

A Path To DOT: Formalizing Fully-Path-Dependent Types

MARIANNA RAPOPORT, University of Waterloo, Canada

ONDŘEJ LHOTÁK, University of Waterloo, Canada


The Dependent Object Types (DOT) calculus aims to formalize the Scala programming language with a focus on *path-dependent types* — types such as $x.a_1 \dots a_n.T$ that depend on the runtime value of a *path* $x.a_1 \dots a_n$ to an object. Unfortunately, existing formulations of DOT can model only types of the form $x.A$ which depend on *variables* rather than general paths. This restriction makes it impossible to model nested module dependencies. Nesting small components inside larger ones is a necessary ingredient of a modular, scalable language. DOT’s variable restriction thus undermines its ability to fully formalize a variety of programming-language features including Scala’s module system, family polymorphism, and covariant specialization.

This paper presents the pDOT calculus, which generalizes DOT to support types that depend on paths of arbitrary length, as well as singleton types to track path equality. We show that naive approaches to add paths to DOT make it inherently unsound, and present necessary conditions for such a calculus to be sound. We discuss the key changes necessary to adapt the techniques of the DOT soundness proofs so that they can be applied to pDOT. Our paper comes with a Coq-mechanized type-safety proof of pDOT. With support for paths of arbitrary length, pDOT can realize DOT’s full potential for formalizing Scala-like calculi.

Additional Key Words and Phrases: DOT, Scala, dependent types, paths

1 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 so that we can focus on the important details. For example, in Scala, we can define trees where the node type remains abstract.

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

If an object x has type `Tree`, then the *path-dependent type* $x.\text{Node}$ 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 concrete definition of `Node` within `RedBlackTree` we have used *aggregation*. If an object r is a `RedBlackTree`, then the path-dependent type $r.\text{root.Color}$ allows us to traverse the composition and access the `Color` type member.

To fulfill their full potential with respect to composition, observe that path-dependent types need to be in fact *dependent types*. For example, if we have apple and orange trees, we want to disallow mixing up their nodes.

```
val appleTree: Tree
val orangeTree: Tree
appleTree.add(orangeTree.root) // expected: appleTree.Node, actual: orangeTree.Node
```

Here, the type system considers `appleTree.Node` and `orangeTree.Node` to be distinct and incompatible types because they depend on the runtime values of different objects.

Furthermore, path-dependent types allow Scala to unify 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: Scala objects correspond to ML modules, classes to functors, and interfaces to signatures [Odersky and Zenger 2005].

The long struggle to formalize path-dependent types recently led to machine-verified soundness proofs for several variants of the Dependent Object Types (DOT) calculus [Amin et al. 2016; Amin and Rompf 2017; Rompf and Amin 2016]. 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. Insights from the proof have influenced the design of Scala 3 and helped uncover soundness bugs in Scala and Java [Amin and Tate 2016].

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 2005], in which modules of a program can be arbitrarily nested to form other, larger modules.

In this paper, we formalize and prove sound a generalization of the DOT calculus [Amin et al. 2016] 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 proof of Rapoport et al. [2017].

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 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, as shown on the right side of Figure 1 in DOT pseudocode (Section 2.1 presents the actual 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 the 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 define `Symbol` inside of the types module:

Scala:	DOT pseudocode:
<pre> package dotted { 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>	<pre> let dotted = new { val core = new { val types = new { type Type type TypeRef = Type & { val symb: core.symbols.Symbol } } } val symbols = new { type Symbol = { val tpe: core.types.Type } } } </pre>

Fig. 1. A simplified excerpt from the Dotty compiler in Scala. This code fragment cannot be expressed in DOT, as shown on the right side in DOT pseudocode

```

val types = new {
  type Type; type Symbol
  type TypeRef = Type & { val symb: this.Symbol }}

