

Pure methods for roDOT (an extended version)

Vlastimil Dort, Yufeng Li, Ondřej Lhoták, Pavel Parížek

Abstract: Object-oriented programming languages typically allow mutation of objects, but pure methods are common too. There is great interest in recognizing which methods are pure, because it eases analysis of program behavior and allows modifying the program without changing its behavior. The roDOT calculus is a formal calculus extending DOT with reference mutability. In this paper, we explore purity conditions in roDOT and pose a SEF guarantee, by which the type system guarantees that methods of certain types are side-effect free. We use the idea from ReIm to detect pure methods by argument types. Applying this idea to roDOT required just a few changes to the type system, but necessitated re-working a significant part of the soundness proof. In addition, we state a transformation guarantee, which states that in a roDOT program, calls to SEF methods can be safely reordered without changing the outcome of the program. We proved type soundness of the updated roDOT calculus, using multiple layers of typing judgments. We proved the SEF guarantee by applying the Immutability guarantee, and the transformation guarantee by applying the SEF guarantee within a framework for reasoning about safe transformations of roDOT programs. All proofs are mechanized in Coq.

This is an extended version of the ECOOP 2024 paper Pure methods for roDOT [13].

This work was supported by the Czech Science Foundation project 23-06506S, and by the Czech Ministry of Education, Youth and Sports project LL2325 of the ERC.CZ programme. This research was also supported by the Natural Sciences and Engineering Research Council of Canada.

1 Introduction

A feature common to many object-oriented programming languages is that execution of a method can have important side effects such as creating new objects on the heap or modifying (mutating) existing objects. For example, a setter method modifies a field of the receiving object. Such effects are also the reason why, in general, execution of a method cannot be treated as evaluation of a function in a mathematical sense, because every call of a method with possible side effects can produce different results.

That being said, many methods in object-oriented programs are actually designed as side-effect-free and meant to work like pure mathematical functions, producing the same result on each invocation. An example of such methods are getters, or generally, computations based solely on the arguments passed into the method. Creating methods without side effects is also often considered to be a good practice, because it reduces hidden dependencies, and these methods can be used more freely without the fear of unwanted interaction of their effects. For example, the program code fragment `val x = computeX() ; val y = computeY()`, which involves two side-effect-free methods, can be transformed to `val y = computeY() ; val x = computeX()` by swapping the order of method calls without any observable change in the program behavior and semantics. Writing side-effect-free methods also enables a greater degree of parallelization (concurrency) and, in general, makes it easier to understand the program behavior. Therefore, the issue of purity is relevant to most mainstream object-oriented programming languages, such as Java, C++, C# and Scala.

However, in common programming languages, pure functions and methods with effects are typically unified under a single concept of a method, and there is no way to express, check and make use of method purity at the language level. The idea that a method is pure can be expressed using an annotation (see, e.g., Checker Framework [14, 10] and Code Contracts [15]), but one must look into the documentation of such an annotation for the exact meaning of purity, and there may be limited possibilities of checking automatically whether the annotation is applied properly.

In the context of Java, ReIm [19] introduced annotations with a formal meaning, which give rise to a type system that allows to recognize side-effect-free methods using the types of their parameters – if all parameters of a method, including the receiver, have read-only types, the method cannot get hold of a writable reference to an existing object, so it is necessarily side-effect free. The advantage of this general approach, based on the usage of static type systems for reasoning about purity and side-effect freedom, is the possibility to prove soundness and consistency of such annotations.

Scala favours a functional programming style, so Scala programs are likely to contain more methods (than Java programs) that can be identified as side-effect-free. Our main objective is to design a type system that guarantees side effect freedom for Scala methods and supports advanced language features present in Scala.

Previous formalization efforts for Scala resulted in the Dependent Object Types (DOT) calculus [2], which captures the essence of Scala’s type system. However, the original DOT calculus does not model mutation of objects, so purity cannot be addressed there, but some variants that do allow mutation have been developed. roDOT [12] is an existing core calculus for Scala with reference mutability. It has mutable fields and a type system feature to distinguish read-only and mutable references. An important rationale behind the design of roDOT, when compared to other possible approaches, is to use existing features of Scala, including its rich type system, as much as possible rather than introducing new forms of types only for reference mutability, to ease adoption of such a type system into the Scala language. In particular, roDOT expresses the mutability of a reference using a specially designated type member in the type of that reference. The type system of roDOT also provides an *immutability guarantee*: an object can only be mutated if there is a path of mutable references to it from the code being executed.

In this paper, we extend the core roDOT calculus from [12] with the concept of side-effect-free methods. Before going into details, we want to emphasize that it was not possible to simply adapt ReIm [19] from Java, because of several challenges specific to Scala and roDOT that we discuss below. However, we use the idea proposed by the authors of ReIm that side-effect-free methods are recognized based on the types of their parameters.

The general concept of purity is, in addition to (1) side-effect-freedom, sometimes understood to comprise more properties: (2) determinism – returning the same value for the same arguments [14], and (3) termination. In this paper, we focus only on the side-effect-free (SEF) property. We will just mention that in regards to determinism, mutable DOT calculi including roDOT have semantics that is deterministic except for instantiation of objects.

Within the context of roDOT, we look at the SEF property from three different perspectives – what a

SEF method does, how it can be recognized using the type system, and how it can be used in programs. We define SEF methods in roDOT as those that do not modify any objects that existed on the heap before the method was called.

As the main result of this paper, we prove the *side-effect-free guarantee* (SEF guarantee), which says that methods with read-only parameters do not modify existing objects.

One important related challenge is that in order to state and prove the SEF guarantee, we needed a way to test whether a given type is read-only. As we will explain, this is not possible in the existing (original) roDOT type system from [12]. Therefore, one of our contributions is an extension of roDOT that makes it possible to recognize read-only types.

Another challenge was defining the SEF guarantee formally and proving it in a calculus that supports a mutable heap (like roDOT). The roDOT operational semantics says that fresh heap addresses are chosen during method execution. Therefore, after calling a SEF method, these heap addresses can be different, yet the heap still has the same overall structure. We formally define a concept of similarity of heaps in roDOT to describe this relation. We prove the SEF guarantee by simulating the execution of a SEF method with a similar execution, where writeable references are removed, and by applying roDOT's immutability guarantee.

Finally, as a corollary of the SEF guarantee, we state and prove a guarantee of safety of a particular code transformation. The transformation guarantee states that swapping two calls to SEF methods anywhere in a program does not affect the result of its execution.

Formalizing safe program transformations has to deal with specific issues, such as mixing of program code and values together on the program heap, or the heap similarity mentioned above. In order to deal with these issues, we design a general framework for reasoning about safe transformations in roDOT. The framework provides a general way to define program transformations, defines what properties a safe program transformation must have, and provides a general theorem about lifting the safety of transformation from execution of a small piece of code to execution of the whole program. Within this framework, we define the specific transformation of swapping two calls of SEF methods. We prove the transformation guarantee using the SEF guarantee and the lifting theorem.

We mechanized all of our formal results, in particular the soundness proof of the extended roDOT calculus and the SEF guarantee, in Coq to enable future formal reasoning to build on them. Note that the soundness of the original roDOT calculus was proved by hand in [12]. We have made our formalization in Coq public as an artifact for this paper.

1.1 Contribution

In summary, the main contributions of this paper are the following:

- a modification of the original roDOT calculus that makes it possible to test whether a type is read-only, which is necessary to state and prove the SEF guarantee;
- a formal definition of side-effect-free methods in the context of roDOT, statement and proof of the SEF guarantee;
- a general framework for defining transformations of roDOT programs and proving that some of them are safe in that they do not change the result of program execution, statement and proof of a transformation guarantee, which states that re-ordering calls to SEF methods is safe in that sense;
- the first mechanization of roDOT and its immutability guarantee, with addition of the SEF and transformation guarantees, and all the proofs in Coq — provided as an artifact.

1.2 Outline

The paper is organized as follows. Section 2 gives an overview of the roDOT calculus, which has been mechanized in Coq and within which we define the SEF condition. Section 3 discusses the definition of pure and SEF methods, looking at several variants. It defines the SEF guarantee and identifies necessary changes to the roDOT type system in order for the guarantee to work. In Section 4 we describe the changes to the calculus in more detail, and discuss a new proof of type soundness of the calculus. In Section 5, we describe how we proved the SEF guarantee, and in Section 6 we define and prove the transformation guarantee within a framework for safe transformations.

$x ::= z, s, r$	Variable	$T ::=$	Type
$ y w$	location, reference	$ \top \perp N$	top, bottom, read-only \perp
$t ::=$	Term	$ \mu(s : T) x_1.B(x_2)$	recursive, selection
$ vx x_1.m x_2$	variable, method call	$ \{a : T_1..T_2\} \{B(r) : T_1..T_2\}$	field, type decl.
$ \text{let } z = t_1 \text{ in } t_2$	let	$ \{m(z : T_1, r : T_3) : T_2\}$	method
$ \text{let } z = \nu(s : T)d \text{ in } t$	let-literal	$ T_1 \wedge T_2 T_1 \vee T_2$	intersection, union
$ x.a x_1.a := x_2$	read, write	$B ::=$	Type member name
$d ::= d_1 \wedge d_2$	Definition	$ A M$	ordinary, mutability
$ \{a = x\} \{A(r) = T\}$	field, type	$\sigma ::= \cdot \text{let } z = \square \text{ in } t :: \sigma$	Stack
$ \{m(z, r) = t\}$	method	$\Sigma ::= \cdot \Sigma, y \rightarrow d$	Heap
$\rho ::= \cdot \rho, w \rightarrow y$	Environment	$c ::= \langle t; \sigma; \rho; \Sigma \rangle$	Configuration

Figure 1: roDOT syntax

$$\begin{array}{c}
\frac{\Gamma; \rho \vdash x_1 : \{m(z : T_1, r : T_3) : T_2\} \quad \Gamma; \rho \vdash x_2 : T_1 \quad \Gamma \text{ vis } x_2}{\Gamma; \rho \vdash x_1 : [x_2/z]T_3 \quad \Gamma \text{ vis } x_1} \quad \frac{\Gamma; \rho \vdash x_1 : T_1 \quad \Gamma \text{ vis } x_1}{\Gamma; \rho \vdash x : \{a : T_1..T_2\} \quad \Gamma \text{ vis } x} \\
\frac{\Gamma; \rho \vdash x_1.m x_2 : [x_1/r][x_2/z]T_2}{\Gamma; \rho \vdash x_1.m x_2 : [x_1/r][x_2/z]T_2} \text{(TT-Call)} \quad \frac{\Gamma; \rho \vdash x : \{M(r) : \perp.. \perp\}}{\Gamma; \rho \vdash x.a := x_1 : T_2} \text{(TT-Write)} \\
\frac{\Gamma; \rho \vdash x : \{a : T_1..T_2\} \quad \Gamma \text{ vis } x \quad \Gamma; \rho \vdash T_2 \text{ ro } T_3 \quad \Gamma; \rho \vdash T_2 \text{ mu}(r) T_4}{\Gamma; \rho \vdash x.a : T_3 \wedge \{M(r) : \perp..(T_4 \vee x.M(r))\}} \text{(TT-Read)} \quad \frac{\Gamma; \rho \vdash T_1 <: T_3 \quad \Gamma; \rho \vdash T_2 <: T_3}{\Gamma; \rho \vdash T_1 \vee T_2 <: T_3} \text{(ST-Or)} \\
\frac{}{\Gamma; \rho \vdash T_1 \wedge (T_2 \vee T_3) <: (T_1 \wedge T_2) \vee (T_1 \wedge T_3)} \text{(ST-Dist)} \\
\frac{\Gamma; \rho \vdash T_3 <: T_1 \quad \Gamma, z : T_3, r : T_6; \rho \vdash T_2 <: T_4 \quad \Gamma, z : T_3; \rho \vdash T_6 <: T_5 \quad T_6 \text{ indep } z \Rightarrow T_5 \text{ indep } z}{\Gamma; \rho \vdash \{m(z : T_1, r : T_5) : T_2\} <: \{m(z : T_3, r : T_6) : T_4\}} \text{(ST-Met)}
\end{array}$$

Figure 2: Selected rules for typing and reduction in roDOT

2 Background – the roDOT calculus

In this section, we present the summary of the roDOT calculus [12], which we use as the baseline for this work. The DOT calculus [2, 33, 30] is a formal calculus, designed to formalize the essence of the types of the Scala programming language. In the basic versions of the DOT calculus, objects have read-only fields (so the objects are immutable), but there are also several versions that allow changing values of the fields of objects (mutation).

