment, and
- $\Gamma \vdash t: T$,

then t is in normal form or there exists a term t' and a store γ' such that $\gamma \mid t \mapsto \gamma' \mid t'$.

$$\frac{\gamma: \Gamma \text{ inert } \Gamma \vdash \Gamma \Diamond \quad \Gamma \vdash t: T \quad \gamma \mid t \mapsto \gamma' \mid t' \quad \Gamma' \vdash t': T}{\gamma': \Gamma' \text{ inert } \Gamma' \vdash \Gamma' \Diamond}$$

$\gamma': \Gamma' \text{ inert } \Gamma' \vdash \Gamma' \Diamond$
(PRESERVATION)

LEMMA 75 (Preservation). Let γ be a store and Γ a typing environment. If all of the following hold:

- $\gamma: \Gamma$,
- Γ is inert,
- Γ is a typed-paths environment,
- $\Gamma \vdash t: T$, and
- $\gamma \mid t \mapsto \gamma' \mid t'$,

then there exists an inert, typed-paths environment Γ' such that $\gamma': \Gamma'$ and $\Gamma' \vdash t': T$.

With the canonical forms and value typing lemmas in place, the proofs of the above two lemmas mirror the proofs of the same lemmas of the simple soundness proof for DOT (Section 5.6).

Progress and Preservation allow us to easily prove type safety which ensures that any well-typed pDOT program does not get stuck, i.e. it either diverges or reduces to a normal form (a path or a value):

$$\frac{\vdash t: T}{(\emptyset \mid t \mapsto^* \gamma \mid t' \wedge t' \rightarrow) \vee t \uparrow}$$

(pDOT SOUNDNESS)

THEOREM 76 (pDOT Type Soundness). If $\vdash t: T$ then either t diverges ($t \uparrow$), or t reduces to a normal form s , i.e. $\emptyset \mid t \mapsto^* \gamma \mid s$ and $\Gamma \vdash s: T$ for some Γ such that $\gamma: \Gamma$.

Since evaluating pDOT programs can result in paths (which are normal form), one might ask whether looking up those paths yields anything meaningful. As mentioned in Section 18.2.3, looking up any well-typed path in the runtime environment results either in a value or an infinite loop. To formulate the final soundness theorem that reasons about both term reduction and path lookup we define the following extended reduction relation \rightarrow :

$$\frac{\gamma \mid t \mapsto \gamma' \mid t'}{\gamma \mid t \twoheadrightarrow \gamma' \mid t'} \qquad \frac{\gamma \vdash s \rightsquigarrow s'}{\gamma \mid s \twoheadrightarrow \gamma \mid s'}$$

We denote the reflexive, transitive closure of extended reduction as \twoheadrightarrow^* .

DEFINITION 77 (Extended Divergence). *A term t extendedly diverges, denoted $t \Uparrow\Uparrow$, if there exists an infinite extended-reduction sequence*

$$\emptyset \mid t \twoheadrightarrow \gamma_1 \mid t_1 \twoheadrightarrow \dots \twoheadrightarrow \gamma_n \mid t_n \twoheadrightarrow \dots$$

Finally, we state the following extended soundness theorem:

THEOREM 78 (Extended Type Soundness). *If $\vdash t : T$ then either t extendedly diverges ($t \Uparrow\Uparrow$), or t reduces to a value, i.e. $\emptyset \mid t \twoheadrightarrow^* \gamma \mid v$.*

A diagram with some of the pDOT-related changes to the DOT soundness proof is presented in Figure 21.7.

$$\frac{\vdash t : T}{\emptyset \mid t \twoheadrightarrow \gamma \mid v} \quad \vee t \Uparrow\Uparrow$$

(EXTENDED SOUNDNESS)