```

However, this transformation requires that all program components be arranged as global members of one top-level object. 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 = new { type Symbol = { val tpe: dotted.core.types.Type }} in
let dotted = new ...

```

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 (see Section 2.2 for details).

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 et al. 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.

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

Section 2.3.1: 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. As we show in Section 2.3.2, if we allow a type to depend on a path, rather than a 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 of cycles in paths, 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 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 referential transparency would suggest that because `t1` and `t2` are both initialized to the same expression, `t1` and `t2` should be the same, and therefore `t1.Node` and `t2.Node` should be the same. But then it would not be possible to 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.

This paper contributes the following:

- 1) The **pDOT** calculus, a generalization of DOT with **path-dependent types of arbitrary length** that lifts DOT's type-selection-on-variables restriction. Section 3 provides an intuition for pDOT's main ideas, and Section 4 presents the calculus in detail.
- 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 2005].
- 3) A **Coq-mechanized type soundness proof** of pDOT that is based on the simple soundness proof by Rapoport et al. [2017]. Our proof maintains the simple proof's modularity properties which makes it easy to extend pDOT with new features. We describe the proof in Section 5.
- 4) **Formalized examples**, presented in Section 6, that illustrate the expressive power of pDOT: 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.

The Coq proof of pDOT can be found under

<https://git.io/dotpaths>

$\mathbf{let} \ y = v(z) \ \{A = T\} \wedge$ $\quad \{a = \lambda(x: z.A) \ x\} \ \mathbf{in} \ \dots$	$\mathbf{val} \ y = \mathbf{new} \ \{ \mathbf{type} \ A = T$ $\quad \mathbf{val} \ a = (x: \mathbf{this}.A) \Rightarrow x \}$
--	---

The type of the field a depends on the type of the object containing it. In DOT, this is expressed using a recursive type. The type of our example object is $\mu(z: \{A: T..T\} \wedge \{a: \forall(x: z.A) z.A\})$.

Given the let binding above, what should be the type of $y.a$? In the recursive type of y , the type of the field a is $\forall(x: z.A) z.A$. However, because the self variable z is in scope only inside the object itself, the type $\forall(x: z.A) z.A$ does not make sense outside the object and cannot be used to type $y.a$. In the field selection, however, we have a name other than z for the object itself, the name y . Therefore, we can open the recursive type by replacing the self variable z with the external name of the object y , giving y the type $\{A: T..T\} \wedge \{a: \forall(x: y.A) y.A\}$. This is achieved by the recursion elimination typing rule. Now the path $y.a$ can have the type $\forall(x: y.A) y.A$. Notice that recursion elimination is possible only when we have a variable such as y as an external name for an object.

Just as we need to apply recursion elimination to the type of y before we can type a field selection $y.a$, we must also do the same before we can use a type-member selection $y.A$ (specifically, to conclude that $T <: y.A <: T$). The recursion elimination is necessary because the type T could also refer to the self variable z , and thus may not make any sense outside of the object. Recursion elimination replaces occurrences of z in the type T with the external name y , so that the resulting type is valid even outside the object. When we add path-dependent types to the calculus, an important consideration will be recursion elimination on paths rather than just variables.

2.2 Path Limitations: 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 :

$\mathbf{let} \ x = v(x)$ $\quad \{c = v(_) \ \{A = x.B\}\} \wedge$ $\quad \{B = \underline{x.c.A}\} \ \mathbf{in} \ \dots$	$\mathbf{val} \ x = \mathbf{new} \ \{ x \Rightarrow$ $\quad \mathbf{val} \ c: \ \{ \mathbf{type} \ A \} = \mathbf{new} \ \{ \mathbf{type} \ A = x.B \}$ $\quad \mathbf{type} \ B = x.c.A$ $\quad \}$
--	--

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.

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 $y.A$:

$\mathbf{let} \ y = \underline{x.c}$	\mathbf{in}	$\mathbf{val} \ y = \underline{x.c}$
$\mathbf{let} \ x = v(x)$		$\mathbf{val} \ x = \mathbf{new} \ \{ x \Rightarrow$
$\quad \{c = v(_) \ \{A = x.B\}\} \wedge$		$\quad \mathbf{val} \ c: \ \{ \mathbf{type} \ A \} = \mathbf{new} \ \{ \mathbf{type} \ A = x.B \}$
$\quad \{B = y.A\}$	$\mathbf{in} \ \dots$	$\quad \mathbf{type} \ B = y.A$
		$\quad \}$

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.

2.3 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 key challenge that makes it difficult to ensure soundness of a DOT calculus.

2.3.1 Bad Bounds. The issue of bad bounds has been discussed thoroughly in many of the previous papers about DOT, and we refer the reader to them for a more detailed presentation.

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 $\text{SEL-}<:_{\text{DOT}}$:

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

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 $S <: U$ 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: $\text{let } x = t \text{ in } u$, $\lambda(x: T) t$, or $v(x: T)d$. 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 when evaluation reaches a context with x in scope, x is bound to a value, and therefore x 's type does not introduce bad bounds.

2.3.2 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 (the Scala version cannot be typechecked):

$v(x : \{a : \{C : \forall(y : \top) \top .. \{c : \top\}\} \wedge \{b : \{c : \top\}\} \{a = x.a\} \wedge \{b = \lambda(y : \top) y\})$	<pre>new { lazy val a: {type C >: Any⇒Any <: {val c: Any}} = a lazy val b: {val c: Any} = (y: Any) ⇒ y }</pre>
--	--

Here, $x.b$ refers to a function $\lambda(y : \top) y$ of type $\forall(y : \top) \top$. According to this hypothetical type system, $\forall(y : \top) \top <: x.a.C <: \{c : \top\}$. 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.

3 MAIN IDEAS

In this Section, we outline the main ideas that have shaped our definition of pDOT. The pDOT calculus that implements these ideas in full detail will be presented in Section 4.

3.1 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:

$\text{let } x = v(y) \{a = v(z) \{B = U\}\} \text{ in } x.a$	<pre>val x = new { a = new { type B = U } } x.a</pre>
---	---

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 $\text{SEL-}<:_{DOT}$ and $<:-\text{SEL}_{DOT}$ rules mentioned in Section 2.3.1. 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:

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

However, before we can use this rule we need to also generalize the recursion elimination rule REC-E_{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:

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

By similar reasoning, we need to generalize all DOT variable-typing rules to path-typing rules.

3.2 Paths as Identifiers

A key design decision of pDOT is to let *paths represent object identity*. To that end, paths in pDOT are *irreducible*, because reducing paths would strip objects of their identity and break preservation.

3.2.1 Variables are Identifiers in DOT. In the DOT calculus by Amin 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, x would reduce to $v = v(y) \{a = v(z) \{B = U\}\}$. 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 REC-EDOT, $\Gamma \vdash x: \{a: \mu(z: \{B: U[x/y]..U[x/y]\})\}$. Does v also have that type? No!

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

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.

3.2.2 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, $x.a$ has the type $\{B: U[x.a/z]..U[x.a/z]\}$. 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 $\mu(z: \{B: U..U\})$.