The roDOT calculus [12] evolved from DOT with mutable fields. The goal was to extend DOT with the ability to control mutation of objects using the type system, while using the existing features of the DOT calculus, dependent types.

In roDOT, write access to a field is controlled by a reference mutability permission. It is based on an idea of a reference capability represented by a special type member M . A reference can only be used to mutate an object if the type of the reference includes this capability, in the form of a type member declaration $\{M : \perp.. \perp\}$. Thanks to that, we can refer to the mutability of a variable x using type selection $x.M$.

Without this capability, the field can only be read, but with it, the field can also be written to. The permission applies transitively, in the sense that reading from a read-only reference always produces read-only references.

2.1 Syntax and typing

The syntax of terms and types in roDOT is in Figure 1. It uses the A-normal form [36] of terms from DOT. To avoid ambiguity, if a variable is used in the position of a term, it is marked as vx . Unlike other versions, the roDOT calculus does not have λ values, but methods are a kind of object member (and cannot be reassigned), so there is a more explicit relationship of a method, the containing object and the

reference used to call the method. Objects are represented by the $\nu(s : R)d$ constructor, appearing as literals in the programs and as items on the heap (R is the type of the object and d is a list of member definitions).

When typing the program or a part of it, free variables are assigned a type in a typing context Γ . There are several kinds of variables. *Abstract variables* are variables bound in terms such as let-in terms and method definitions. When the program executes, objects are created on the heap, and variables referring to concrete objects on the heap are substituted in place of the abstract variables. Each object on the heap has a unique location y and one or more references w . In an object on the heap, the values of fields are locations of other objects. In terms, only references may appear. The kind of the variable has no effect on execution or typing. In roDOT, references are a separate concept from locations in order to allow references to the same object to have different types (specifically, different mutabilities). While the run-time stack and focus of execution work with references that have their own mutabilities, the heap only works with locations, and mutability is determined by field types.

The types form a lattice, with the top, bottom, union and intersection types. Objects can contain multiple members – fields, methods and type members. Types of objects are formed by intersection of individual declaration types for each member.

The type members $\{A : T..T\}$ specify lower and upper bounds, and they introduce a new dependent type $x.A$ that has a subtyping relationship with those bounds. This is relevant because roDOT uses a type member for mutability. The ability to create dependent types in this manner is the defining feature of the DOT calculus.

The declarations of an object’s members are wrapped in a recursive type, so several declarations in one object type can reference each other, using a member type selection $s.A$ involving the self-reference s . An example of a type of an object without mutability is $\mu(s : \{A : T..T\} \wedge \{a : s.A..s.A\} \wedge \{m(r : T, z : T) : T\})$. An object of this type has a type member A with bounds T , a field a with a self-referential type $s.A$, and a method m .

In roDOT, dependent types are also used to express the mutability of a reference, by selecting the special type member M . When accessing an object through a reference which does not have this capability, for example $\{a : T..T\}$, the field can only be read. With it, for example $\{a : T..T\} \wedge \{M : \perp..\perp\}$, the field can also be written to.

In the declaration of the type member M , the lower bound is always \perp , and the upper bound determines the mutability. If the upper bound is also \perp , it means the reference is mutable. Otherwise, it is read-only. This way, mutable references are subtypes of read-only references, so a mutable reference can be used anywhere a read-only reference is expected, but not vice versa. We will use \mathcal{M}_T as a shorthand for the type member declaration $\{M : \perp..T\}$, or just \mathcal{M} when the bound is not important. The mutability of a reference applies to the whole object – a mutable reference allows writing to all fields.

An example of a type of an object with a type member A , a field a , method m and a mutability declaration is $\mu(s : \{A : T..T\} \wedge \{a : T..T\} \wedge \{m(r : T, z : T) : T\}) \wedge \{M : \perp..\perp\}$.

A declaration of a method allows specifying a type of the receiving reference $r : T$, which can be more precise than the type of the recursive self parameter s in the defining object. This allows the type of the method to require that the receiver be writeable, or allow it to be read-only. It is similar to the ability to annotate the type of `this` parameter in Java, used by the Checker Framework [18, 9]. For this reason, every method in roDOT has two parameters: a normal parameter z and the receiver r , which is a reference to the object containing the method, like `this` in Scala. In roDOT, the type of r can be dependent on z . The parameter r is special in how it gets its type, but in terms of semantics, behaves the same as z .

Several rules in roDOT need a read-only version of a type. For that, there are two type-level operations: $T \text{ ro } U$ means that U is a readonly version of T , $T \text{ mu } U$ means that U is a mutability bound of type T (rules are shown in Figure 11 in the appendix). A special type N is defined to be the read-only version of the least type in the subtyping lattice, \perp .

The typing rules (selected in Figure 2, full set in Figures 8 to 12 in the appendix) describe correctly formed programs. In addition to the typing context Γ , which assigns types to variables, the left side of the typing judgment includes an environment ρ that connects references in the terms to locations of objects on the heap.

The write term, typed by TT-Write, is guarded by a check of the mutability permission on the receiving reference. The premise $\Gamma; \rho \vdash x : \{M : \perp..\perp\}$ ensures that only mutable references can be used for writing.

Reading a field, typed by TT-Read, is always possible, but the type of the result is changed to read-only if the source reference is read-only. This type operation is called viewpoint adaptation, and ensures

that read-onlyness is transitive, which is required for the immutability guarantee of roDOT and for our SEF guarantee. This is achieved by taking a read-only version of the field's type, and adding a mutability that is a union of the mutabilities of the source reference and of the field type. For example, if a reference w has type $\{a : T_1.. \mu(s : \dots) \wedge \mathcal{M}_U\}$, then the term $w.a$ has type $\mu(s : \dots) \wedge \mathcal{M}_{w.a \vee U}$.

With the `vis` judgment (Figure 9 in the appendix), roDOT hides captured variables in methods – to access a value from outside, it must be stored in a field of the containing object, so viewpoint adaptation applies to it.

Variables appearing in terms and definitions have types given by the typing and subtyping rules in Figures 9 and 10 in the appendix. Selected rules are shown in Figure 2: ST-Met is a subtyping rule for method declarations. The part highlighted in grey is not part of roDOT, but our modification, which we will describe in Section 4.2. Rules ST-Or and ST-Dist are examples of subtyping rules for union types, which are relevant in Section 4.1.

2.2 Semantics

The operational semantics of roDOT is defined as a small step semantics, with machine configurations (Figure 1) consisting of a term in the focus of execution t , a stack σ , a heap Σ and an environment ρ . The environment ρ maps references to locations and the heap Σ maps locations to objects. The stack σ is used to evaluate terms of the form `let $z = t_1$ in t_2` . The stack is a list of frames of the form `let $z = \square$ in t_2` , where \square represents t_1 while it is being evaluated in the focus of execution. When t_1 is evaluated to a value, that value is substituted for the square in the top frame of the stack, and the t_2 from that frame then becomes the new focus of execution.

Execution starts with the program, an empty stack, empty heap and an empty environment, and proceeds by steps defined in Figure 13 in the appendix, until it reaches an answer configuration, which has an empty stack and the focus of execution is a single variable. During execution, new items are added to the heap and the environment (there is no garbage collection). Calling a method copies its body to the focus of execution, while the receiver and argument are substituted.

The semantics is generally deterministic – there is no way to express a nondeterministic choice. However, there is one source of non-determinism: locations of objects on the heap. Allocating objects must be regarded as a non-deterministic operation because even if the new objects are initially equal, they may take on different values due to subsequent mutation.

2.3 Properties

The roDOT calculus has the type soundness property (Theorem 1, Theorem 7 in [12]) – a term that has a type in an empty context can be executed and either reduces to an answer, or executes indefinitely. DOT and roDOT do not include explicit checks for error conditions, but trying to access (read, write or call) a non-existing member of an object is an error. In such a case, a reduction step is not defined and the execution “gets stuck”. The soundness theorem guarantees this does not happen for typed programs.

Theorem 1 (Type Soundness).

$\text{If } \vdash t_0 : T,$	$\text{then either } \exists w, j, \Sigma, \rho: \langle t_0; \cdot; \cdot \rangle \mapsto^j \langle vw; \cdot; \rho; \Sigma \rangle,$
$\text{or } \forall j: \exists t_j, \sigma_j, \Sigma_j, \rho_j: \langle t_0; \cdot; \cdot \rangle \mapsto^j \langle t_j; \sigma_j; \rho_j; \Sigma_j \rangle.$	$\text{The initial term } t_0 \text{ is well typed,}$ $\text{then execution terminates in } j \text{ steps}$ $\text{with answer } w,$ $\text{or continues indefinitely.}$

Type soundness and other properties are based on the fact that during execution, the type of the configuration is preserved. Rules for typing a machine configuration are in Figure 14 in the appendix. As the program executes and new objects are added to the heap, new locations and reference variables are used to refer to the objects. To give the configurations a type, these variables are added to the typing context. Their type is the type of the object, and has a fixed form – it is a recursive type containing declarations of all the object's members, intersected with a declaration of mutability. A typing context that only contains types of this form is called an *inert context*. Under an inert context, stronger claims can be made about types of variables [30], and it plays an important role in the proof of soundness.

The essential property of roDOT is the immutability guarantee (Theorem 2, Theorem 9 in [12]): in order for an object to be mutated, a writeable reference to it must exist, or it must be possible to reach it by a path of writeable fields, starting from a writeable reference – the object must be mutably reachable, defined formally in Figure 3.

$$\begin{array}{c}
\Gamma \vdash \langle t; \sigma; \rho; \Sigma \rangle \mathbf{mreach} y_1 \\
y_1 \rightarrow \dots \{a = y_2\} \dots \in \Sigma \\
\Gamma; \rho \vdash y_1 : \{a : \perp.. \{M(r) : \perp.. \perp\}\} \\
\hline
\Gamma \vdash \langle t; \sigma; \rho; \Sigma \rangle \mathbf{mreach} y_2 \quad (\text{Rea-Fld})
\end{array}
\qquad
\begin{array}{c}
t \mathbf{tfree} w \vee \sigma \mathbf{tfree} w \\
w \rightarrow y \in \rho \\
\Gamma; \rho \vdash w : \{M(r) : \perp.. \perp\} \\
\hline
\Gamma \vdash \langle t; \sigma; \rho; \Sigma \rangle \mathbf{mreach} y \quad (\text{Rea-Term})
\end{array}$$

Figure 3: roDOT mutable reachable references

Theorem 2 (Immutability Guarantee).