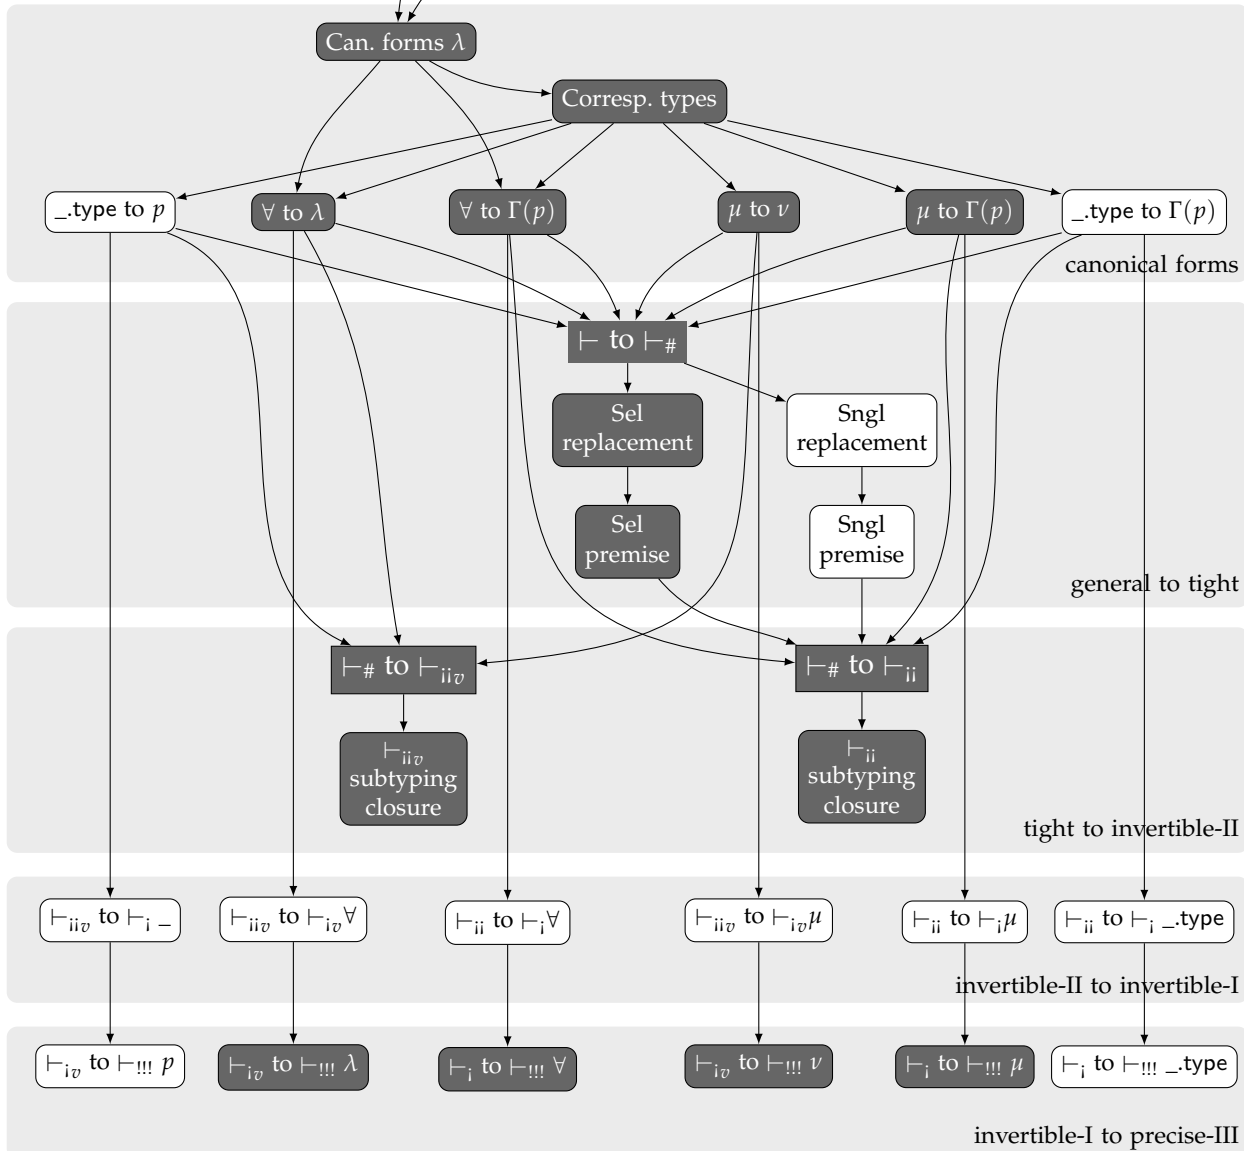


Figure 21.7: An instance of the dependency graph from Figure 5.4 showing the main lemmas in the pDOT proof as an extension of the simple DOT proof (Part I). Gray nodes denote the pDOT lemmas that have similar analogues in the DOT proof. White nodes denote pDOT specific lemmas. We omit the Precise-II and Precise-I related lemmas as well as additional lemmas required to prove Corresponding Types and the conversions between tight and invertible-II, and between invertible-II and invertible-I typing.

RELATED WORK

This section reviews the work related to formalizing Scala with support for fully path-dependent types.

22.1 EARLY CLASS-BASED SCALA FORMALIZATIONS

Several predecessors of the DOT calculus support path-dependent types on paths of arbitrary length. The first Scala formalization, νObj (Odersky, Cremet, et al., 2003), is a nominal, class-based calculus with a rich set of language features that formalizes object-dependent type members. Two subsequent calculi, Featherweight Scala (FS_{alg}) (Cremet et al., 2006) and Scalina (Moors, Piessens, and Odersky, 2008), build on νObj to establish Scala formalizations with algorithmic typing and with full support for higher-kinded types. All three calculi support paths of arbitrary length, singleton types, and abstract type members. Whereas FS_{alg} supports type-member selection directly on paths, νObj and Scalina allow selection $T\#A$ on types. A path-dependent type $p.A$ can thus be encoded as a selection on a singleton type: $p.\text{type}\#A$. νObj is the only of the above calculi that comes with a type-safety proof. The proof is non-mechanized.

Both pDOT and these calculi prevent type selections on non-terminating paths. νObj achieves this through a *contraction* requirement that prevents a term on the right-hand side of a definition from referring to the definition's self variable. At the same time, recursive calls can be encoded in νObj by referring to the self variable from a nested class definition. FS_{alg} ensures that paths are normalizing through a cycle detection mechanism that