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

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.

3.2.3 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 terminates, or p is a path that cyclically aliases other paths. Additionally, as we will see in Section 5.2.3, 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.

When we make paths a normal form, we also have to generalize DOT's ANF syntax to use paths wherever DOT uses variables. For example, function application has the form $p q$ where p and q are paths rather than $x y$ where x and y are variables. As a result, all the DOT reduction and typing rules that operate on variables are generalized to paths in pDOT.

3.3 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 Section 1, t_3 aliases t_2 , but t_1 does not alias t_2 , even though they identify syntactically equal objects.

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. We will precisely define the replacement operation in Section 4.2.

3.4 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 [Rapoport et al. 2017].

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 $p.\mathbf{type}$, is a type that is inhabited *only* with the value that is represented by p . In the tree example from Section 1, to tell the type system that t_3 aliases t_2 , we ensure that $\mathbf{val} \ t_3 : t_2.\mathbf{type}$. This information is used to allow subtyping between aliased paths, and to allow such paths to be typed with the same types, as we will see in Section 4.2.

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 Section 6, we show a pDOT encoding of an example that motivates this Scala feature.

3.5 Distinguishing Fields and Methods

Scala distinguishes between fields (*vals*) and methods (*defs*), while 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 2018]. 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:

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

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:

$$\mathbf{let} \ a' = t_a \quad \mathbf{in} \\ \mathbf{let} \ x = v(x) \{a = a'\} \wedge \{b = \lambda(_). t_b\} \mathbf{in} \dots$$

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.

3.6 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 $v(x: \{A: T..T\}) \{A = T\}$, 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 $v(x: \{A: \top..\perp\}) \{A = T\}$, 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 have $\{A: T..T\}$ in its self type.
- (2) An object containing a definition $\{a = v(x: T)d\}$ must have $\{a: \mu(x: T)\}$ in its self type.
- (3) An object containing a definition $\{a = \lambda(x: T)U\}$ must have $\{a: \forall(x: S)V\}$ in its self type.
- (4) An object containing a definition $\{a = p\}$ must have $\{a: p.type\}$ 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

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

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

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

then x also has all of the following types:

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

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, the first two typings cannot be typed. In Section 4.2, we show one simple type rule that generalizes pDOT to open and abstract types of term members nested deeply inside an object. In Section 6, we encode several examples from previous DOT papers in pDOT and show that the real-world compiler example from Section 1 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.

4 FROM DOT TO PDOT

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 $\lambda(x: T)t$).

	Variable	$s :=$	Stable Term	$S, T, U :=$	Type
x, y, z				\top	top type
a, b, c	Term member	p	path	\perp	bottom type
A, B, C	Type member	v	value	$\{a: T\}$	field declaration
$p, q, r :=$	Path	$v(x: T)d$	Value	$\{A: S..T\}$	type declaration
x	variable	$\lambda(x: T) t$	object	$S \wedge T$	intersection
$p.a$	field selection	$\{a = s\}$	lambda	$\mu(x: T)$	recursive type
$t, u :=$	Term	$\{A = T\}$	Definition	$\forall(x: S) T$	depend. function
s	stable term	$d \wedge d'$	field definition	$p.A$	type projection
$p q$	application		type definition	$p.type$	singleton type
$\text{let } x = t \text{ in } u$	let binding		aggregate definition		

Fig. 2. Abstract syntax of pDOT

4.1 Syntax

Figure 2 shows the abstract syntax of pDOT which is based on the DOT calculus of Amin et al. [2016]. Differences from that calculus are indicated by shading.

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 $x y$ are done on paths: $p.a$ and $p q$. 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 3.4, the pDOT calculus formalizes Scala's singleton types. A singleton type $p.type$ 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 like a path-dependent type $p.A$ depends on the value of p , a singleton type $q.type$ depends on the value of q . Singleton types are therefore a second form of dependent types in the calculus.

4.2 pDOT Typing Rules

The typing and subtyping rules of pDOT are shown in Figures 3 and 4. The type system is based on the DOT of Amin et al. [2016], and all changes are highlighted in gray.

4.2.1 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 3.1. The key rules that make DOT and pDOT interesting are the type-selection rules $\leftarrow\text{-SEL}$ and $\text{SEL}\leftarrow$. 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 subtyping constraints $S <: p.A$ and $p.A <: U$. Thanks to these two rules, pDOT supports fully-path-dependent types.

4.2.2 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 3.6, the rules assign a precise self type for objects, ensuring that paths are

	$\Gamma ::= \emptyset \mid \Gamma, x : T$	Type environment
Term typing		
$\frac{\Gamma(x) = T}{\Gamma \vdash x : T} \quad (\text{VAR})$	$\frac{\Gamma \vdash p : \{a : T\}}{\Gamma \vdash p.a : T} \quad (\text{FLD-E})$	$\frac{\Gamma \vdash p : q.\text{type} \quad \Gamma \vdash q.a}{\Gamma \vdash p.a : q.a.\text{type}} \quad (\text{SINGL-E})$
$\frac{\Gamma, x : T \vdash t : U \quad x \notin \text{fv}(T)}{\Gamma \vdash \lambda(x : T) t : \forall(x : T) U} \quad (\text{ALL-I})$	$\frac{\Gamma \vdash p.a : T}{\Gamma \vdash p : \{a : T\}} \quad (\text{FLD-I})$	$\frac{\Gamma \vdash p : T [p/x]}{\Gamma \vdash p : \mu(x : T)} \quad (\text{REC-I})$
$\frac{\Gamma \vdash p : \forall(z : S) T \quad \Gamma \vdash q : S}{\Gamma \vdash p q : T [q/z]} \quad (\text{ALL-E})$	$\frac{\Gamma \vdash t : T}{\Gamma, x : T \vdash u : U \quad x \notin \text{fv}(U)} \quad (\text{LET})$	$\frac{\Gamma \vdash p : \mu(x : T)}{\Gamma \vdash p : T [p/x]} \quad (\text{REC-E})$
$\frac{x; \Gamma, x : T \vdash d : T}{\Gamma \vdash v(x : T) d : \mu(x : T)} \quad (\{\}-\text{I})$	$\frac{\Gamma \vdash p : q.\text{type} \quad \Gamma \vdash q : T}{\Gamma \vdash p : T} \quad (\text{SINGL-TRANS})$	$\frac{\Gamma \vdash p : T \quad \Gamma \vdash p : U}{\Gamma \vdash p : T \wedge U} \quad (\wedge\text{-I})$
		$\frac{\Gamma \vdash t : T \quad \Gamma \vdash T <: U}{\Gamma \vdash t : U} \quad (\text{SUB})$
Definition typing		
$p; \Gamma \vdash \{A = T\} : \{A : T..T\} \quad (\text{DEF-TYP})$		$\frac{\Gamma \vdash q}{p; \Gamma \vdash \{a = q\} : \{a : q.\text{type}\}} \quad (\text{DEF-PATH})$
$\frac{p; \Gamma \vdash \lambda(x : T) t : \forall(x : U) V}{p; \Gamma \vdash \{a = \lambda(x : T) t\} : \{a : \forall(x : U) V\}} \quad (\text{DEF-ALL})$		
$\frac{p.a; \Gamma \vdash d [p.a/y] : T [p.a/y] \quad \text{tight } T}{p; \Gamma \vdash \{a = v(y : T) d\} : \{a : \mu(y : T)\}} \quad (\text{DEF-NEW})$		$\frac{p; \Gamma \vdash d_1 : T_1 \quad p; \Gamma \vdash d_2 : T_2 \quad \text{dom}(d_1), \text{dom}(d_2) \text{ disjoint}}{p; \Gamma \vdash d_1 \wedge d_2 : T_1 \wedge T_2} \quad (\text{ANDDEF-I})$
Typeable paths		
$\frac{\Gamma \vdash p : T}{\Gamma \vdash p} \quad (\text{WF})$	Tight bounds	
	$\text{tight } T = \begin{cases} U = V & \text{if } T = \{A : U..V\} \\ \text{tight } U & \text{if } T = \mu(x : U) \text{ or } \{a : U\} \\ \text{tight } U \text{ and tight } V & \text{if } T = U \wedge V \\ \text{true} & \text{otherwise} \end{cases}$	

Fig. 3. pDOT typing rules

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.

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 $\{\}-\text{I}$ 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 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,