If $y \rightarrow d \in \Sigma_1$ and $\Gamma \vdash \langle t_1; \sigma_1; \rho_1; \Sigma_1 \rangle : T$,
and $\langle t_1; \sigma_1; \rho_1; \Sigma_1 \rangle \mapsto^k \langle t_2; \sigma_2; \rho_2; \Sigma_2 \rangle$,
then either $y \rightarrow d \in \Sigma_2$,
or $\Gamma \vdash \langle t_1; \sigma_1; \rho_1; \Sigma_1 \rangle \mathbf{mreach} y$.

For an object at some point during well-typed execution,
at any later point,
either the object does not change,
or it was reachable by mutable references.

3 Method Purity for roDOT

Here we informally define the meaning of side-effect freedom in roDOT, and informally state the main results of this paper: the SEF guarantee and the transformation guarantee.

We structure our work around an observation that (in any programming language or calculus), we can look at side-effect freedom from different perspectives:

1. (Static) Recognize which methods are SEF statically at compile time, using types.
2. (Runtime) Define what events can (or cannot) happen when a SEF method is executed.
3. (Usage) Differentiate SEF methods from general methods based on how they can be safely used in programs.

For each of these perspectives, we will state a *SEF condition*, each giving a different definition of SEF methods in roDOT. First we do it informally in this section, and then formalize the definitions in the following sections. The guarantees then form connections between different SEF conditions.

3.1 Runtime SEF condition

Saying that a method is side-effect-free is informally understood as saying that the execution of the method will not perform any actions that are considered to be side effects. This view corresponds to the second perspective on our list.

This perspective is most directly related to the semantics. In roDOT, this means looking at the small step semantics, defining the beginning and end of execution of a method, and defining the SEF condition in terms of the state of execution or the steps performed between the beginning and the end. When looking at the effects caused by method execution, the only relevant part of the machine configuration is the heap (the focus of execution is the part being evaluated, the stack cannot be changed, and the mapping from references to locations is only relevant for typing). The heap can only be modified by two kinds of execution steps: instantiation of an object and writing a value to a field of an object on the heap.

The condition of side-effect-freedom can be stated in multiple versions of varying strength. In the strictest sense, we could say that a SEF method cannot have any effect on the heap at all, meaning no instantiations and no writes. That would, however, be overly restrictive, as object instantiation is one of the basic operations in object-oriented programming. It is therefore usually (such as in [34, 38, 19, 14]) allowed that a SEF method can instantiate new objects, and also write to the fields of those newly instantiated objects. In turn, the only forbidden action is writing to fields of previously existing objects.

Another choice in the definition is when the change to the heap is detected, which leads to different answers to questions such as: (a) Is it allowed to write to a field of an existing object, if the value written is the same as the current value so the object does not actually change? (b) Is it allowed to write to a field of an existing object, if the field is restored to the previous value before the end of the method execution? We choose to allow (a) but not (b), so our definition observes the state of the heap at every moment during the execution of the method. Allowing (b) would lead to a weaker condition, which would check the state of the heap only at the end of the method call. Forbidding (a) would lead to a stronger condition, defined in terms of allowed steps of execution rather than in terms of the state.

Informal statement of Definition 15 (Run-time SEF condition, in Section 5.1). An execution of a method is side-effect free, when at every step of execution until returning from the method, the heap contains all the objects from the start of execution in an unchanged state.

3.2 Static SEF condition

The static perspective (the first in our list) is useful because it provides a way to check that a method is SEF by looking at the code. We must, however, accept that statically, it will not be possible to recognize all methods that are pure from the second (and third) perspective.

In ReIm [19], SEF methods are recognized by the mutability of the parameters. roDOT uses the same notion of transitive read-only references, therefore it should be possible to use an analogous condition in roDOT.

Informal statement of Definition 11 (static SEF condition). A method has a SEF type, if both its parameter and its receiver parameter have read-only types.

This condition will be formally defined in Section 4.1. Example 3 and Example 4 illustrate its ability to recognize SEF methods.

Example 3. A getter defined as $\{m_{get}(r, z) = z.a\}$ can be typed with $\{m_{get}(r : \top, z : \{a : \top.. \top\}) : \top\}$. Both \top and $\{a : \top.. \top\}$ are read-only types, and therefore the getter is SEF.

Example 4. The method m_{sef} defined by $\{m_{sef}(r, z_a) = (\text{let } x = \nu(r_o : R_o) \dots \text{in } z_a.m_a x)\}$ calls a method of its argument, passing a newly allocated object to it. This method has type $\{m_{sef}(r : \top, z_a : T_z) : \top\}$, where $T_z = \{m_a(r : \top, z : \mu(r_o : R_o) \wedge \{M : \perp.. \perp\}) : \top\}$. By Definition 11, m_{sef} is SEF, because it has read-only parameters, even though it calls m_a , which may mutate the heap.

Example 5 shows how viewpoint adaptation transitively ensures that read-only parameters cannot be used to modify existing objects. Example 6 shows how a dependent type can change whether the method is SEF or possibly not.

Example 5. The method defined by $\{m_{va}(r, z) = (\text{let } x = z.a \text{ in } x.b := r)\}$ mutates an object stored in a field of the argument z , and therefore is not SEF. This method cannot be typed with a read-only type for the parameter z , because even if the field a has a mutable type, by viewpoint adaptation of fields in roDOT, the variable x would also have a read-only type, so the subsequent write would not be allowed.

Example 6. A method with a type $\{m_{dep}(r : \top, z_a : \{a : \top.. \top\} \wedge x.A) : \top\}$ has a parameter with a type dependent on the variable x , which can decide the mutability. This method is recognized as SEF only in contexts where $N <: x.A$. When $x.A <: \mathcal{M}_\perp$, then the method can (indirectly) mutate the argument.

3.3 SEF guarantee

For the **SEF guarantee** (Theorem 16), we want to be able to claim that a method is SEF based on the type of the method declaration. The SEF guarantee makes the connection from the first to the second perspective.

Informal statement of Theorem 16 (SEF guarantee, in Section 5.2). Let c_1 be a well-typed machine configuration just prior to executing a method call step $w_1.mw_2$. If, by typing of the receiving reference w_1 , the method m has a SEF type, then the execution of the method will be side-effect free.

3.4 Using pure methods in roDOT

Finally, the third perspective shows why SEF methods are useful. It is, however, a view from outside of the method, and does not tell us how to construct a SEF method or check it.

The practical use of a type system with SEF methods comes when it allows us to look at the code, and based on what we see (from the first perspective) gives us a guarantee about its behavior (second perspective) and how it can be used (the third perspective). An example of this is allowing safe transformations of the program, which can be applied at coding time using IDE-provided code transformations, or at compile time as optimizations. For example, calls to SEF methods can be safely reordered.

To keep the problem simple, we will look at one particular case of such reordering: swapping two calls to SEF methods. With SEF methods, the code $x1.m1() ; x2.m2()$ is equivalent $x2.m2() ; x1.m1()$.

Informal definition (Call-swapping transformation of programs). A program t_1 is transformed into t_2 by SEF call-swapping, when the programs are the same except in one place, where t_1 calls two methods in succession, but t_2 calls them in the opposite order. Furthermore, within the contexts of typing these method calls, both methods have the same read-only types, and allow both programs to be typed in the same manner.

Example 7. A chain of calls let $x_1 = x_{o1}.m_1x_{a1}$ in let $x_2 = x_{o2}.m_2x_{a2}$ in t , can transformed by call swapping into let $x_2 = x_{o2}.m_2x_{a2}$ in let $x_1 = x_{o1}.m_1x_{a1}$ in t .

The static condition from the first perspective is already a part of the definition of the transformation. The transformation guarantee then states that this transformation is safe – it does not change the behavior of the program. By that, the guarantee connects the static condition (first perspective) with the call-swapping transformation (third perspective). We use the run-time condition (second perspective) as a connecting step between them in the proof of this guarantee.

Informal statement of Theorem 27. The call-swapping transformation is safe, in the sense that if for any programs t_1 and t_2 related by this transformation, provided that t_1 terminates with an answer c_1 , then t_2 also terminates with an answer c_2 , which is the same as c_1 , except for certain unavoidable differences in variable names and in method bodies.

The formal definition of the transformation, formal statement of the transformation guarantee and an outline of its proof are provided in Section 6.

4 Recognizing SEF methods by type in modified roDOT

In this section, we formalize the static SEF condition in roDOT given informally in Section 3.2. Although the notion of read-only types, used by this condition, was already defined in roDOT, we identify issues with that definition in regards to this new use.

We fix them by updating the calculus with small changes, which comprise adding one new subtyping rule and one type splitting rule, and one restriction added to the method subtyping rule. The updated calculus is neither a subset nor a superset of the original, so it is necessary to update the proof of soundness and the immutability guarantee, which were proven by hand for the original roDOT [12]. The soundness proof followed the scheme from [30] and uses an auxiliary definition of invertible typing, which allows doing proofs by induction on the typing of variables. This is possible thanks to eliminating possible cycles in the derivation, by forcing the derivation to follow the syntactic structure of the target type.

One of the new subtyping rules, however, breaks this soundness proof, because it introduces new possibilities to derive types in cycles, which cannot be repaired by simply handling additional cases in the original proof. In the presence of cycles, we cannot use the straightforward inductive hypothesis to prove properties necessary for type safety, because a derivation for typing a variable can involve derivations of arbitrarily complex types.

We implemented a new proof based on a different auxiliary typing definition, which avoids cycles by forcing the derivation to arrive at the target type by adding type constructors in a fixed order (for example, all unions in the type are handled before intersections). Compared to the original invertible typing, which was single typing judgment with many rules, the new approach leads to a definition in several layers, where each layer has a small number of typing rules. We call this set of judgments *layered typing*. In layered typing, we re-prove important properties of invertible typing, so that the new definition fits into the rest of the existing soundness proof, and also prove new properties required for the SEF guarantee.

The rest of this section is structured as follows: in Section 4.1, we formalize static SEF condition, and discuss the meaning of read-only types in roDOT. In Section 4.2, we propose small changes to the roDOT calculus to make definitions work for the SEF guarantee. We give a short overview of the structure of the original soundness proof for roDOT, and show how this proof breaks with the new changes. In Section 4.3, we describe the new layered typing that replaces invertible typing in the updated proof and show its important properties.

4.1 Static SEF condition in roDOT

In the **SEF guarantee** (Theorem 16), we claim that a method is SEF based on the type of the method declaration. Our SEF guarantee follows the approach of ReIm [19] and requires the parameters to have

read-only types.

4.1.1 Read-only types in roDOT

The check whether a type is read-only was also present in roDOT, but it had a limited purpose – to ensure that recursive types are read-only in VT-RecI (Figure 9 in the appendix). It was not based on subtyping, but rather on the relation ro , which makes a read-only version of a type using a syntax-based type splitting.

This definition did not guarantee that all supertypes of a read-only type are also read-only. As we will explain in Section 4.1.3, this would be a critical problem for the SEF guarantee.

We solve this problem by using a different notion of read-only types, based on subtyping with the “read-only bottom” type \mathbb{N} .