ensures that a field selection can appear only once as part of a path. Scalina avoids type selection $T\#A$ on a non-terminating type T by explicitly requiring T to be of a *concrete* kind, which means that T expands to a structural type R that contains a type member A . Although Scalina allows A to have upper and lower bounds, bad bounds are avoided because A also needs to be immediately initialized with a type U that conforms to A 's bounds, which is more restrictive than DOT. In pDOT, it is possible to create cyclic paths but impossible to do a type selection on them because as explained in Section 18.6, cyclic paths that can appear in a concrete execution context cannot be typed with a type-member declaration.

A difference between pDOT and the above calculi is that to ensure type soundness, paths in pDOT are normal form. This is necessary to ensure that each object has a name, as explained in Section 18.2.2. νObj and FS_{alg} achieve type safety in spite of reducing paths by allowing

field selection only on variables. This way, field selections always occur on named objects. Scalina does not require objects to be tied to names. In particular, its field selection rule E_Sel allows a field selection $new\ T.a$ on an object if T contains a field definition $\{a = s\}$. The selection reduces to $s\ [^{new\ T}/this]$, i.e. each occurrence of the self variable is replaced with a copy of $new\ T$.

A second difference to pDOT is the handling of singleton types. In order to reason about a singleton type $p.type$, νObj , FS_{alg} , and Scalina use several recursively defined judgments (membership, expansion, and others) that rely on analyzing the shape and well-formedness of the type that p expands to. By contrast, pDOT contains one simple Sngl-Trans rule that allows a path to inherit the type of its alias. On the other hand, pDOT has the shortcoming that singleton typing is not reflexive. Unlike in the above systems and in Scala, pDOT lacks a type axiom $\Gamma \vdash p : p.type$. Such a rule would undermine the anti-symmetry of path aliasing which is essential to the safety proof.