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

Fig. 4. pDOT subtyping rules

when typechecking the object assigned to $x.a$ in the expression

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

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

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

An alternative design of the DEF-NEW rule can be to introduce a fresh variable y into the context (similarly to the $\{\text{-I}$ rule). However, we would have to assign y the type $x.a.\text{type}$ 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.

4.2.3 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 denotes 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

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

the path $x.a$ should have the type $x.\text{type}$ 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, $\mu(x: \{a: x.\text{type}\} \wedge \{b: S\})$. Through applying recursion elimination (REC-E), field selection (FLD-E), and finally subsumption (SUB) with the intersection subtyping rule $\text{AND}_1\text{-<:}$, we will obtain that $\Gamma \vdash x.a: x.\text{type}$.

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 be also 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 $\Gamma \vdash x.a.b : x.b.type$.

Singleton Type Transitivity. 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 $\Gamma \vdash x.b : S$.

Equivalent Types. As described in Section 3.3, 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 $\Gamma \vdash p : q.type$, 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:

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

The pDOT subtyping rules `SNGLpq-<` and `SNGLqp-<`: allow us to assign these types to r by establishing subtyping between equivalent types. Specifically, if we know that $\Gamma \vdash p : q.type$ 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.

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 replacement operation implicitly defines a total order on type's paths. The statement $T [q/p, n] = U$ denotes that

- T is a type that contains at least n paths,
- the n th path in T contains p as a prefix,
- the result of replacing that occurrence of p in T with q yields the type U .

The precise definition of the replacement operation is presented in Figure 11 in the Appendix.

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

$$\frac{\dots \quad \Gamma \vdash r : \{a : p.type\} \wedge \{b : p.type\}}{\text{REC-E}} \quad \frac{\dots \quad \Gamma \vdash p : q.type \quad \frac{\{a : p.type\} \wedge \{b : p.type\} [q/p, 1] = \{a : p.type\} \wedge \{b : q.type\}}{\text{SNGL}_{pq}-<} \quad \text{REPL-AND}_2}{\Gamma \vdash \{a : p.type\} \wedge \{b : p.type\} <: \{a : p.type\} \wedge \{b : q.type\}} \quad \text{SUB}}{\Gamma \vdash r : \{a : p.type\} \wedge \{b : q.type\}}$$

To replace several occurrences of a path with another, we repeatedly apply `SNGLpq-<`: or `SNGLqp-<`:

4.2.4 Abstracting Over Field Types. Finally, we describe one of the most interesting pDOT rules that adds significant expressivity to pDOT.

Consider a function $f = \lambda(x : \{a : T\}) \dots$ and a path p that refers to the object $v(x : \{a : q.type\}) \{a = q\}$, where $\Gamma \vdash q : T$. Since $\Gamma \vdash p : \mu(x : \{a : q.type\})$, by `REC-E`, $\Gamma \vdash p : \{a : q.type\}$.³ 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 , $q.type$ is *not* a subtype of T , and therefore $\{a : q.type\}$ is not a subtype of $\{a : T\}$.

³Assume that q does not start with x .

$$\begin{array}{c}
\frac{\gamma(x) = v(x : T) \dots \{a = t\} \dots}{\gamma \mid x.a \mapsto \gamma \mid t} \text{ (PROJ}_{\text{DOT}}\text{)} \\
\frac{\gamma(x) = \lambda(z : T) t}{\gamma \mid xy \mapsto \gamma \mid t[y/z]} \text{ (APPLY}_{\text{DOT}}\text{)} \\
\gamma \mid \mathbf{let } x = y \mathbf{ in } t \mapsto \gamma \mid t[y/x] \text{ (LET-VAR}_{\text{DOT}}\text{)} \\
\gamma \mid \mathbf{let } x = v \mathbf{ in } t \mapsto \gamma, x \mapsto v \mid t \text{ (LET-VALUE}_{\text{DOT}}\text{)} \\
\frac{\gamma \mid t \mapsto \gamma' \mid t'}{\gamma \mid \mathbf{let } x = t \mathbf{ in } u \mapsto \gamma' \mid \mathbf{let } x = t' \mathbf{ in } u} \text{ (CTX}_{\text{DOT}}\text{)}
\end{array}
\qquad
\begin{array}{c}
\gamma ::= \emptyset \mid \gamma, x \mapsto v \quad \text{Store} \\
\frac{\gamma \vdash p \rightsquigarrow^* \lambda(z : T) t}{\gamma \mid pq \mapsto \gamma \mid t[q/z]} \text{ (APPLY)} \\
\gamma \mid \mathbf{let } x = p \mathbf{ in } t \mapsto \gamma \mid t[p/x] \text{ (LET-PATH)} \\
\gamma \mid \mathbf{let } x = v \mathbf{ in } t \mapsto \gamma, x \mapsto v \mid t \text{ (LET-VALUE)} \\
\frac{\gamma \mid t \mapsto \gamma' \mid t'}{\gamma \mid \mathbf{let } x = t \mathbf{ in } u \mapsto \gamma' \mid \mathbf{let } x = t' \mathbf{ in } u} \text{ (CTX)}
\end{array}$$

Fig. 5. Operational semantics of DOT and pDOT

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 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:

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

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 $\{a : \mu(x : T)\}$.

4.3 Reduction Semantics

The operational semantics of pDOT is presented in Figure 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 3.2.2);
- *path lookup*: pDOT uses the reflexive, transitive closure of the *path lookup* operation \rightsquigarrow that generalizes variable lookup in value environments to paths.

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

Fig. 6. Value-environment path lookup

The path lookup operation is presented in Figure 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 value environment 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 value environment: for example, if $\gamma \vdash p \rightsquigarrow q$ and $\gamma \vdash q \rightsquigarrow v$, we want to conclude that 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 $\lambda(z: \top) z$:

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

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 value environment 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 5.2.3, 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.

5 TYPE SAFETY

We implemented the type-safety proof of pDOT in Coq as an extension of the simple DOT soundness proof by Rapoport et al. [2017]. This section presents an overview of the key challenges and insights of proving pDOT sound.

Type safety 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):

THEOREM 5.1 (TYPE SOUNDNESS). *If $\vdash t: T$ then either t diverges, i.e. there exists an infinite reduction sequence $\emptyset \mid t \mapsto \gamma_1 \mid t_1 \mapsto \dots \mapsto \gamma_n \mid t_n \mapsto \dots$ starting with t , 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 3.2.3, 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 :

$$\frac{\gamma \mid t \mapsto \gamma' \mid t'}{\gamma \mid t \rightarrow \gamma' \mid t'} \qquad \frac{\gamma \vdash s \rightsquigarrow s'}{\gamma \mid s \rightarrow \gamma \mid s'}$$

We denote the reflexive, transitive closure of extended reduction as \rightarrow^* . Finally, we state the following extended soundness theorem:

THEOREM 5.2 (EXTENDED TYPE SOUNDNESS). *If $\vdash t : T$ then either t diverges, i.e. there exists an infinite reduction sequence $\emptyset \mid t \rightarrow \gamma_1 \mid t_1 \rightarrow \dots \rightarrow \gamma_n \mid t_n \rightarrow \dots$ starting with t , or t reduces to a value, i.e. $\emptyset \mid t \rightarrow^* \gamma \mid v$.*

Our proof follows the syntax-based approach to type soundness by Wright and Felleisen [1994]. The two central lemmas of the proof are Progress and Preservation:

LEMMA 5.3 (PROGRESS). *Let γ be a value environment and Γ a typing environment. If all of the following holds:*

- 1) $\gamma : \Gamma$ (i.e. if $\gamma = \overline{(x_i, v_i)}$ then $\Gamma = \overline{(x_i, T_i)}$ and $\Gamma \vdash v_i : T_i$),
- 2) Γ is inert (i.e. all types in Γ are the precise types of values, see Section 5.2.2),
- 3) Γ is well-formed (i.e. all paths in the types of Γ are typeable in Γ , see Section 5.2.1), and
- 4) $\Gamma \vdash t : T$,