The purpose of \mathbb{N} in the original roDOT was to be the read-only version of the type \perp for defining the ro relation. Because the bottom type \perp is a subtype of all types, it is inherently mutable. For that reason, the type \mathbb{N} was added and made a lower bound of read-only types. That allows us to define read-only types as supertypes of \mathbb{N} .

Definition 8 (Read-only types). *A type T is read-only, if $\Gamma; \rho \vdash \mathbb{N} <: T$.*

With Definition 8 settled, we discovered a few problems related to read-only types, which would not allow us to state the SEF guarantee in the original roDOT.

Our proof of the SEF guarantee, specifically Lemma 19 in Section 5.4, relies on the idea that if a reference has some read-only type, then any other reference to the same object has that type too. Note that because of subsumption, a variable of a mutable type also has the corresponding read-only types. This essentially means that in any place where a reference is used by virtue of its read-only type, it can be replaced with a read-only version of that reference. With the new Definition 8 of read-only types, this can be stated as:

Lemma 9 (Read-only types are shared by all references). *If $\Gamma \sim \rho$ and $\Gamma; \rho \vdash y : T$ and $\Gamma; \rho \vdash \mathbb{N} <: T$, then $\Gamma; \rho \vdash w : T$ for any w such that $\rho(w) = y$.*

This key lemma, however, does not hold in the original roDOT, because of union types.

Union types were not a part of DOT, but were added to roDOT in order to be used to define view-point adaptation (union types are already a part of Scala’s type system), along with the subtyping rules ST-Or, ST-Or1, ST-Or2, ST-Dist, which are shown in Figure 10 in the appendix. However, using unions, it is possible to construct a type that is a supertype of both \mathbb{N} and a mutability declaration:

Example[†] 10 (In the original roDOT, counter-example to Lemma 9). Let $T_{\text{am}} := \{a : T_a\} \vee \mathcal{M}_{\perp}$ be a union of some field declaration with a declaration of mutability, and $T_{\text{bm}} := \{b : T_b\} \wedge \mathcal{M}_{\perp}$ be a type of a writable reference to some other field b .

The type T_{am} is not mutable, because it is not a subtype of \mathcal{M}_{\perp} . It is read-only, because $\Gamma \vdash \mathbb{N} <: T_{\text{am}}$, by the rules of subtyping of union types and by rule ST-N-Fld (Figure 10 in the appendix).

Let y_1 be a location of type T_{bm} , and w_2 be a reference to y_1 with type $T_b := \{b : T_{\text{bm}}\}$. By subtyping of unions and intersections, $\Gamma \vdash T_{\text{bm}} <: T_{\text{am}}$, so by subsumption, y_1 has type T_{am} . By Lemma 9, w_2 should have also type T_{am} , but in the original roDOT, it does not.

Example[†] 10 is marked with the [†] sign, which we use in this chapter to identify properties of the original roDOT from prior work, in contrast to the modified roDOT in this paper.

We observe that the read-only type T_{am} in this counter-example is a union of disjoint declarations, so it does not allow accessing the field a_{am} , or any other member. Therefore, T_{am} is no more useful for typing programs than \top . In order to make Lemma 9 work, we decided to extend the type system with new subtyping rules to make types like this equivalent to \top . These changes will be described in Section 4.2.

4.1.2 The SEF condition

A method is statically SEF if the types of its receiver and parameters are read-only according to Definition 8 i.e., they are supertypes of \mathbb{N} . Thanks to subsumption and subtyping of method types, the type $\{m(z : \mathbb{N}, r : \mathbb{N}) : \top\}$ is a type bound for methods named m and requires that both the argument and receiver have read-only types.

$$\frac{\Gamma; \rho \vdash N <: T}{\Gamma; \rho \vdash T \text{ ro } T} \text{(TS-N)} \quad \frac{}{\Gamma; \rho \vdash \top <: N \vee \{M(r) : T_1..T_2\}} \text{(ST-NM)}$$

Figure 4: New rules for roDOT

Definition 11 (Static SEF condition). *A method is statically SEF if it has a type $\{m(z : N, r : N) : \top\}$*

In Section 5, we will show that this condition works because a method must access all objects through the argument or the receiver (capturing values is modeled using fields of the receiver), so the method will not be able to get a writeable reference to any existing object.

4.1.3 Subtyping of method types

In order for the SEF guarantee (Theorem 16) to work with Definition 11, it is critical that all subtypes of a SEF method type are also SEF. The reason is at the site of a method call, the observed static type of the method is a supertype of the actual type of the method within its containing object, so this is needed to make the connection from the SEF type at a call site to the SEF type of the actual method.

That is why Definition 8 needs to be based on subtyping, so that all supertypes of read-only types are read-only (method types are contravariant in their parameter types).

Still, the type system required one more change related to a possible dependency between the types of method parameters. In roDOT, the type of the receiver r can be a dependent type referring to the other parameter z of the method. If, however, the receiver type depended on the mutability of z , then while typing the body of the method, it would be possible to derive that z is mutable, even if its type is read-only in the sense of Definition 8. If r has the type $\{A : z.M.. \perp\}$, one can use the typing rules ST-SelL and ST-SelU (Figure 10 in the appendix) to derive $z.M <: \perp$. The change to the rules TT-Call and ST-Met in Figure 2 prevents this issue by disallowing using method types where the receiver depends on the mutability of the parameter.

4.2 The updated roDOT calculus

In the previous section, we defined the static SEF condition, but identified several reasons why this definition would not work as intended in roDOT as-is. We fix these issues by changes to the roDOT calculus, which amount to two new and one modified typing rule:

- A new subtyping rule ST-NM (Figure 4) is added, which makes the union of a mutability declaration and the read-only lower bound N a top-like type (the other direction of subtyping was already a part of the type system).
- A new rule TS-N (Figure 4) is added to type splitting, making it so that all types that are read-only by Definition 8 are unaffected by the splitting operation.
- The typing rule TT-Call and subtyping rule ST-Met have a new premise (shown highlighted in Figure 2), which disallows introducing a dependency between the receiver type and the parameter in method subtyping. This fixes the problem described in Section 4.1.3.

The new rule ST-NM fixes the counter-example to Lemma 9, because now we have $\Gamma; \rho \vdash \top <: T_{\text{am}}$, derived from $\Gamma; \rho \vdash \top <: N \vee \mathcal{M}_{\top}$ and $\Gamma; \rho \vdash N <: \{a : T_a\}$. By subsumption and $\Gamma; \rho \vdash w_2 : \top$, that also means that $\Gamma; \rho \vdash w_2 : T_{\text{am}}$.

Additionally, we can now improve the type splitting relation $\vdash \text{ ro }$, by extending it with a new rule TS-N, shown in Figure 4. With that, the condition in VT-RecI (Figure 9 in the appendix) that recursive types are read-only, $\Gamma; \rho \vdash T \text{ ro } T$, becomes equivalent to Definition 8:

Lemma 12 (Read-only types). $\Gamma; \rho \vdash N <: T \Leftrightarrow \Gamma; \rho \vdash T \text{ ro } T$.

4.2.1 Updating the safety proof

The changes described above require updating the type safety proof of the calculus, to show that the changes did not allow invalid programs to be typed. The new subtyping rule ST-NM has a significant effect on the soundness proof, because it makes it possible to derive many additional union types, such as the now top-like type $\mathcal{M} \vee N$.




Figure 5: Dependencies (\rightarrow) and equivalences (\Leftrightarrow) between definitions of typing in roDOT

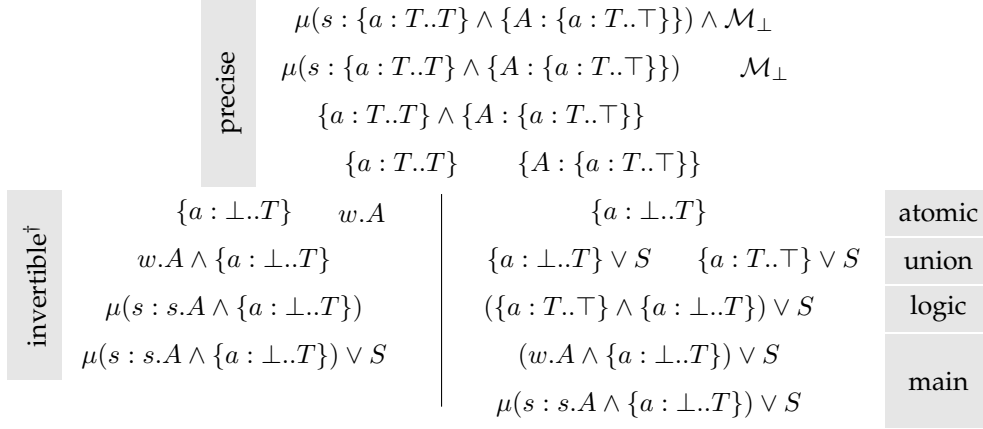


Figure 6: Example derivation of a type by invertible typing (left) and layered typing (right). Assuming that a variable w has the type $\mu(s : \{a : T..T\} \wedge \{A : \{a : T..\top\}\}) \wedge \mathcal{M}_\perp$ in the typing context, w has all the types shown here, ordered from types that are simple to derive at the top, to more complex derivations, which make use of derivations above. S is an arbitrary type.

The proof of soundness of roDOT before these changes followed the structure of the proof of DOT [30]. The core part of this proof is to show that if a reference w has some declaration type D (such as a field $\{a : T\}$), then the type associated with w in the typing context Γ is an object type containing D or a more precise declaration of the same member. That means, for $\Gamma; \rho \vdash w : D$, where D is a declaration type, because the types in Γ correspond to the object on the heap ($\Gamma \sim \Sigma$), the actual object referred to by w must contain a corresponding member definition in Σ , and therefore it is safe to access that member.

The proof was based on two alternative definitions of typing for variables – tight typing and invertible typing. Figure 5 shows the relations between the different versions of typing.

Tight typing is used as an intermediate step in equivalence of general and invertible typing. It is very similar to general typing – it has the same rules, except that subtyping rules involving selection types (ST-SelL and ST-SelU in Figure 10 in the appendix) use precise typing, a simpler version of variable typing, which does not have subsumption.

Updating tight typing for our modified rules is straightforward – we apply the same changes as to general typing, and the proof of equivalence between general and tight typing still works. However, we will show that updating invertible typing poses a challenge, as it cannot be easily extended with the additional rules.

4.2.2 Invertible Typing[†]

The main utility of invertible typing was providing a simple path of derivation of a variable’s type, starting from the type given to it by the typing context, and ending with a type that was used to access a member at some particular point in the program. This direct path then allowed induction-based proofs of properties of the typing relation.

This task would be especially hindered if the typing rules allowed cycles in the derivation, which would allow the derivation to go through unnecessarily complicated types. For example, with general typing, it is possible to derive $\Gamma \vdash x : T$ from $\Gamma \vdash x : T \wedge T$ and vice versa. Therefore, a derivation of type T can start with $\Gamma \vdash x : T$, go through arbitrarily complicated types such as $(T \wedge T) \wedge T$, and come back to T . This inhibits arguments by induction on the derivation of a general typing.

Invertible typing in DOT prevented this by ensuring that the derivation closely follows the syntactic structure of the target type.

The original roDOT adopted the invertible typing from DOT [30], where it has two layers, which we present using an example derivation of a type for a variable w in Figure 6.