None of the other calculi have to confront the problem of bad bounds. Unlike DOT, pDOT, and Scala, νObj and FS_{alg} do not support lower bounds of type members and have no unique upper and lower bounds on types. Scalina does have top and bottom types and supports bounds through interval kinds, but it avoids bad bounds by requiring types on which selection occurs to be concrete. In addition, it is unknown whether Scalina and FS_{alg} are sound.

Finally, the three type systems are nominal and class based, and include a large set of language features that are present in Scala. DOT is a simpler and smaller calculus that abstracts over many of the design decisions of the above calculi. Since DOT aims to be a base for experimentation with new language features, it serves well as a minimal core calculus for languages with type members, and the goal of pDOT is to generalize DOT to fully path-dependent types.

22.2 DOT-LIKE CALCULI

Amin, Moors, and Odersky, (2012) present the first version of a DOT calculus. It includes type intersection, recursive types, unique top and bottom types, type members with upper and lower bounds, and path-dependent types on paths of arbitrary length. This version of DOT has explicit support for fields (vals) and methods (defs). Fields must be initialized to variables, which prevents the creation of non-terminating paths (since that would require initializing fields to paths), but it also limits expressivity. Specifically, just like in DOT by Amin, Grütter, et al., (2016), path-dependent types cannot refer to nested modules because modules have to be created through methods, and method invocations cannot be part of a path-dependent type. The calculus is not type-safe, as illustrated by multiple counterexamples in the paper.

In particular, this version of DOT does not track path equality which, as explained in the paper, breaks preservation.

To be type-safe, DOT must ensure that path-dependent types are invoked only on terminating paths. A possible strategy to ensure a sound DOT with support for paths is to investigate the conditions under which terms terminate, and to impose these conditions on the paths that participate in type selections. To address these questions, Wang and Rompf, (2017) present a Coq-mechanized semantic proof of strong normalization for the $D_{<}$ calculus. $D_{<}$ is a generalization of System $F_{<}$ with lower- and upper-bounded type tags and variable-dependent types. The paper shows that recursive objects constitute the feature that enables recursion and hence Turing-completeness in DOT. Since $D_{<}$ lacks recursive objects, it is strongly normalizing. Furthermore, the lack of objects and fields implies that this version of $D_{<}$ can only express paths that are variables.

Hong, Park, and Ryu, (2018) present π DOT, a strongly normalizing version of a $D_{<}$ without top and bottom types but with support for paths of arbitrary length. π DOT keeps track of path aliasing through path-equivalence sets, and the paper also mentions the possibility of using singleton types to formalize path equality. Like the calculus by Wang and Rompf, (2017), this version of $D_{<}$ is strongly normalizing due to the lack of recursive self variables. This guarantees that paths are acyclic. It also ensures that due to the lack of recursion elimination, reducing paths preserves soundness (unlike in pDOT, as explained in Section 18.2.2). π DOT comes with a non-mechanized soundness proof.

By contrast with these two papers, our work proposes a Turing-complete generalization with paths of arbitrary length of the full DOT calculus, which includes objects and type intersections.

22.2.1 Other Related Languages and Calculi

Scala’s module system shares many similarities with languages of the ML family. Earlier presentations of ML module systems Dreyer, Crary, and Harper, (2003), Harper and Lillibridge, (1994), and Leroy, (1994) allow fine-grained control over type abstractions and code reuse but do not support mutually recursive modules, and separate the language of terms from the language of modules. *MixML* extends the essential features of these type systems with the ability to do both *hierarchical* and *mixin* composition of modules (Rossberg and Dreyer, 2013). The language supports *recursive* modules which can be packaged as *first-class values*. The expressive power of *MixML*’s module system, plus support for decidable type-checking requires a set of constraints on the *linking* (module mixin) operation that restrict recursion between modules, including a total order on *label paths*, and yields a complex type system that closely models actual implementations of ML.

Rossberg, Russo, and Dreyer, (2014) and Rossberg, (2018) address the inherent complexity of ML module systems by presenting encodings of an ML-style module language into System F_ω . The latter paper presents *1ML*, a concise version of ML that fully unifies the language of modules with the language of terms. However, both formalizations exclude recursive modules.