then t is in normal form or there exists a term t' and a value environment γ' such that $\gamma \mid t \mapsto \gamma' \mid t'$.

LEMMA 5.4 (PRESERVATION). *Let γ be a value environment and Γ a typing environment. If all of the following holds:*

- 1) $\gamma : \Gamma$,
- 2) Γ is inert,
- 3) Γ is well-formed,
- 4) $\Gamma \vdash t : T$, and
- 5) $\gamma \mid t \mapsto \gamma' \mid t'$,

then there exists an inert, well-formed typing environment Γ' such that $\gamma' : \Gamma'$ and $\Gamma' \vdash t' : T$.

The pDOT proof follows the design principles laid out by Rapoport et al. [2017] of separating the reasoning about types, variables (paths), and values from each other to ensure modularity and facilitate future extensions of DOT.

5.1 Main Ideas of the DOT Safety Proof

The DOT type-safety proof addresses two main challenges:

- 1) *Rule out bad bounds:* Although bad bounds give rise to DOT typing contexts in which undesirable subtyping relationships hold, the proof needs to show that all reachable run-time states can be typed in well-behaved contexts.
- 2) *Induction on typing:* The DOT typing rules are very flexible, mirroring intuitive notions about which types a term ought to have. This flexibility requires rules that are opposites of each other and thus admit cycles in a derivation. The possibility of cycles impedes inductive reasoning about typing derivations.

The existing type safety proof defines a notion of *inert* types and typing contexts, a simple syntactic property of types that rules out bad bounds. Specifically, an inert type must be either a dependent function type or a recursive object type in which type members have equal bounds (i.e. $\{A : T..T\}$ rather than $\{A : S..U\}$). Crucially, the preservation lemma shows that reduction preserves

$$\text{inert } \forall(x: T) U \quad \frac{\text{record } T}{\text{inert } \mu(x: T)} \quad \text{record } \{A: T..T\} \quad \text{record } \{a: q.\text{type}\} \quad \frac{\text{inert } T}{\text{record } \{a: T\}}$$

Fig. 7. Inert Types in pDOT

inertness: that is, when $\gamma \mid t \mapsto \gamma' \mid t'$ there is an *inert* typing environment Γ' that corresponds to γ' and in which t' has the required type.

The proof also 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. In particular, 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.

5.2 Type Safety: From DOT to pDOT

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.

5.2.1 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, $S <: U$.

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, as shown in the inertness definition in Figure 7, 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.

Both the DOT and pDOT preservation lemmas must ensure that reduction preserves inertness of typing contexts, but pDOT also requires preservation of a second property, well-formedness of typing contexts. A pDOT type can depend on paths rather than only variables, and the type makes sense only if the paths within it make sense; more precisely, well-formedness 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.

5.2.2 Stratifying Typing Rules in pDOT. The language features that pDOT adds to DOT also 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.

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 . A further 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. 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 .⁴ 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 $\Gamma \vdash p_i: p_{i+1}.\text{type}$ or $\Gamma \vdash p_{i+1}: p_i.\text{type}$, 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 $\Gamma \vdash p_0: q.\text{type}$ and $\Gamma \vdash p_n: q.\text{type}$.⁵ Therefore, we further stratify the proof recipe into two typing judgments the first of which accounts for the $\text{SNG}_{pq}\text{-<}$ rule, and the second for the $\text{SNG}_{qp}\text{-<}$ 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 7 typing relations rather than the 4 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.

5.2.3 Canonical Forms in pDOT. Like many type soundness proofs, 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 are implied directly by the changes to the 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 $q.\text{type}$. But in order for p to have a function type, there cannot be a cycle of paths in

⁴An exceptional case is when p is declared to have type $q.\text{type}$ and q is declared to have type $p.\text{type}$. Fortunately, this case of a cycle turns out to be harmless because neither path is declared to have any other type other than its singleton type, and therefore neither path can be used in any interesting way.

⁵In degenerate cases, the middle path q might actually be p_0 or p_n .

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 5.5. *Let γ be a value environment and Γ be an inert, well-formed environment such that $\gamma : \Gamma$. If $\Gamma \vdash \underline{p} : \forall(x : T) U$ then there exists a type T' and a term t such that*

- 1) $\gamma \vdash \underline{p} \rightsquigarrow^* \lambda(x : T') t$,
- 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.

6 EXAMPLES

In this section, 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 $v(x \Rightarrow d')$:

- 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 definitions, 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 for 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

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

where $v(y \Rightarrow d')$ is the abbreviated version of $v(y : U)d$.

6.1 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 :

$v(p \Rightarrow \\ A = \mu(\mathbf{this} : \{\text{foo} : \forall(_) U\}); \\ \text{newA} : \forall(x : U) p.A \\ = v(\mathbf{this}) \{\text{foo} = \lambda(_) . x\})$	<pre> package p { class A(x : U) { def foo : U = x } } </pre>
---	--

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

$v(p \Rightarrow$ $B = p.A \wedge \{C: \perp..T\};$ $\text{newB}: \forall(x: U) p.B$ $= v(\mathbf{this}) \{ \text{foo} = \lambda(_).x; C = \dots \})$	<pre> package p { class B(x: U) extends A(x: U) { type C } } </pre>
--	---

6.2 Lists

As an example to illustrate that pDOT supports the type abstractions of DOT we formalize the covariant-list library by Amin et al. [2016] in pDOT, presented in Figure 8 a). 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 3.5. 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 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 $\perp <: \{A: T.. \perp\}$, we could derive $T <: \text{nil.head.A} <: \perp$.

6.3 Mutually Recursive Modules

The second example, presented in Figure 8 b), illustrates pDOT's ability to use path-dependent types of arbitrary length. It formalizes the compiler example from Section 1 in which the nested classes Type and Symbol recursively reference each other.

6.4 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 [2005]. An example from that paper introduces a class C with an incr method that increments a mutable integer field and returns the object itself (this). A class D extends C and defines an analogous decr method. By declaring the return types of the incr and decr methods as singleton types, we can invoke a chain of method calls $d.\text{incr}.\text{decr}$ on an object d of type D . If incr were declared with type C this would not have been possible since C does not have a decr member.

Our formalization of the example is displayed in Figure 8 c). Note that it excludes the mutation part of the original example since pDOT does not support mutation, and that side effect is orthogonal to the ability of chaining methods on objects of different types.

7 RELATED WORK

This section reviews the work related to formalizing Scala with support for fully path-dependent types. An detailed account on the history and comparison of DOT-related calculi and other Scala formalizations can be found in the recent DOT papers such as Amin et al. [2016]; Rapoport et al. [2017].

7.1 Early Class-based Scala Formalizations

Several predecessors of the DOT calculus support path-dependent types on paths of arbitrary length. The first Scala formalization, vObj [Odersky et al. 2003], is a nominal, class-based calculus with

a) A **covariant list** library in pDOT