The first layer, **precise typing**, only derives types by deconstructing the type of the variable given by the typing context. For each reference w , its type in the typing context is an intersection of a mutability declaration with a recursive type containing an intersection of declarations. Precise typing allows opening this recursive type and extracting the declarations from the intersection, but does not support subtyping. The top of Figure 6 shows individual steps of this process.

The second layer, **invertible typing**[†], combines both variable typing and subtyping into a single layer. In DOT and the original roDOT, it has fewer rules than general typing and subtyping, because it only has rules that construct the target type syntactically “bottom-up”, such as closing recursive types (akin to VT-Recl), or deriving intersection and union types. Thus the derivations of invertible typing are unambiguously guided by the syntax of the target type. The left side of Figure 6 shows individual steps of this process in building up the type $\mu(s : s.A \wedge \{a : \perp..T\}) \vee S$ for w .

As per Figure 5, invertible typing was equivalent to tight typing. That required invertible typing to be closed under tight subtyping (Lemma[†] 13).

Lemma[†] 13 (In original roDOT, invertible typing is closed under subtyping).

If $\Gamma; \rho \vdash_{\#\#} x : T_1$, and $\Gamma; \rho \vdash_{\#} T_1 <: T_2$, where $\Gamma \sim \rho$, then $\Gamma; \rho \vdash_{\#\#} x : T_2$.

The addition of ST-NM, together with the rules ST-Or and ST-Dist (Figure 2), breaks this. In Lemma[†] 13, the case for ST-Or relies on case analysis of deriving union types (Lemma[†] 14).

Lemma[†] 14 (In original roDOT, typing with union types can be inverted).

If $\Gamma; \rho \vdash_{\#\#} x : T_3 \vee T_4$, then either $\Gamma; \rho \vdash_{\#\#} x : T_3$ or $\Gamma; \rho \vdash_{\#\#} x : T_4$.

However, the rule ST-NM adds new ways of deriving union types such as $N \vee \mathcal{M}_\perp$, and the distributivity rule ST-Dist actually allows deriving arbitrarily large types of the form $T_N \vee T_M$, where the two parts can consist of arbitrary intersections and unions of various types that contain N and M somewhere within them.

For example, with the variables from Example[†] 10, we have $\Gamma; \rho \vdash w_2 : \{a : T_a\} \vee \mathcal{M}_\perp$, but $\Gamma; \rho \not\vdash w_2 : \{a : T_a\}$ and $\Gamma; \rho \not\vdash w_2 : \mathcal{M}_\perp$. Such a type cannot be derived in a syntactically bottom-up manner that invertible typing is based on. Rather than trying to fix invertible typing by adding complicated rules, we replace it by a new auxiliary typing judgment, which derives types in different way.

4.3 Layered Typing

In *layered typing*, we avoid the need for Lemma[†] 14 by organizing the derivation of a type not bottom-up, but by handling different type constructors in separate layers of typing judgments. All union type constructors are derived before intersection types, recursive types and type selections. Additionally, we derive union types on two layers:

First, the *basic layer* derives the newly top-like types possible by the rule ST-NM. Because intersections, recursive types and selections are out of the picture at this layer, these types have a simple form of possibly nested union types, where one of the sides contains N and the other \mathcal{M} , where \mathcal{M} is a declaration $\{M : \perp..T\}$ for some bound T . We will write that as $\vdash N \triangleleft T_N$ and $\vdash \mathcal{M} \triangleleft T_M$. Second, the *union layer* derives types possible by the rules ST-Or1 and ST-Or2, allowing nesting a known type of w in a union with any other type. This way, the layers retain the information about how a union type has been derived and those cases can be handled separately when inverting the derivation.

Intersection types can be handled in analogy to how any logical formula can be derived by starting from conjunctive normal form (CNF) and pushing conjunctions down. Any type constructed from a mixture of union and intersection types can be derived by starting from an intersection of union types and pushing the intersections down.

The *logic layer* sitting above the union layer can derive any mixture of unions and intersections using the LTL-And rule shown in Figure 7. It takes derivations of two types that may have some parts in common but differ in one place. The common part C^\vee is a syntactical context which combines the argument into a union with other types. For example, we can write the two types $\{a_1 : T_1\} \vee \{a_2 : T_2\}$ and $\{a_1 : T_1\} \vee \mathcal{M}_\perp$ as $C^\vee[\{a_2 : T_2\}]$ and $C^\vee[\mathcal{M}_\perp]$. If we view these two types as an intersection, then the rule pushes the intersection down to the place where the two types differ. In the example derivation on the right of Figure 6, we derived two union types on the union layer. (The type $\{a : T..T\}$ was

Typing layer	Relevant type constructors	Relevant rules
Atomic layer	$\{a : T..U\}, \{A : T..U\}, \{m(S, T) : U\}$	ST-Met, ST-Fld, ST-Typ
Basic layer	$N \vee M$	ST-NM
Union layer	$\top, T \vee U$	ST-Or1, ST-Or2, ST-Top
Logic layer	$T \wedge U$	ST-Dist, ST-And, VT-AndI
Main layer	$\mu(s : T), x.A$	VT-Recl, ST-SelL, ST-N-Rec

Table 1: The layers of layered typing

$$\frac{\Gamma; \rho \vdash_1 x : C^\vee[T_1] \quad \Gamma; \rho \vdash_1 x : C^\vee[T_2]}{\Gamma; \rho \vdash_1 x : C^\vee[T_1 \wedge T_2]} \text{(LTL-And)} \quad \frac{\Gamma; \rho \vdash_m x : C^{\wedge\vee}[[v_3/r]T_1] \quad \Gamma \vdash_1 v_2 : \{B(r) : T_1..T_2\}}{\Gamma; \rho \vdash_m x : C^{\wedge\vee}[v_2.B(v_3)]} \text{(LTM-Sel)}$$

Figure 7: Selected rules of layered typing

derived on the previous layer and S is an arbitrary type.) On the logic layer, we combined them into one type, pushing the intersection down to the left.

The rest of the type constructors are handled either below the basic layer or above the logic layer. Subtyping between declarations is handled in an *atomic layer* positioned before the basic layer. This layer only deals with types that are single declarations.

Recursive types of the form $\mu(s : T)$ and selection types can “wrap around” or replace any part of the derived type (in general typing by VT-Recl and ST-SelL, Figures 9 and 9 in the appendix), which may both appear under unions and intersections, and also contain them within. Therefore, they are handled above the logic layer in a final, *main layer*. In the rule LTM-Sel in Figure 7, the syntactic context $C^{\wedge\vee}$ can consist of a mixture of unions and intersections. The rules have premises that correspond to conditions in the relevant rules of tight typing. For example, in Figure 6, the left side of the union is wrapped under a recursive type in the last step.

The layers are summarized in Table 1, showing the relevant type constructors and the connection to rules of general typing. Selected rules are shown in Figure 7; full definitions are in Figures 15–19 in the appendix.

- Typing on the **atomic layer** ($\Gamma; \rho \vdash_a x : T$) only gives variables single declaration types – the declarations derived by precise typing, and their supertypes (subtyping rules between declarations are handled here).
- The **basic layer** ($\Gamma; \rho \vdash_b x : T$) additionally gives all variables top-like types of the form $T_N \vee T_M$ and $T_M \vee T_N$, where $\vdash N \triangleleft T_N$ and $\vdash M \triangleleft T_M$.
- The **union layer** ($\Gamma; \rho \vdash_u x : T$) handles \top and unions of known and arbitrary types.
- The **logic layer** ($\Gamma; \rho \vdash_l x : T$) handles intersections and distributivity. The rule LTL-And takes two types, preserves their common part, and combines the differing parts using an intersection type – pushing the intersection down from the top to its target place.
- The **main layer** ($\Gamma; \rho \vdash_m x : T$) closes recursive types and handles type selections.

For layered typing, we also proved the following properties:

- If a location has a declaration type by layered typing, then it also has a declaration type by precise typing, with the same or tighter bounds. This property has three variants, for field, type and method declarations.
- Layered typing is equivalent to general typing. As in the original proof, we use tight typing as a step between general and layered typing, and separately prove both directions of equivalence between tight and layered typing (Lemma 31 and Lemma 37 in the appendix).
- We also use layered typing to prove Lemma 9 – if a location has some read-only type in layered typing, then all references to that location have that type too.

With these properties, the safety proof from roDOT, with invertible typing replaced by the new layered typing definition, works as a safety proof of the updated calculus. Formal statements of these and other selected properties are given in Appendix A.3 in the appendix.

5 The SEF Guarantee

In Section 3.3, we informally stated the SEF guarantee, which provides the connection between a static typing condition (Definition 11) and run-time behavior of the method.

In this section, we present the formal definitions of the run-time SEF condition (Definition 15 in Section 5.1) and the SEF guarantee (Section 5.2). We then outline the proof of the guarantee (Section 5.3) and discuss some details of the proof (Section 5.4).

5.1 The run-time SEF condition

We informally stated the run-time SEF condition in Section 3.1, where we mentioned that several possible versions of such a condition could be defined. In our approach, we allow a pure method to create new objects and to modify just these new objects, which are under full control of the method.

The main SEF condition is that the method must not modify any **existing objects** that are already on the heap when the method starts executing. We can state such a condition in three variants, depending on the way in which it is checked that an object was not modified. Here we will use the variant that guarantees that existing objects on the heap do not change. In such a case, we say that a given execution of a method, starting from a method call start and reaching method call end in k steps, has the Sef-I property (Definition 15). The other possible variants are stated as Definition 42 and Definition 43 in the appendix.

Definition 15 (Sef-I). *A method execution $\langle w_1.m w_2; \sigma; \rho_1; \Sigma_1 \rangle \mapsto^k \langle v w_3; \sigma; \rho_2; \Sigma_2 \rangle$ is Sef-I when for every $j \leq k$ and $\langle w_1.m w_2; \sigma; \rho_1; \Sigma_1 \rangle \mapsto^j \langle t_3; \sigma_3; \rho_3; \Sigma_3 \rangle$, Σ_1 is a prefix of Σ_3 .*

5.1.1 Method call limits

Because we are defining a condition on what can happen while a method is executing, we need to understand what it means in roDOT that a method starts and ends its execution.

In roDOT, a method is called by a term $w_1.m w_2$. A *method call start* is a configuration of the form $\langle w_1.m w_2; \sigma; \rho_1; \Sigma_1 \rangle$, where w_1 is the receiver, m is the called method, w_2 is the argument, σ is the *continuation stack*, and Σ_1 is the *existing heap* (the environment ρ_1 does not have a special significance here).

The execution proceeds by replacing the call term $w_1.m w_2$ with the body of the method. Then, the body is executed. Unless there is an infinite loop, the body of the method will eventually evaluate to a single value. The machine will reach a configuration $\langle v w_3; \sigma; \rho_2; \Sigma_2 \rangle$, where w_3 is the result of the call and σ is the same stack as at the method call start.

The **first** such configuration after a method call start is the corresponding *method call end*. Another such configuration could possibly occur later in a completely unrelated way, but only the first such configuration is the method call end.

When a method call end is reached, the execution will either terminate, or proceed by popping a frame from the stack.

5.2 The SEF guarantee

The SEF guarantee, informally stated in Section 3.3, says that a SEF method does not modify existing objects in the heap during its execution. Theorem 16 is based on Definition 15, and speaks about the state of the heap at every point during the call. It is not the strongest possible purity guarantee, because this allows writing the value that already is in the field. On the other hand, it does not allow the value of fields to be changed and then changed back.