A type system that distinguishes types based on the runtime values of their enclosing objects was first introduced by Ernst, (2001) in the context of *family polymorphism*. Notably, family polymorphism is supported by *virtual classes*, which can be inherited and overridden within different objects and whose concrete implementation is resolved at runtime. Virtual classes are supported in the *Beta* and *gbeta* programming languages Ernst, (1999) and Madsen and Møller-Pedersen, (1989) (but not in Scala in which classes are statically resolved at compile time) and formalized by the *vc* and *Tribe* calculi Clarke et al., (2007) and Ernst, Ostermann, and Cook, (2006). Paths in *vc* are relative to this and consist of a sequence of out keywords, which refer to enclosing objects, and field names. To track path equality, *vc* uses a normalization function that converts paths to a canonical representation, and to rule out cyclic paths it defines a partial order on declared names. *Tribe*'s paths can be both relative or absolute: they can start with a variable, and they can intermix class and object references. The calculus uses singleton types to track path equality and rules out cyclic paths by disallowing cyclic dependencies in its inheritance relation.

A difference between pDOT and all of *vc*, *Tribe*, and the ML formalizations is that pDOT does not impose any orderings on paths, and fully supports recursive references between objects and path-dependent types. In addition, pDOT's ability to define type members with both lower and upper bounds introduces a complex source of unsoundness in the form of *bad bounds* (alas, the cost for its expressiveness is that pDOT's type system is likely not decidable). Yet, by being mostly structurally typed, without having to model initialization and inheritance, pDOT remains general and small. Finally, by contrast to the above, pDOT comes with a mechanized type-safety proof.

22.2.2 Decidability

As discussed in Chapter 8, typechecking the DOT calculus is conjectured to be undecidable. The open question of decidability of DOT needs to be resolved before we can consider decidability of pDOT.

We believe that pDOT does not introduce additional sources of undecidability into DOT. One feature of pDOT that might call this into question is singleton types. In particular, Stone and Harper, (2006) study systems of singleton kinds that reason about types with non-trivial reduction rules, yet it remains decidable which types reduce to the same normal form. The singleton types of both Scala and pDOT

are much simpler and less expressive in that only assignment of an object between variables and paths is allowed, but the objects are not arbitrary terms and do not reduce. Thus, the Scala and pDOT singleton types only need to track sequences of assignments. Therefore, although decidability of pDOT is unknown because it is unknown for DOT, the singleton types that we add in pDOT are unlikely to affect decidability because they are significantly less expressive than the singleton types studied by Stone and Harper.

CONCLUSION

The DOT calculus was designed as a small core calculus to model Scala’s type system with a focus on path-dependent types. However, DOT can only model types that depend on variables, which significantly under-approximates the behaviour of Scala programs. Scala and, more generally, languages with type members need to rely on fully path-dependent types to encode the possible type dependencies in their module systems without restrictions. Until now, it was unclear whether combining the fundamental features of languages with path-dependent types, namely bounded abstract type members, intersections, recursive objects, and paths of arbitrary length is type-safe.

We propose pDOT, a calculus that generalizes DOT with support for paths of arbitrary length. The main insights of pDOT are to represent object identity through paths, to ensure that well-typed paths without cyclic aliasing always represent values, to track path equality with singleton types, and to eliminate type selections on cyclic paths through precise object typing. pDOT allows us to use the full potential of path-dependent types. pDOT comes with a type-safety proof and motivating examples for fully-path dependent types and singleton types that are mechanized in Coq.