```

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

```

b) Mutually recursive types in a compiler package: **fully-path-dependent types**

```

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

```

c) **Chaining** method calls using **singleton types**

```

let pkg = v(p ⇒ C = μ(this: {incr: this.type});
            D = μ(this: p.C ∧ {decr: this.type});
            newD: ∀(⊥) p.D = λ_.
              let result = v(this ⇒ incr = this; decr = this)
              in result)
in let d = pkg.newD _
in d.incr.decr

```

Fig. 8. Example pDOT encodings

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 et al. 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 to refer 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 3.6, 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 3.2.2. 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 $\text{new } T.a$ on an object if T contains a field definition $\{a = s\}$. The selection reduces to s [$\text{new } T/\text{this}$], i.e. each occurrence of the self variable is replaced with a copy of $\text{new } T$.

A second difference to pDOT is the handling of singleton types. In order to reason about a singleton type $p.\text{type}$, 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 $\Gamma \vdash p : p.\text{type}$. Such a rule would undermine the anti-symmetry of singleton types 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.

7.2 DOT-like Calculi

Amin et al. [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 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 et al. [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 3.2.2). 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.

8 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.

This paper proposes 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 acyclic paths 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.

ACKNOWLEDGMENTS

We would like to thank Martin Odersky for suggesting the idea of extending DOT with paths of arbitrary length, and for the helpful discussions on early variants of pDOT. We thank Yaoyu Zhao for his help with the Coq proof. We greatly appreciate Jaemin Hong’s proofreading of our paper; Jaemin provided a number of insightful questions and comments that helped us improve the presentation and fix multiple errors, and he pointed out an expressivity problem related to the absence of singleton reflexivity which we had not noticed. Abel Nieto gave us invaluable feedback and ideas for the writing and the presentation of the material; our discussions helped us formulate the definition of path replacement. Our conversations with Paolo Giarrusso improved our understanding of the the expressiveness and limitations of DOT; Paolo also gave us the idea to formulate an Extended Soundness Theorem. We had other helpful discussions on DOT with Ifaz Kabir, Sukyoung Ryu, Jason Zhu, Ilya Sergey, and Prabhakar Ragde. This research was supported by the Natural Sciences and Engineering Research Council of Canada.

REFERENCES

- Nada Amin, 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*. 249–272. https://doi.org/10.1007/978-3-319-30936-1_14
- Nada Amin, Adriaan Moors, and Martin Odersky. 2012. Dependent Object Types. In *International Workshop on Foundations of Object-Oriented Languages (FOOL 2012)*.
- Nada Amin 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*. 666–679. <http://dl.acm.org/citation.cfm?id=3009866>
- Nada Amin 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*. 838–848. <https://doi.org/10.1145/2983990.2984004>
- Kim B. Bruce, Martin Odersky, and Philip Wadler. 1998. A Statically Safe Alternative to Virtual Types. In *ECOOP’98 - Object-Oriented Programming, 12th European Conference*. 523–549.
- Vincent Cremet, François Garillot, Sergueï Lenglet, and Martin Odersky. 2006. A Core Calculus for Scala Type Checking. In *Mathematical Foundations of Computer Science, 31st International Symposium, Slovakia*.
- Scala Documentation. 2018. Paths. Retrieved February 26, 2019 from <https://www.scala-lang.org/files/archive/spec/2.11/03-types.html#paths>
- Erik Ernst. 2001. Family Polymorphism. In *ECOOP 2001 - Object-Oriented Programming, 15th European Conference, Budapest, Hungary, June 18–22, 2001, Proceedings*. 303–326. https://doi.org/10.1007/3-540-45337-7_17
- Jaemin Hong, 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)*. ACM, New York, NY, USA, 35–39. <https://doi.org/10.1145/3241653.3241657>
- Adriaan Moors, Frank Piessens, and Martin Odersky. 2008. Safe type-level abstraction in Scala. In *International Workshop on Foundations of Object-Oriented Languages (FOOL 2008)*.
- Martin Odersky, 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*. 201–224. https://doi.org/10.1007/978-3-540-45070-2_10
- Martin Odersky and Matthias Zenger. 2005. 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*. 41–57. <https://doi.org/10.1145/1094811.1094815>
- Marianna Rapoport, Ifaz Kabir, Paul He, and Ondřej Lhoták. 2017. A simple soundness proof for dependent object types. *PACMPL* 1, OOPSLA (2017), 46:1–46:27.
- Tiark Rompf and Nada Amin. 2016. 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*. 624–641. <https://doi.org/10.1145/2983990.2984008>
- Fei Wang 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.

x, y, z	Variable	$v :=$	Value	$S, T, U :=$	Type
a, b, c	Term member	$v(x: T)d$	object	\top	top type
A, B, C	Type member	$\lambda(x: T) t$	lambda	\perp	bottom type
$t, u :=$	Term	$d :=$	Definition	$\{a: T\}$	field declaration
x	variable	$\{a = t\}$	field definition	$\{A: S..T\}$	type decl.
$x.a$	field selection	$\{A = T\}$	type definition	$x.A$	type projection
v	value	$d \wedge d'$	aggregate definition	$S \wedge T$	intersection
$x y$	application			$\mu(x: T)$	recursive type
let $x = t$ in u	let binding			$\forall(x: S) T$	dependent fun.

Fig. 9. Abstract syntax of DOT by Amin et al. [2016]

Andrew K. Wright and Matthias Felleisen. 1994. A Syntactic Approach to Type Soundness. *Inf. Comput.* 115, 1 (1994), 38–94.

A APPENDIX

A.1 DOT by Amin et al. [2016]

In Figure 9, we present the abstract syntax, and in Figure 10, the typing rules of the DOT calculus by Amin et al. [2016].

A.2 Path Replacement

The definition of the path-replacement rules that define equivalent types is presented in Figure 11. In the figure, the notation $p.\bar{b}$ denotes a path with a possibly empty sequence of field selections, e.g. $p.b_1.b_2$. The path replacement operation depends on the recursively defined pathnum function that counts the number of paths in a type.

A.3 Proof Recipe of pDOT

In this section, we present an overview of the proof recipe (see Section 5.2.2) and the auxiliary typing judgments that are involved in it. The purpose of the proof recipe is to reason about the most precise inert type assigned to a path by the typing environment, given the path's general type. The proof recipe works by performing a sequence of type transformations starting with General typing and ending in Elimination typing.