Theorem 16. *Let the configuration $c_1 := \langle w_1.m w_2; \sigma_1; \rho_1; \Sigma_1 \rangle$ be well-typed in a context Γ . Further assume that $\Gamma \vdash w_1 : \{m(z : \mathbb{N})(r : \mathbb{N}) : \top\}$. Then for any k steps of execution:*

1. Either the method call has finished executing. *There is $j < k$ for which $c_1 \mapsto^j \langle v w_3; \sigma_1; -; - \rangle$.*
2. Or, the method call has not finished executing and in this period existing objects in the heap are unchanged. *For each $c_1 \mapsto^k c_2$, all heap locations in c_1 also exist in c_2 and moreover they are unchanged in c_2 .*

5.3 Overview of the proof

The SEF guarantee talks about objects not being modified during the execution of methods, based on the mutability of method parameters. We base our proof of the SEF guarantee on the immutability guarantee (IG, Theorem 2), which states that individual objects can only be modified through mutable references. This guarantee was proven for roDOT [12] and is included in our mechanization in Coq.

However, the immutability guarantee cannot be applied at the start of the method, because there may be many mutable references to objects on the heap. Also, IG guarantees immutability until the end of execution of the whole program, but the SEF guarantee only until the end of the method.

These differences can be bridged by taking the machine configuration at the method start, and constructing a different configuration that will execute the same way until the end of the method, but removing the parts that prevent the IG from applying.

First, note that the stack is not relevant to how the method executes and stays the same from the method start until its end. We therefore remove this stack entirely, and get an execution isolated from the rest of the program. This execution proceeds through the same steps, but stops at the method end. By removing the stack, we rid the configuration of any references to objects that might be used after the method call returns. If we apply the IG to this configuration, it will guarantee that objects are not modified until the end of the method, exactly what is needed for the SEF guarantee.

Removing the stack is not enough for the IG to apply though, because a SEF method can be called with arguments that are mutable references. We do not want to prevent that from happening, because even when a method is SEF, it can be useful to pass mutable references to the method and have it return one of these references with its mutability intact.

What is special about a SEF method is that (because of its declared parameter types) it cannot use the mutability during its execution. Therefore, when called with mutable arguments, it should execute in exactly the same way as if called with read-only arguments.

So the second modification to the configuration after removing the stack is to change the mutability of the arguments to read-only. That way, the alternative configuration contains no writeable references, and IG guarantees that no objects that were on the heap at the start will be modified. Still, this alternative configuration executes the same steps as the original, meaning the original method execution also does not modify any existing object on the heap.

5.4 Proof of the SEF Guarantee

The strategy of the proof of Theorem 16 is to focus on the second case of the SEF guarantee by using the immutability guarantee to show the theorem for a configuration c_2 obtained by temporarily truncating the stack of c_1 from Theorem 16.

Lemma 17 (SEF guarantee without stack). *For c_1 satisfying the conditions of Theorem 16, let $c'_1 := \langle w_1.m w_2; \cdot; \rho_1; \Sigma_1 \rangle$. If $c'_1 \mapsto^n c'_2$ for some n and c'_2 , then the heap of c'_2 contains all objects of c'_1 without modification.*

It is easy to prove the full SEF theorem with this result for c'_1 . The premise of the immutability guarantee is that c'_1 is well-typed in some context Γ_2 and there are no mutably reachable objects in c'_1 with respect to the typing of Γ_2 . Clearly c'_1 is well typed in the original context Γ . But for the part about mutably reachable objects, we cannot just take Γ as Γ_2 because for this, Γ_2 must assign read-only types to w_i . Even though we have $(r : \mathbb{N})$ in the typing $\Gamma \vdash w_1 : \{m(z : \mathbb{N})(r : \mathbb{N}) : \top\}$, this does not necessarily mean $\Gamma(w_1)$ is a read-only type. For example, w_1 might be mutable in Γ but m might not make use of its mutability. Therefore, instead of using Γ as Γ_2 , we construct Γ_2 and show that c'_1 is well-typed in Γ_2 like so:

1. *Reference elimination.* Remove bindings for w_i from Γ and replace all occurrences of w_i with the corresponding location y_i in both Γ and c'_1 .
2. *Read-only weakening.* Add back bindings for w_i , where the new type bound to w_i is the type bound to y_i except with the mutable part set to read-only.

We do this instead of changing the types of w_i , because we only know that w_i are used in a read-only way in the focus, while on the heap, w_i might be used as a part of dependent types referring to their mutability. Changing their types would break the typing of the heap.

The second step is essential, because only references can be read-only, while locations always have mutable types. The references w_i are added in the same order as in the original context, to ensure that

types in the typing context only refer to preceding variables in case the types are dependent. The two steps above correspond to the following two lemmas.

Lemma 18 (Reference elimination). *Let c_1 and Γ satisfy the conditions of Theorem 16, and $\rho_1[w_i] = y_i$. Define Γ' as the context obtained from Γ by first removing bindings for w_i and then replacing w_i with y_i . Define ρ' as the environment obtained from ρ_1 by removing bindings for w_i . Then the term $y_1.m y_2$ is well-typed under $\Gamma';\rho'$ and we have heap correspondence $\Gamma';\rho' \sim [y_i/w_i]_i \Sigma_1$.*

Proof. Because w_i is a reference to y_i , types assigned to y_i and w_i by Γ differ only by mutability, and y_i has a mutable type. So $\Gamma(y_i)$ is a subtype of $\Gamma(w_i)$, and the result follows by substitutivity. \square

Lemma 19 (Read-only weakening). *Let c_1 and Γ satisfy the conditions of Theorem 16, and y_1, y_2 be such that $\rho_1[w_i] = y_i$. Then there is a context Γ_2 binding w_i to read-only types such that the configuration $c'_1 := \langle w_1.m w_2; \cdot; \rho_1; [y_i/w_i]_i \Sigma_1 \rangle$, is well-typed in Γ_2 (formally $\Gamma_2 \vdash c'_1 : \top$).*

By Lemma 19 along with the immutability guarantee, we have SEF established for $c'_1 := \langle w_1.m w_2; \cdot; \rho_1; [y_i/w_i]_i \Sigma_1 \rangle$. However, we need SEF in particular for the configuration $c'_1 := \langle w_1.m w_2; \cdot; \rho_1; \Sigma_1 \rangle$ in Lemma 17, where there is no substitution $[y_i/w_i]_i$ in the heap. Nevertheless, this substitution can be ignored in the sense that execution can only change fields of objects, but in roDOT, fields always store locations while w_i are references. (When a reference is assigned to a field, the corresponding location is stored, in order to make the field type determine the mutability of the stored value.) Because c'_1 and c''_1 are the same everywhere except for $[y_i/w_i]_i$ on the heap, the SEF property of c''_1 can be carried over to c'_1 . The first part of carrying the SEF property over to c'_1 is to relate each k -th step of execution starting from c''_1 and the k -th step of execution starting from c'_1 . We need to show that the two executions are almost the same, except some specific variables that appear in the machine configuration can differ. In particular, the references w_i can be replaced by the locations y_i . Additionally, the locations and references of newly created objects can differ, because variable names are not assigned deterministically.

For that reason, we define the *similarity* relation, which formalizes structural equivalence of syntactic elements such as terms, objects or whole configurations that differ only in names of variables. It relates two such elements using a *renaming* relation over variables. A formal definition of similarity and its basic properties is given in Appendix A.5.

Similarity has two important properties with respect to program execution: (1) it is preserved by reduction, and (2) if a machine configuration can reduce, then all similar configurations can reduce too (and the results are similar). That means that if we start with two similar configurations, where the execution of one reaches an answer state, then the execution of the other will reach a similar answer state. With this definition of similarity, Lemma 20 formalises the idea that c''_1 is similar to c'_1 up to renaming w_i to y_i :

Lemma 20 (Similarity for eliminated references). *Let c'_1 satisfy the conditions of Lemma 17 and c''_1 the conditions of Lemma 19. Then $c'_1 \stackrel{(w_1, y_1), (w_2, y_2)}{\approx} c''_1$.*

The final part of carrying over the SEF property to c'_1 is to recognize that reduction only changes values of fields (while the structure of the object, methods and type members are immutable).

Definition 21 (Objects identical except fields). *For objects o_1 and o_2 , we write $o_1 \stackrel{fid}{\approx} o_2$ to mean they are identical except for possibly the values of fields.*

Lemma 22 (Reduction only changes fields). *If $\langle -; -; -; \Sigma \rangle \mapsto^n \langle -; -; -; \Sigma' \rangle$ and y is a location in Σ , then $\Sigma(y) \stackrel{fid}{\approx} \Sigma'(y)$.*

And with this, we can finish the proof of Theorem 16.

Proof. By classical reasoning, assume the condition 1 is false so that the goal is to prove the condition 2. That is, assume that there is no $j < k$ such that the top-most frame of c_1 is popped after execution by j steps: $c_1 \mapsto^j \langle \mathbf{v}w_3; -; -; \Sigma_2 \rangle$. Then, the sequence of reductions $c_1 \mapsto \dots \mapsto c_2$ corresponds to a sequence of reductions $c'_1 \mapsto \dots \mapsto c'_2$ because even though c'_1 has no awaiting frames, there are no frame pops in this execution sequence by the current assumption. By Lemma 17, the condition 2 follows. \square

6 Transformations

In Section 3.4, we informally stated the transformation guarantee, which connects the static SEF condition with a practical application – that calls to SEF methods can be safely swapped.

Defining the call-swapping and the guarantee in a formal way requires dealing with several technicalities particular to the roDOT calculus (or DOT calculi in general). In order to separate the common problems from the specific case of swapping calls, we build a general framework in which various transformations of roDOT programs can be defined and proven safe. We instantiate it here only with the call swapping transformation, but it could represent the general part of a proof of safety for other transformations, such as reordering field reads or removing dead code.

In Section 6.1, we present the framework for defining and reasoning about safe program transformations in roDOT and similar calculi, including a general Theorem 25 about safety of transformations. In Section 6.2, we define the transformation that swaps two calls to methods that are statically determined to be side-effect free (Definition 26) and state the transformation guarantee.

6.1 Transformation framework

The framework defines a general form for roDOT program transformations, defines the precise meaning of a *safe transformation* that “does not affect the behavior of the program”, and provides a way of proving this property, while helping to deal with common technical issues of DOT formalizations.

Our general approach is to define a transformation that applies to an initial program, and prove it safe by showing that if the original and transformed program are executed side-by-side, they will either eventually reach the “same” answer, or both not terminate.

In the initial program, a transformation, such as swapping calls, can be located anywhere, including inside a body of a method of an object literal, such as $\text{let } x = \nu(r : R)\{m(r, z) = t_m\}$ in t_2 . We must consider that in roDOT, terms are in A-normal form, and that a term can have multiple different types. Also code (terms) and values (objects) are mixed with each other during execution of a program – the program contains object literals, and methods on the heap contain program code.

During execution, objects are created on the heap, including copies of the code of their methods, which can be affected by the transformation. It would be too restrictive to require that a transformed program produce the exact same output value as the original program since the output value may be an object that may contain the transformed code. To facilitate this, we define each transformation using a local relation which relates two terms that differ only locally, and the framework provides lifting operators, which allow this transformation to occur anywhere in a program or in a machine configuration.