BIBLIOGRAPHY

- Amin, Nada, Samuel Grütter, Martin Odersky, Tiark Rompf, and Sandro Stucki (2016). "The Essence of Dependent Object Types." In: *A List of Successes That Can Change the World - Essays Dedicated to Philip Wadler on the Occasion of His 60th Birthday*, pp. 249–272.
- Amin, Nada, Adriaan Moors, and Martin Odersky (2012). "Dependent Object Types." In: *International Workshop on Foundations of Object-Oriented Languages (FOOL 2012)*.
- Amin, Nada and Tiark Rompf (2017). "Type soundness proofs with definitional interpreters." In: *Proceedings of the 44th ACM SIGPLAN Symposium on Principles of Programming Languages, POPL 2017, Paris, France, January 18–20, 2017*, pp. 666–679. URL: <http://dl.acm.org/citation.cfm?id=3009866>.
- Amin, Nada, Tiark Rompf, and Martin Odersky (2014). "Foundations of path-dependent types." In: *Proceedings of the 2014 ACM International Conference on Object Oriented Programming Systems Languages & Applications, OOPSLA 2014, part of SPLASH 2014, Portland, OR, USA, October 20–24, 2014*, pp. 233–249.
- Amin, Nada and Ross Tate (2016). "Java and Scala's type systems are unsound: the existential crisis of null pointers." In: *Proceedings of the 2016 ACM SIGPLAN International Conference on Object-Oriented Programming, Systems, Languages, and Applications, OOPSLA 2016, part of SPLASH 2016, Amsterdam, The Netherlands, October 30 - November 4, 2016*, pp. 838–848.
- Appel, Andrew W. and Amy P. Felty (2000). "A Semantic Model of Types and Machine Instructions for Proof-Carrying Code." In: *POPL 2000, Proceedings of the 27th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, pp. 243–253.
- Bell, CJ, Robert Dockins, Aquinas Hobor, and David Walker (2008). "Comparing semantic and syntactic methods in mechanized proof frameworks." In: *International Workshop on Proof-Carrying Code (PCC)*. Citeseer.
- Bruce, Kim B., Martin Odersky, and Philip Wadler (1998). "A Statically Safe Alternative to Virtual Types." In: *ECOOP'98 - Object-Oriented Programming, 12th European Conference*, pp. 523–549.
- Clarke, Dave, Sophia Drossopoulou, James Noble, and Tobias Wrigstad (2007). "Tribe: a simple virtual class calculus." In: *Proceedings of the 6th International Conference on Aspect-Oriented Software Development, AOSD 2007, Vancouver, British Columbia, Canada, March 12–16, 2007*, pp. 121–134.
- Cremet, Vincent, Francois Garillot, Sergueï Lenglet, and Martin Odersky (2006). "A Core Calculus for Scala Type Checking." In: *Mathe-*

- mathematical Foundations of Computer Science, 31st International Symposium, Slovakia.*
- Documentation, Dotty (2018a). *Intersection Types*. URL: <https://dotty.epfl.ch/docs/reference/intersection-types.html>.
- Documentation, Scala (2018b). *Paths*. URL: <https://www.scala-lang.org/files/archive/spec/2.11/03-types.html#paths>.
- Dreyer, Derek, Karl Cray, and Robert Harper (2003). “A type system for higher-order modules.” In: *Conference Record of POPL 2003: The 30th SIGPLAN-SIGACT Symposium on Principles of Programming Languages, New Orleans, Louisiana, USA, January 15-17, 2003*, pp. 236–249.
- Ernst, Erik (1999). “gbeta – a Language with Virtual Attributes, Block Structure, and Propagating, Dynamic Inheritance.” PhD thesis. Department of Computer Science, University of Aarhus, Århus, Denmark.
- (2001). “Family Polymorphism.” In: *ECOOP 2001 - Object-Oriented Programming, 15th European Conference, Budapest, Hungary, June 18-22, 2001, Proceedings*, pp. 303–326.
- Ernst, Erik, Klaus Ostermann, and William R. Cook (2006). “A virtual class calculus.” In: *Proceedings of the 33rd ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL 2006, Charleston, South Carolina, USA, January 11-13, 2006*, pp. 270–282.
- Giarrusso, Paolo, Léo Stefanescu, Amin Timany, and Lars Birkedal (2019). *Towards Semantic Type Soundness for Dependent Object Types and Scala with Logical Relations in Iris*. URL: <https://www.dropbox.com/s/ljn89labnf9g5kp/beamer-intro-double-delft-2019-02-13-03.13.pdf>.
- Harper, Robert and Mark Lillibridge (1994). “A Type-Theoretic Approach to Higher-Order Modules with Sharing.” In: *Conference Record of POPL’94: 21st ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, Portland, Oregon, USA, January 17-21, 1994*, pp. 123–137.
- Hong, Jaemin, Jihyeok Park, and Sukyoung Ryu (2018). “Path Dependent Types with Path-equality.” In: *Proceedings of the 9th ACM SIGPLAN International Symposium on Scala*. Scala 2018. St. Louis, MO, USA, pp. 35–39. ISBN: 978-1-4503-5836-1.
- Hu, Jason and Ondřej Lhoták (2019). “Undecidability of $D_{<}$ and Its Decidable Fragments.” In: *CoRR abs/1908.05294*. URL: <http://arxiv.org/abs/1908.05294>.
- Hughes, John (1989). “Why Functional Programming Matters.” In: *Comput. J.* 32.2, pp. 98–107.
- Kabir, Ifaz and Ondřej Lhoták (2018). “ κ DOT: scaling DOT with mutation and constructors.” In: *Proceedings of the 9th ACM SIGPLAN International Symposium on Scala*. ACM, pp. 40–50.