The following diagram summarizes the typing relations that participate in the proof recipe:

$$\text{General } (\vdash) \rightarrow \text{Tight } (\vdash_{\#}) \rightarrow \text{Introduction-qp } (\vdash_{qp}) \rightarrow \text{Introduction-pq } (\vdash_{pq}) \rightarrow \text{Elim-III } (\vdash_{e3}) \rightarrow \text{Elim-II } (\vdash_{e2}) \rightarrow \text{Elim-I } (\vdash_{e1})$$

In addition, Table 1 shows which typing rules of pDOT are handled by each of the auxiliary typing relations. The exact definitions of the typing relations are presented in Figures 12 to 15. The proof recipe starts by translating general into *tight* typing, denoted $\Gamma \vdash_{\#} p: T$, which is equivalent to general typing but does not admit the introduction of new subtyping relationships. Tight typing is presented in Figure 12. The following lemma states that in an inert context, the general type of a term is the same as its tight type:

LEMMA A.1 (\vdash TO $\vdash_{\#}$). *If $\Gamma \vdash t: T$ and Γ is inert then $\Gamma \vdash_{\#} t: T$.*

Term typing

$$\begin{array}{c}
\frac{\Gamma(x) = T}{\Gamma \vdash x : T} \quad (\text{VAR}_{DOT}) \\
\frac{\Gamma, x : T \vdash t : U \quad x \notin \text{fv}(U)}{\Gamma \vdash \mathbf{let } x = t \mathbf{ in } u : U} \quad (\text{LET}_{DOT}) \\
\frac{\Gamma, x : T \vdash t : U \quad x \notin \text{fv}(T)}{\Gamma \vdash \lambda(x : T) t : \forall(x : T) U} \quad (\text{ALL-I}_{DOT}) \\
\frac{\Gamma \vdash x : \forall(z : S) T \quad \Gamma \vdash y : S}{\Gamma \vdash x y : T [y/z]} \quad (\text{ALL-E}_{DOT}) \\
\frac{\Gamma, x : T \vdash d : T}{\Gamma \vdash \nu(x : T) d : \mu(x : T)} \quad (\{\!\!\}\text{-I}_{DOT}) \\
\frac{\Gamma \vdash x : \{a : T\}}{\Gamma \vdash x.a : T} \quad (\text{FLD-E}_{DOT}) \\
\frac{\Gamma \vdash t : T}{\Gamma \vdash \mathbf{let } x = t \mathbf{ in } u : U} \quad (\text{LET}_{DOT}) \\
\frac{\Gamma \vdash x : T}{\Gamma \vdash x : \mu(x : T)} \quad (\text{REC-I}_{DOT}) \\
\frac{\Gamma \vdash x : \mu(z : T)}{\Gamma \vdash x : T [x/z]} \quad (\text{REC-E}_{DOT}) \\
\frac{\Gamma \vdash x : T \quad \Gamma \vdash x : U}{\Gamma \vdash x : T \wedge U} \quad (\text{AND-I}_{DOT}) \\
\frac{\Gamma \vdash t : T \quad \Gamma \vdash T <: U}{\Gamma \vdash t : U} \quad (\text{SUB}_{DOT})
\end{array}$$

Definition typing

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

Subtyping

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

Fig. 10. DOT Type Rules [Amin et al. 2016]

Each of the remaining proof-recipe typing relations comes in two versions: for paths and for values. We present only the versions for paths; the formulations of the lemmas for values are similar. After obtaining a tight type of a path, we can convert it into an Introduction- qp type:

LEMMA A.2 ($\vdash_{\#}$ TO \vdash_{qp}). *If $\Gamma \vdash_{\#} p : T$ and Γ is inert then $\Gamma \vdash_{qp} p : T$.*

$$\begin{array}{c}
\frac{}{p.\bar{b}.A[q/p, 0] = q.\bar{b}.A} \text{ (REPL-PATH)} \\
\frac{}{p.\bar{b}.\text{type}[q/p, 0] = q.\bar{b}.\text{type}} \text{ (REPL-SNGL)} \\
\frac{T_1[q/p, n] = T_2}{(T_1 \wedge U)[q/p, n] = T_2 \wedge U} \text{ (REPL-AND}_1\text{)} \\
\frac{T_1[q/p, n] = T_2 \quad \text{pathnum } U = m}{(U \wedge T_1)[q/p, m+n] = U \wedge T_2} \text{ (REPL-AND}_2\text{)} \\
\frac{T[q/p, n] = U}{\mu(x: T)[q/p, n] = \mu(x: U)} \text{ (REPL-REC)} \\
\frac{T_1[q/p, n] = T_2}{(\forall(x: T_1)U)[q/p, n] = \forall(x: T_2)U} \text{ (REPL-ALL}_1\text{)} \\
\frac{T_1[q/p, n] = T_2 \quad \text{pathnum } U = m}{(\forall(x: U)T_1)[q/p, m+n] = \forall(x: U)T_2} \text{ (REPL-ALL}_2\text{)} \\
\frac{T_1[q/p, n] = T_2}{\{a: T_1\}[q/p, n] = \{a: T_2\}} \text{ (REPL-FLD)} \\
\frac{T_1[q/p, n] = T_2}{\{A: T_1..U\}[q/p, n] = \{A: T_2..U\}} \text{ (REPL-TYP}_1\text{)} \\
\frac{T_1[q/p, n] = T_2 \quad \text{pathnum } U = m}{\{A: U..T_1\}[q/p, m+n] = \{A: U..T_2\}} \text{ (REPL-TYP}_2\text{)}
\end{array}$$

$$\text{pathnum } T = \begin{cases} 1 & \text{if } T = p.A \text{ or } p.\text{type} \\ \text{pathnum } U & \text{if } T = \{a: U\} \text{ or } \mu(x: U) \\ \text{pathnum } S + \text{pathnum } U & \text{if } T = \{A: S..U\}, \\ & S \wedge U, \text{ or } \forall(x: S)U \\ 0 & \text{otherwise} \end{cases}$$

Fig. 11. Replacement of the n th path p in a type by q

Relation	Type rules	Inlined subtyping rules (SUB+...)
Elimination	\vdash_{e1} VAR, FLD-E, REC-E	AND ₁ -<:, AND ₂ -<:
	\vdash_{e2} SNGL-E	
	\vdash_{e3} SNGL-TRANS	
Introduction	\vdash_{pq} &-I	SNGL _{pq} -<:, <:-AND, TOP, ALL-<:-ALL, FLD-<:-FLD, TYP-<:-TYP
	\vdash_{qp} $\hat{\&}$ -I, FLD-I, REC-I	SNGL _{qp} -<:, <:-AND, <:-SEL
Tight	$\vdash_{\#}$ all	all, tight versions of SEL rules
General	\vdash all	all

Table 1. Auxiliary typing relations that make up the proof recipe of pDOT

Introduction- qp typing, denoted $\Gamma \vdash_{qp} p: T$, contains introduction rules (such as recursion introduction) and inlined versions of the SNGL_{qp} subtyping rule (see Figure 13).

The remaining proof recipe lemmas are specialized to function and singleton types, and we will only present the ones for function types. The following lemma establishes that if a path p has a function type in Introduction- qp typing then it has the same type in Introduction- pq typing:⁶

⁶The versions of Lemma A.3 for singleton and object types are more complicated: the relationship between a path's Introduction- qp as Introduction- pq type will be equivalence (i.e. possible replacements of aliased paths) rather than equality.

Tight term typing

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

Tight subtyping

$$\begin{array}{c}
\Gamma \vdash_{\#} T <: \top \quad (\text{TOP}_{\#}) \\
\\
\Gamma \vdash_{\#} \perp <: T \quad (\text{BOT}_{\#}) \\
\\
\Gamma \vdash_{\#} T <: T \quad (\text{REFL}_{\#}) \\
\\
\frac{\Gamma \vdash_{\#} S <: T \quad \Gamma \vdash_{\#} T <: U}{\Gamma \vdash_{\#} S <: U} \quad (\text{TRANS}_{\#}) \\
\\
\Gamma \vdash_{\#} T \wedge U <: T \quad (\text{AND}_1\text{-<:}_{\#}) \\
\\
\Gamma \vdash_{\#} T \wedge U <: U \quad (\text{AND}_2\text{-<:}_{\#}) \\
\\
\frac{\Gamma \vdash_{\#} S <: T \quad \Gamma \vdash_{\#} S <: U}{\Gamma \vdash_{\#} S <: T \wedge U} \quad (\text{<:-AND}_{\#}) \\
\\
\frac{\Gamma \vdash_{e3} p : \{A : S..S\}}{\Gamma \vdash_{\#} S <: p.A} \quad (\text{<:-SEL}_{\#}) \\
\\
\frac{\Gamma \vdash_{e3} p : \{A : S..S\}}{\Gamma \vdash_{\#} p.A <: S} \quad (\text{SEL-<:}_{\#}) \\
\\
\frac{\Gamma \vdash_{e3} p : q.\text{type} \quad \Gamma \vdash_{e2} q}{\Gamma \vdash_{\#} T <: T [q/p, n]} \quad (\text{SINGL}_{pq}\text{-<:}_{\#}) \\
\\
\frac{\Gamma \vdash_{e3} p : q.\text{type} \quad \Gamma \vdash_{e2} q}{\Gamma \vdash_{\#} T <: T [p/q, n]} \quad (\text{SINGL}_{qp}\text{-<:}_{\#}) \\
\\
\frac{\Gamma \vdash_{\#} T <: U}{\Gamma \vdash_{\#} \{a : T\} <: \{a : U\}} \quad (\text{FLD-<:-FLD}_{\#}) \\
\\
\frac{\Gamma \vdash_{\#} S_2 <: S_1 \quad \Gamma \vdash_{\#} T_1 <: T_2}{\Gamma \vdash_{\#} \{A : S_1..T_1\} <: \{A : S_2..T_2\}} \quad (\text{TYP-<:-TYP}_{\#}) \\
\\