The general safety Theorem 25 is based on executing the two programs and observing that the intermediate states are also related by the transformation (lifted to whole configurations and allowed to occur at multiple places), except the moments when the directly affected part of the program is executing. When the two answers are reached, they will be similar except that bodies of methods on the heap may differ as the transformation permits.

A more detailed description of the framework design and its definitions are given in Appendix A.7. The following text describes the most important parts.

6.1.1 Transformations of roDOT programs in general

A transformation of a program is defined as a binary relation on terms – the original program and the transformed one. For example, the call-swapping transformation is defined as a relation that relates a program containing two method calls with a program that only differs in the order of those calls.

Because the safety of the transformation depends on typing information, we define the transformation as a binary relation between triples: the term, its type and a typing context. This generalizes to other syntactic elements – stacks, heaps and machine configurations, though for each kind of element, the exact meaning of “typing context” and “type” may differ. For terms, the typing context is actually paired with a runtime environment $\Gamma; \rho$.

Definition 23 (Transformation). *A transformation τ is a relation between triples consisting of typing contexts $\Gamma_{1,2}$, types $T_{1,2}$ and typeable elements $X_{1,2}$. We write $\langle \Gamma_1 \vdash X_1 : T_1 \rangle \rightarrow_\tau \langle \Gamma_2 \vdash X_2 : T_2 \rangle$ and say that X_1 is transformed into X_2 .*

Like any binary relation, transformations can be symmetric, reflexive or transitive, and we can construct transformations using iteration, composition, union and inversion.

Additionally, a transformation is *type-safe*, if the syntactic elements on both sides are correctly typed under the respective contexts. Another useful property of a transformation is being *type-identical*, where both the types and typing contexts are the same on both sides.

To facilitate the possibility of the transformation being located anywhere in a term, it is useful to define the transformation in two steps: (1) A *local transformation*, which only allows swapping calls at the root of the term. (2) A lifting operator $\text{lift } \tau$ which takes a local transformation τ and allows it to be located at one place anywhere in a term.

Such a local transformation of a term can be further lifted by $\text{cfg } \tau$ to a whole run-time configuration, where τ applies at exactly one place in the focus of execution, in the stack or on the heap. To allow multiple occurrences, we can apply the iteration operator to the lifted transformation. Having a definition that only allows one occurrence is useful in the proof of Theorem 25 in Appendix A.7.4, where we want to look at each occurrence individually. More details about lifting are in Appendix A.7.3; the definitions of the lifting operators are shown in Figure 23 and Figure 24.

6.1.2 Safe transformations

The definition of a safe transformation must allow the answers of the two programs to contain different variable names and allow for the fact that the transformation can occur in the heap of the program answer, possibly at multiple places. Therefore a local transformation is safe if execution of the transformed program reaches answers related by an iteration of this transformation. To deal with non-deterministic location names, the transformation is in union with a *similarity transformation* \approx , which allows one-to-one variable renaming.

Definition 24 (Safe transformation). *A transformation τ is safe if $\langle \Gamma_1 \vdash c_1 : T \rangle \rightarrow_{\tau} \langle \Gamma_2 \vdash c_2 : T \rangle$ and $c_1 \rightarrow^k c_3$, where c_3 is an answer typed as $\Gamma_3 \vdash c_3 : T$, implies that there exist c_4, Γ_4 and j such that $c_2 \rightarrow^j c_4$, $\Gamma_4 \vdash c_4 : T$ and $\langle \Gamma_3 \vdash c_3 : T \rangle \rightarrow_{(\tau \cup \approx)^*} \langle \Gamma_4 \vdash c_4 : T \rangle$.*

Thanks to being able to define a transformation by applying a general lifting to a local transformation, the safety proof of such a transformation can be also divided into a theorem that will apply to any local transformation with certain local properties, and then proving those local properties for the particular local transformation.

This approach makes it possible to state the call-swapping guarantee presented here (Theorem 27), or analogous guarantees for other local transformations.

Theorem 25 states that if a local term transformation does not change typing of the term, is compatible with properties such as weakening, narrowing and substitution, does not change whether the term is an answer or not, and if execution of just the transformed term will eventually reach similar configurations, then transforming a program by this transformation anywhere will not change its result. Full definitions of the premises are in Appendix A.7 as Definitions 49, 50, 52, 56 and 54.

Theorem 25 (General safety for local transformations). *If τ is a transformation that is type-identical, type-safe, compatible with weakening, narrowing and substitution, preserves answers, and eventually reduces to similarity, then $(\text{cfg } \tau \cup \approx)^*$ is safe.*

6.2 The call-swapping transformation

The specific transformation guarantee that we want to achieve should state that swapping two calls will not change the outcome of the program, in the sense of Definition 24.

Call-swapping is defined as a local transformation that transforms one program containing two successive calls into another program in which the calls are swapped. Due to the A-normal form of terms, two successive calls in the program have the form $\text{let } x_{c1} = x_{o1}.m_1x_{a1}$ in $\text{let } x_{c2} = x_{o2}.m_2x_{a2}$ in t . In the transformation, the two calls $x_{o1}.m_1x_{a1}$ and $x_{o2}.m_2x_{a2}$ appear in the opposite order, but the continuation t is the same.

The transformation is only safe if both the methods are SEF, so it has several typing premises, analogous to the ones of Theorem 16 in Section 5.2.

Definition 26 (Local call swapping). *The local call-swapping transformation csw is a transformation of terms that relates $\langle \Gamma \vdash \text{let } x_{c1} = x_{o1}.m_1x_{a1}$ in $\text{let } x_{c2} = x_{o2}.m_2x_{a2}$ in $t : T \rangle \rightarrow_{\text{csw}} \langle \Gamma \vdash \text{let } x_{c2} = x_{o2}.m_2x_{a2}$ in $\text{let } x_{c1} = x_{o1}.m_1x_{a1}$ in $t : T \rangle$ when*

- $x_{c1,2}$ are distinct from $x_{a1,2}$ and $x_{o1,2}$,
- $\Gamma \vdash x_{o1}.m_1x_{a1} : T_{c1}, \Gamma \vdash x_{o2}.m_2x_{a2} : T_{c2}$, and $\Gamma, x_{c1} : T_{c1}, x_{c2} : T_{c2} \vdash t : T$,
- $\Gamma \vdash x_{o1} : \{m_1(r_1 : \mathbb{N}, z_1 : T_{a1}) : \top\}$, and $\Gamma \vdash x_{o2} : \{m_2(r_2 : \mathbb{N}, z_2 : T_{a2}) : \top\}$,
- $\Gamma \vdash x_{a1} : T_{a1}$, and $\Gamma \vdash x_{a2} : T_{a2}$,
- $\Gamma \vdash \mathbb{N} <: T_{a1}$, and $\Gamma \vdash \mathbb{N} <: T_{a2}$.

As the the final form of the transformation guarantee, we apply Definition 24 to Definition 26, and specialize the theorem to initial programs. The proof is given in Appendix A.8.

Theorem 27 (Transformation guarantee). *If $\langle \vdash t_1 : T \rangle \rightarrow_{\text{lift } csw} \langle \vdash t_2 : T \rangle$ and $\langle t_1; \cdot; \cdot; \cdot \rangle \rightarrow^k c_3$, where c_3 is an answer typed as $\Gamma_3 \vdash c_3 : T$, then there exists c_4, Γ_4 and j such that $\langle t_2; \cdot; \cdot; \cdot \rangle \rightarrow^j c_4$, c_4 is an answer typed as $\Gamma_4 \vdash c_4 : T$ and $\langle \Gamma_3 \vdash c_3 : T \rangle \rightarrow_{(cfg \ csw \cup \approx)^*} \langle \Gamma_4 \vdash c_4 : T \rangle$.*

7 Related work

Since the topic of this work includes both the DOT calculus and method purity, here we discuss previous work related to these concepts. Prior to this work, many variants of the DOT calculus were published, some including mechanized proofs. Also the issue of purity in object-oriented languages is of great research interest, and it is approached from different angles of automation and precision. We give details about the existing work in the following subsections. As far as we know, our work is the first one to consider the issue of purity within a DOT calculus.

7.1 Mechanizations of DOT calculi

The first appearance of a DOT calculus [3] did not include a proof of soundness, but was followed by several versions with proofs in Coq [33, 30] and Iris [17]. In particular, WadlerFest DOT [2], thanks to its simplicity and its proof of soundness based on invertible typing [30], was used as a baseline for numerous extensions [32, 22, 31, 23], including roDOT. While objects are immutable in WadlerFest DOT, it was extended to support mutation using mutable slots in Mutable WadlerFest DOT [32], and more directly by allowing changing values of fields in kDOT [22]. A simplified version [21] with mutable fields, but without the specific kDOT feature of constructors, was used as a base for the mechanization of roDOT.

The differences between the mechanization of roDOT and those of previous DOT calculi mainly stem from the differences in how roDOT handles variables – namely, typing of variables and terms being separated from each other, using different definitions of typing contexts to support variable hiding, using the runtime environment to map references to locations, and using typing information in its definition of operational semantics.

The mechanization of roDOT includes a feature to ease further extensions to the calculus. The definitions and theorems are parameterized by a “typing mode”, which allows selecting type system features that are supported. Using this feature, our proofs work for roDOT both with and without the changes described in this paper.

7.2 Purity in other languages

Purity in programming is such an important concept that in many languages, functions are pure by default. This approach is typically associated with functional programming, but an object-oriented system can also be pure [1] when the objects are immutable. That is also the case in the basic DOT calculus.

In pure functional languages, effects must typically be explicitly declared in the program using monadic types. This style of programming has been shown to be as powerful as other styles and is used in practical programming languages such as Haskell.

Regarding purity in object oriented languages with mutable fields, many publications [37, 34, 38, 40, 6, 16, 29] focus on Java and languages with similar type systems, such as C#. For Scala, a type system for purity was developed, but not based on the DOT calculus [35].

When approached from a practical standpoint, the definition of purity in these languages has to include considerations other than modification of object fields, such as accessing global variables or

synchronization. This leads to different definitions of purity. The term “pure” is sometimes used to mean the same as “side-effect-free”, without requiring determinism.

Observational purity [25, 5] is a weaker property that allows side effects as long as they are not observable from certain parts of the code. This definition is based on classes and access control, features which are not modeled in DOT calculi.

Purity is of great use to program verification and specification frameworks, where it enables inserting run-time checks without changing behavior, and allows more precise analysis. Code Contracts [15], JML [20] and Checker Framework [14] allow annotating a method as pure. Code Contracts do not check that this annotation is correctly applied, and JML and Checker Framework use simpler checks, where pure methods are not allowed to call impure methods. Checker Framework uses the fact that side-effect free methods do not invalidate flow-sensitive types of local variables.

To avoid imposing an annotation burden on the programmer, purity can be inferred by automatic program analysis [26, 34], and side-effect analysis can be used for program optimization [11].

ReIm [19] provides both a type system for reference mutability and a way to automatically infer mutability types. It can therefore automatically find pure methods, which have all parameters read-only. We adopted this way of recognizing pure methods by parameter types for roDOT in this work. While in ReIm, mutability is attached to parameter types as a qualifier in the style of the Checker Framework, roDOT uses the special member type `M` to include the mutability in the parameter type using intersection types. In ReIm, mutability qualifiers are subjected to qualifier polymorphism and viewpoint adaptation. roDOT can express the equivalent of polymorphic qualifiers using dependent types and implements viewpoint adaptation using union and intersection types [12].