- Leroy, Xavier (1994). “Manifest Types, Modules, and Separate Compilation.” In: *Conference Record of POPL’94: 21st ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, Portland, Oregon, USA, January 17-21, 1994*, pp. 109–122.
- Mackay, Julian, Hannes Mehnert, Alex Potanin, Lindsay Groves, and Nicholas Robert Cameron (2012). “Encoding Featherweight Java with assignment and immutability using the Coq proof assistant.” In: *FTfFP 2012*, pp. 11–19.
- Madsen, Ole Lehrmann and Birger Møller-Pedersen (1989). “Virtual Classes: A Powerful Mechanism in Object-Oriented Programming.” In: *Conference on Object-Oriented Programming: Systems, Languages, and Applications (OOPSLA’89), New Orleans, Louisiana, USA, October 1-6, 1989, Proceedings*. Pp. 397–406.
- Moors, Adriaan, Frank Piessens, and Martin Odersky (2008). “Safe type-level abstraction in Scala.” In: *International Workshop on Foundations of Object-Oriented Languages (FOOL 2008)*.
- Nieto, Abel (2017). “Towards algorithmic typing for DOT (short paper).” In: *Proceedings of the 8th ACM SIGPLAN International Symposium on Scala, SCALA@SPLASH 2017, Vancouver, BC, Canada, October 22-23, 2017*, pp. 2–7.
- (2018). *Scala with explicit nulls*. URL: <https://github.com/abeln/dotty/wiki/scala-with-explicit-nulls>.
- Norris, Rob (2015). *Returning the “Current” Type in Scala*. URL: <https://tpolecat.github.io/2015/04/29/f-bounds.html>.
- Odersky, Martin (2016). *Type projection is unsound #1050*. URL: <https://github.com/lampepfl/dotty/issues/1050>.
- Odersky, Martin, Vincent Cremet, Christine Röckl, and Matthias Zenger (2003). “A Nominal Theory of Objects with Dependent Types.” In: *ECOOP 2003 - Object-Oriented Programming, 17th European Conference, Darmstadt, Germany, July 21-25, 2003, Proceedings*, pp. 201–224.
- Odersky, Martin, Guillaume Martres, and Dmitry Petrashko (2016). “Implementing higher-kinded types in Dotty.” In: *Proceedings of the 7th ACM SIGPLAN Symposium on Scala, SCALA@SPLASH 2016*, pp. 51–60.
- Odersky, Martin and Matthias Zenger (2005a). “Independently extensible solutions to the expression problem.” In: *FOOL 2005*. Vol. 12.
- (2005b). “Scalable component abstractions.” In: *Proceedings of the 20th Annual ACM SIGPLAN Conference on Object-Oriented Programming, Systems, Languages, and Applications, OOPSLA 2005, October 16-20, 2005, San Diego, CA, USA*, pp. 41–57.
- Osvald, Leo, Grégory M. Essertel, Xilun Wu, Lilliam I. González Alayón, and Tiark Rompf (2016). “Gentrification gone too far? affordable 2nd-class values for fun and (co-)effect.” In: *Proceedings of the 2016 ACM SIGPLAN International Conference on Object-Oriented Programming, Systems, Languages, and Applications, OOPSLA 2016*,