\frac{\Gamma \vdash_{\#} S_2 <: S_1 \quad \Gamma, x : S_2 \vdash T_1 <: T_2}{\Gamma \vdash_{\#} \forall(x : S_1) T_1 <: \forall(x : S_2) T_2} \quad (\text{ALL-<:-ALL}_{\#})
\end{array}$$

Fig. 12. Tight typing

Introduction- qp path typing

$$\begin{array}{c}
\frac{\Gamma \vdash_{pq} p: T}{\Gamma \vdash_{qp} p: T} \quad (\text{PATH}_{qp}) \\
\\
\frac{\Gamma \vdash_{qp} p: S \quad \Gamma \vdash_{qp} p: T}{\Gamma \vdash_{qp} p: S \wedge T} \quad (\text{AND}_{qp}) \\
\\
\frac{\Gamma \vdash_{qp} p: T [p/x]}{\Gamma \vdash_{qp} p: \mu(x: T)} \quad (\text{REC-I}_{qp}) \\
\\
\frac{\Gamma \vdash_{qp} p.a: T}{\Gamma \vdash_{qp} p: \{a: T\}} \quad (\text{FLD-I}_{qp}) \\
\\
\frac{\Gamma \vdash_{qp} p: T}{\Gamma \vdash_{e1} q: \{A: T..T\}} \quad (\text{TYP-SEL}_{qp}) \\
\\
\frac{\Gamma \vdash_{qp} r: \mu(x: T) \quad \Gamma \vdash_{e2} q}{\Gamma \vdash_{qp} r: \mu(x: T) [p/q, n]} \quad (\text{SNGL-REC}_{qp}) \\
\\
\frac{\Gamma \vdash_{qp} r: r'.A \quad \Gamma \vdash_{e2} q}{\Gamma \vdash_{qp} r: r'.A [p/q, n]} \quad (\text{SNGL-SEL}_{qp}) \\
\\
\frac{\Gamma \vdash_{qp} r: r'.\text{type} \quad \Gamma \vdash_{e2} q}{\Gamma \vdash_{qp} r: r'.\text{type} [p/q, n]} \quad (\text{SNGL-SNGL}_{qp})
\end{array}$$

Introduction $_{qp}$ value typing

$$\begin{array}{c}
\frac{\Gamma \vdash_{pq} v: T}{\Gamma \vdash_{qp} v: T} \quad (\text{PATH}_{qp-v}) \\
\\
\frac{\Gamma \vdash_{qp} v: S \quad \Gamma \vdash_{qp} p: T}{\Gamma \vdash_{qp} v: S \wedge T} \quad (\text{AND}_{qp-v}) \\
\\
\frac{\Gamma \vdash_{qp} v: T}{\Gamma \vdash_{qp} v: \mu(x: T)} \quad (\text{SNGL-REC}_{qp-v}) \\
\\
\frac{\Gamma \vdash_{qp} v: T [p/x]}{\Gamma \vdash_{qp} v: \mu(x: T)} \quad (\text{REC-I}_{qp-v}) \\
\\
\frac{\Gamma \vdash_{qp} v: T \quad \Gamma \vdash_{e1} q: \{A: T..T\}}{\Gamma \vdash_{qp} v: q.A} \quad (\text{TYP-SEL}_{qp-v}) \\
\\
\frac{\Gamma \vdash_{qp} v: r'.A \quad \Gamma \vdash_{e2} q}{\Gamma \vdash_{qp} v: r'.A [p/q, n]} \quad (\text{SNGL-SEL}_{qp-v})
\end{array}$$

Fig. 13. Introduction $_{qp}$ Typing

LEMMA A.3 (\vdash_{qp} TO \vdash_{pq} (\forall)). *If $\Gamma \vdash_{qp} p: \forall(x: T) U$ and Γ is inert then $\Gamma \vdash_{pq} p: \forall(x: T) U$.*

Introduction- pq denoted $\Gamma \vdash_{pq} p: T$, typing contains introduction rules and inlined versions of the SNGL_{pq} subtyping rule (see Figure 14).

To convert a function type from Introduction- pq typing into Elimination-III typing, we use the following lemma.

LEMMA A.4 (\vdash_{pq} TO \vdash_{e3} (\forall)). *If $\Gamma \vdash_{pq} p: \forall(x: T) U$ and Γ is inert then there exist types T' and U' such that $\Gamma \vdash_{\#} T <: T'$, $\Gamma, x: T \vdash_{\#} U' <: U$, and $\Gamma \vdash_{e3} p: \forall(x: T') U'$.*

If p has type T in Elimination-III typing, denoted $\Gamma \vdash_{e3} p: T$, 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 $\Gamma \vdash_{e3} p: T$, then one of the following is true:

- 1) $\Gamma \vdash_{e1} p: T$, i.e. T is the most precise type that Γ assigns to p (modulo possible recursion and intersection elimination), or

Introduction- pq path typing

$$\begin{array}{c}
\frac{\Gamma \vdash_{e3} p : T}{\Gamma \vdash_{pq} p : T} \quad (\text{PATH}_{pq}) \\
\\
\frac{\Gamma \vdash_{pq} p : \{a : S\} \quad \Gamma \vdash_{\#} S < : T}{\Gamma \vdash_{pq} p : \{a : T\}} \quad (\text{FLD-}<_{pq}) \\
\\
\frac{\Gamma \vdash_{pq} p : \{A : T..U\} \quad \Gamma \vdash_{\#} T' < : T \quad \Gamma \vdash_{\#} U < : U'}{\Gamma \vdash_{pq} p : \{A : T'..U'\}} \quad (\text{TYP}_{pq}) \\
\\
\frac{\Gamma \vdash_{pq} p : \forall(x : S) T \quad \Gamma \vdash_{\#} S' < : S \quad \Gamma, x : S' \vdash T < : T'}{\Gamma \vdash_{pq} p : \forall(x : S') T'} \quad (\text{ALL}_{pq}) \\
\\
\frac{\Gamma \vdash_{pq} p : S \quad \Gamma \vdash_{pq} p : T}{\Gamma \vdash_{pq} p : S \wedge T} \quad (\text{AND}_{pq}) \\
\\
\frac{\Gamma \vdash_{pq} p : T}{\Gamma \vdash_{pq} p : \top} \quad (\text{TOP}_{pq}) \\
\\
\frac{\Gamma \vdash_{pq} r : \mu(x : T) \quad \Gamma \vdash_{e1} p : q.\text{type} \quad \Gamma \vdash_{e2} q}{\Gamma \vdash_{pq} r : \mu(x : T) [q/p, n]} \quad (\text{SINGL-REC}_{pq}) \\
\\
\frac{\Gamma \vdash_{pq} r : r'.A \quad \Gamma \vdash_{e1} p : q.\text{type} \quad \Gamma \vdash_{e2} q}{\Gamma \vdash_{pq} r : r'.A [q/p, n]} \quad (\text{SINGL-SEL}_{pq}) \\
\\
\frac{\Gamma \vdash_{pq} r : r'.\text{type} \quad \Gamma \vdash_{e1} p : q.\text{type} \quad \Gamma \vdash_{e2} q}{\Gamma \vdash_{pq} r : r'.\text{type} [q/p, n]} \quad (\text{SINGL-SINGL}_{pq})
\end{array}$$

Introduction- pq value typing

$$\begin{array}{c}
\frac{\Gamma \vdash_{e1} v : T}{\Gamma \vdash_{pq} v : T} \quad (\text{PATH-}v_{pq}) \\
\\
\frac{\Gamma \vdash_{pq} v : \forall(x : S) T \quad \Gamma \vdash_{\#} S' < : S \quad \Gamma, x : S' \vdash T < : T'}{\Gamma \vdash_{pq} v : \forall(x : S') T'} \quad (\text{ALL-}v_{pq}) \\
\\
\frac{\Gamma \vdash_{pq} v : S \quad \Gamma \vdash_{pq} v : T}{\Gamma \vdash_{pq} v : S \wedge T} \quad (\text{AND-}v_{pq}) \\
\\
\frac{\Gamma \vdash_{pq} v : \mu(x : T) \quad \Gamma \vdash_{e1} p : q.\text{type} \quad \Gamma \vdash_{e2} q}{\Gamma \vdash_{pq} v : \mu(x : T) [q/p, n]} \quad (\text{REC-SINGL-}v_{pq}) \\
\\
\frac{\Gamma \vdash_{pq} v : T}{\Gamma \vdash_{pq} v : \top} \quad (\text{TOP-}v_{pq})
\end{array}$$

Fig. 14. Introduction- pq Typing

- 2) $p = p'.\bar{b}$ (i.e. p' is a prefix of p), $\Gamma \vdash_{e1} p' : q.\text{type}$ (i.e. $\Gamma \vdash_{e2} p'.\bar{b} : q.\bar{b}.\text{type}$), and either
- $T = q.\bar{b}.\text{type}$, or
 - $\Gamma \vdash_{e3} q.\bar{b} : T$.

The Elimination typing rules are presented in Figure 15. For the proof recipe, we do not need to further convert p 's type into Elimination-II and -I typings. However, these typing relations are needed to prove the lemmas of the proof recipe.

Elimination typing for values

$$\frac{\Gamma, x: T \vdash t: U \quad x \notin \text{fv}(T)}{\Gamma \vdash_{e1} \lambda(x: T) t: \forall(x: T) U} \quad (\text{ALL-I}_{e1})$$

$$\frac{x; \Gamma, x: T \vdash d: T}{\Gamma \vdash_{e1} \nu(x: T) d: \mu(x: T)} \quad (\{\}-\text{I}_{e1})$$

Elimination typing for paths \vdash_{e1}

$$\frac{\Gamma(x) = T}{\Gamma \vdash_{e1} x: T} \quad (\text{VAR}_{e1})$$

$$\frac{\Gamma \vdash_{e1} p: T \wedge U}{\Gamma \vdash_{e1} p: U} \quad (\text{AND}_2\text{-E}_{e1})$$

$$\frac{\Gamma \vdash_{e1} p: \mu(z: T)}{\Gamma \vdash_{e1} p: T[p/z]} \quad (\text{REC-E}_{e1})$$

$$\frac{\Gamma \vdash_{e1} p: T \wedge U}{\Gamma \vdash_{e1} p: T} \quad (\text{AND}_1\text{-E}_{e1})$$

$$\frac{\Gamma \vdash_{e1} p: \{a: T\}}{\Gamma \vdash_{e1} p.a: T} \quad (\text{FLD-E}_{e1})$$

 \vdash_{e2}

$$\frac{\Gamma \vdash_{e1} p: T}{\Gamma \vdash_{e2} p: T} \quad (\text{PATH}_{e2})$$

$$\frac{\Gamma \vdash_{e2} p: q.\text{type} \quad \Gamma \vdash_{e2} q.a}{\Gamma \vdash_{e2} p.a: q.a.\text{type}} \quad (\text{SINGL-E}_{e2})$$

 \vdash_{e3}

$$\frac{\Gamma \vdash_{e2} p: T}{\Gamma \vdash_{e3} p: T} \quad (\text{PATH}_{e3})$$

$$\frac{\Gamma \vdash_{e2} p: q.\text{type} \quad \Gamma \vdash_{e3} q: U}{\Gamma \vdash_{e3} p: U} \quad (\text{SINGL-TRANS}_{e3})$$

Fig. 15. Elimination Typing Rules