- part of *SPLASH 2016, Amsterdam, The Netherlands, October 30 - November 4, 2016*, pp. 234–251.
- Pierce, Benjamin C. (1992a). “Bounded Quantification is Undecidable.” In: *Proceedings of the 19th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*. POPL ’92. Albuquerque, New Mexico, USA: ACM, pp. 305–315. ISBN: 0-89791-453-8.
- (1992b). “Bounded Quantification is Undecidable.” In: *Conference Record of the Nineteenth Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, Albuquerque, New Mexico, USA, January 19-22, 1992*, pp. 305–315.
- (2002). *Types and programming languages*. MIT Press. ISBN: 978-0-262-16209-8.
- Rapoport, Marianna, Ifaz Kabir, Paul He, and Ondřej Lhoták (2017). “A simple soundness proof for dependent object types.” In: *PACMPL 1.OOPSLA*, 46:1–46:27.
- Rapoport, Marianna and Ondřej Lhoták (2016). “Mutable WadlerFest DOT.” In: *CoRR abs/1611.07610*. arXiv: [1611.07610](https://arxiv.org/abs/1611.07610). URL: <http://arxiv.org/abs/1611.07610>.
- Rapoport, Marianna and Ondřej Lhoták (2017). “Mutable WadlerFest DOT.” In: *FTfJP 2017*, 7:1–7:6.
- Reynolds, John C (1998). “Definitional interpreters for higher-order programming languages.” In: *Higher-order and symbolic computation* 11.4, pp. 363–397.
- Rompf, Tiark and Nada Amin (2015). “From F to DOT: Type Soundness Proofs with Definitional Interpreters.” In: *CoRR abs/1510.05216v1*. URL: <http://arxiv.org/abs/1510.05216v1>.
- (2016a). “From F to DOT: Type Soundness Proofs with Definitional Interpreters.” In: *CoRR abs/1510.05216v2*. URL: <http://arxiv.org/abs/1510.05216v2>.
- (2016b). “Type soundness for dependent object types (DOT).” In: *Proceedings of the 2016 ACM SIGPLAN International Conference on Object-Oriented Programming, Systems, Languages, and Applications, OOPSLA 2016, part of SPLASH 2016, Amsterdam, The Netherlands, October 30 - November 4, 2016*, pp. 624–641.
- Rossberg, Andreas (2018). “1ML — Core and modules united.” In: *Journal of Functional Programming* 28, e22.
- Rossberg, Andreas and Derek Dreyer (2013). “Mixin’ Up the ML Module System.” In: *ACM Trans. Program. Lang. Syst.* 35.1, 2:1–2:84.
- Rossberg, Andreas, Claudio V. Russo, and Derek Dreyer (2014). “F-ing modules.” In: *J. Funct. Program.* 24.5, pp. 529–607.
- Scalas, Alceste and Nobuko Yoshida (2016). “Lightweight Session Programming in Scala.” In: *30th European Conference on Object-Oriented Programming, ECOOP 2016, July 18-22, 2016, Rome, Italy*, 21:1–21:28.
- Stone, Christopher A. and Robert Harper (2006). “Extensional equivalence and singleton types.” In: *ACM Trans. Comput. Log.* 7.4, pp. 676–722.

- Wang, Fei and Tiark Rompf (2017). "Towards Strong Normalization for Dependent Object Types (DOT)." In: *31st European Conference on Object-Oriented Programming, ECOOP 2017, June 19-23, 2017, Barcelona, Spain*, 27:1–27:25.
- Wright, Andrew K. and Matthias Felleisen (1994). "A Syntactic Approach to Type Soundness." In: *Inf. Comput.* 115.1, pp. 38–94.

COLOPHON

This document was typeset using the typographical look-and-feel classicthesis developed by André Miede and Ivo Pletikosić. The style was inspired by Robert Bringhurst's seminal book on typography *"The Elements of Typographic Style"*.

Final Version as of October 9, 2019 Ÿ.