7.3 Capability and Effect Systems

There are other ways to express the permitted side-effects of functions using types, which have been developed in recent work on formal type systems.

The principle of **capabilities** [28, 27] is to require every operation that can have a side effect to take an extra value, called a capability, as a parameter. Then, if some function or method does not have the capability value corresponding to a particular effect, we can conclude that it does not perform that effect. Capabilities are well suited for coarse-grained effects, such as performing input/output in general or accessing some specific file, where a single capability value can guard a set of related operations. To apply such an approach to reasoning about a fine-grained effect such as writing to a field of a specific object, we would need large numbers of such capability values, one new capability value for each existing object. For each reference passed to a parameter or stored in a field, a corresponding capability would need to be passed or stored, thus multiplying the number of parameters and fields.

Wyvern’s effect system [24] expresses possible effects by type members of objects. That is syntactically similar to how roDOT represents mutability, but the meaning of the type members is different. In roDOT, the type member of an object reference defines the bounds on the mutability of the reference, the knowledge about whether a reference may be used for mutation, in the type of that reference. In contrast, in Wyvern, the effect member represents a permission to perform an effect, such as `file.Write`, where the effect can be independent of the object that contains the effect member. Thus, Wyvern effect members are more similar to the capability-based approach.

Another successful direction is to use types to express sets of possible variables captured or aliased by values in the program. **Capture Types** [7] follow from a capability based approach, and enable reasoning about where capability values may be stored in the heap or captured in closures, in order to more precisely reason about where effects may occur. **Reachability Types** [4] annotate the type of an expression with a set of variables, which are values that are possibly reachable from the result of that expression. This can be used in conjunction with effect qualifiers as in Graph IR [8], where a function type declares a set of variables that can be read or written, describing the possible effects in a fine-grained way. The types can also be extended to support qualifier polymorphism [39]. This work is defined in the context of a higher order functional formalism, whereas roDOT is an object-oriented calculus. Also, both Wyvern and Reachability Types express effects using new constructs added to the type system, while roDOT aims to encode mutability using the existing DOT constructs of dependent types, unions and intersections.

8 Conclusion

To conclude, our paper confirms that the reference mutability system provided by roDOT can be mechanically proven sound, and with a few changes can be used to guarantee side-effect freedom of methods, and to justify safe transformations of programs.

References

- [1] Martín Abadi and Luca Cardelli. *A Theory of Objects*. Monographs in Computer Science. Springer, 1996. doi:10.1007/978-1-4419-8598-9.
- [2] Nada Amin, Samuel Grütter, Martin Odersky, Tiark Rompf, and Sandro Stucki. The essence of dependent object types. In Sam Lindley, Conor McBride, Philip W. Trinder, and Donald Sannella, editors, *A List of Successes That Can Change the World - Essays Dedicated to Philip Wadler on the Occasion of His 60th Birthday*, volume 9600 of *Lecture Notes in Computer Science*, pages 249–272. Springer, 2016. doi:10.1007/978-3-319-30936-1_14.
- [3] Nada Amin, Tiark Rompf, and Martin Odersky. Foundations of path-dependent types. In Andrew P. Black and Todd D. Millstein, editors, *Proceedings of the 2014 ACM International Conference on Object Oriented Programming Systems Languages & Applications, OOPSLA 2014, part of SPLASH 2014, Portland, OR, USA, October 20-24, 2014, OOPSLA '14*, pages 233–249. ACM, 2014. doi:10.1145/2660193.2660216.
- [4] Yuyan Bao, Guannan Wei, Oliver Bračevac, Yuxuan Jiang, Qiyang He, and Tiark Rompf. Reachability types: Tracking aliasing and separation in higher-order functional programs. *Proc. ACM Program. Lang.*, 5(OOPSLA), oct 2021. doi:10.1145/3485516.
- [5] Mike Barnett, David A Naumann, Wolfram Schulte, and Qi Sun. 99.44% pure: Useful abstractions in specifications. In *ECOOP workshop on formal techniques for Java-like programs (FTJFP)*, 2004.
- [6] William C. Benton and Charles N. Fischer. Mostly-functional behavior in Java programs. In Neil D. Jones and Markus Müller-Olm, editors, *Verification, Model Checking, and Abstract Interpretation, 10th International Conference, VMCAI 2009, Savannah, GA, USA, January 18-20, 2009. Proceedings*, volume 5403 of *Lecture Notes in Computer Science*, pages 29–43. Springer, 2009. doi:10.1007/978-3-540-93900-9_7.
- [7] Aleksander Boruch-Gruszecki, Martin Odersky, Edward Lee, Ondřej Lhoták, and Jonathan Brachthäuser. Capturing types. *ACM Trans. Program. Lang. Syst.*, 45(4), nov 2023. doi:10.1145/3618003.
- [8] Oliver Bračevac, Guannan Wei, Songlin Jia, Supun Abeysinghe, Yuxuan Jiang, Yuyan Bao, and Tiark Rompf. Graph IRs for impure higher-order languages: Making aggressive optimizations affordable with precise effect dependencies. *Proc. ACM Program. Lang.*, 7(OOPSLA2), oct 2023. doi:10.1145/3622813.
- [9] The Checker Framework Manual: Custom pluggable types for Java. <https://checkerframework.org/manual/#initialization-checker>, 2022.
- [10] The Checker Framework Manual: Custom pluggable types for Java. <https://checkerframework.org/manual/#purity-checker>, 2022.
- [11] Lars Ræder Clausen. A Java bytecode optimizer using side-effect analysis. *Concurrency: Practice and Experience*, 9(11):1031–1045, 1997. doi:10.1002/(SICI)1096-9128(199711)9:11<1031::AID-CPE354>3.0.CO;2-0.
- [12] Vlastimil Dort and Ondřej Lhoták. Reference mutability for DOT. In Robert Hirschfeld and Tobias Pape, editors, *34th European Conference on Object-Oriented Programming, ECOOP 2020, November 15-17, 2020, Berlin, Germany (Virtual Conference)*, volume 166 of *LIPICs*, pages 18:1–18:28. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2020. doi:10.4230/LIPICs.ECOOP.2020.18.
- [13] Vlastimil Dort, Yufeng Li, Ondřej Lhoták, and Pavel Parížek. Pure methods for roDOT. In *ECOOP 2024, LIPICs*. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2024 (forthcoming).

- [14] Michael D. Ernst. Annotation type Pure. <https://checkerframework.org/api/org/checkerframework/dataflow/qual/Pure.html>, 2022.
- [15] Manuel Fähndrich, Michael Barnett, and Francesco Logozzo. Embedded contract languages. In Sung Y. Shin, Sascha Ossowski, Michael Schumacher, Mathew J. Palakal, and Chih-Cheng Hung, editors, *Proceedings of the 2010 ACM Symposium on Applied Computing (SAC), Sierre, Switzerland, March 22-26, 2010*, pages 2103–2110. ACM, 2010. doi:10.1145/1774088.1774531.
- [16] Matthew Finifter, Adrian Mettler, Naveen Sastry, and David A. Wagner. Verifiable functional purity in Java. In Peng Ning, Paul F. Syverson, and Somesh Jha, editors, *Proceedings of the 2008 ACM Conference on Computer and Communications Security, CCS 2008, Alexandria, Virginia, USA, October 27-31, 2008*, pages 161–174. ACM, 2008. doi:10.1145/1455770.1455793.
- [17] Paolo G. Giarrusso, Léo Stefanescu, Amin Timany, Lars Birkedal, and Robbert Krebbers. Scala step-by-step: soundness for DOT with step-indexed logical relations in Iris. *Proc. ACM Program. Lang.*, 4(ICFP):114:1–114:29, 2020. doi:10.1145/3408996.
- [18] James Gosling, Bill Joy, Guy Steele, Gilad Bracha, and Alex Buckley. The Java® language specification, Java SE 8 edition. <https://docs.oracle.com/javase/specs/jls/se8/html/jls-8.html#jls-8.4.1>, 2022.
- [19] Wei Huang, Ana Milanova, Werner Dietl, and Michael D. Ernst. ReIm & ReImInfer: checking and inference of reference immutability and method purity. In *Proceedings of the 27th Annual ACM SIGPLAN Conference on Object-Oriented Programming, Systems, Languages, and Applications, OOPSLA 2012, part of SPLASH 2012, Tucson, AZ, USA, October 21-25, 2012, OOPSLA '12*, pages 879–896. Association for Computing Machinery, 2012. doi:10.1145/2384616.2384680.
- [20] JML reference manual: Class and interface member declarations. https://www.cs.ucf.edu/~leavens/JML/jmlrefman/jmlrefman_7.html#SEC60, 2022.
- [21] Ifaz Kabir. themaplelab / dot-public: A simpler syntactic soundness proof for dependent object types. <https://github.com/themaplelab/dot-public/tree/master/dot-simpler>.
- [22] Ifaz Kabir and Ondřej Lhoták. κ DOT: scaling DOT with mutation and constructors. In *Proceedings of the 9th ACM SIGPLAN International Symposium on Scala, SCALA@ICFP 2018, St. Louis, MO, USA, September 28, 2018*, pages 40–50, 2018. doi:10.1145/3241653.3241659.
- [23] Ifaz Kabir, Yufeng Li, and Ondrej Lhoták. ι DOT: a DOT calculus with object initialization. *Proc. ACM Program. Lang.*, 4(OOPSLA):208:1–208:28, 2020. doi:10.1145/3428276.
- [24] Darya Melicher, Anlun Xu, Valerie Zhao, Alex Potanin, and Jonathan Aldrich. Bounded abstract effects. *ACM Trans. Program. Lang. Syst.*, 44(1), jan 2022. doi:10.1145/3492427.
- [25] David A. Naumann. Observational purity and encapsulation. In Maura Cerioli, editor, *Fundamental Approaches to Software Engineering, 8th International Conference, FASE 2005, Held as Part of the Joint European Conferences on Theory and Practice of Software, ETAPS 2005, Edinburgh, UK, April 4-8, 2005, Proceedings*, volume 3442 of *Lecture Notes in Computer Science*, pages 190–204. Springer, 2005. doi:10.1007/978-3-540-31984-9_15.
- [26] Jens Nicolay, Quentin Stiévenart, Wolfgang De Meuter, and Coen De Roover. Purity analysis for JavaScript through abstract interpretation. *Journal of Software: Evolution and Process*, 29(12), 2017. doi:10.1002/smr.1889.
- [27] Martin Odersky, Aleksander Boruch-Gruszecki, Jonathan Immanuel Brachthäuser, Edward Lee, and Ondřej Lhoták. Safer exceptions for Scala. In *Proceedings of the 12th ACM SIGPLAN International Symposium on Scala, SCALA 2021*, page 1–11, New York, NY, USA, 2021. Association for Computing Machinery. doi:10.1145/3486610.3486893.
- [28] Martin Odersky, Aleksander Boruch-Gruszecki, Edward Lee, Jonathan Brachthäuser, and Ondřej Lhoták. Scoped capabilities for polymorphic effects, 2022. arXiv:2207.03402.

- [29] David J. Pearce. JPure: A modular purity system for Java. In Jens Knoop, editor, *Compiler Construction - 20th International Conference, CC 2011, Held as Part of the Joint European Conferences on Theory and Practice of Software, ETAPS 2011, Saarbrücken, Germany, March 26–April 3, 2011. Proceedings*, volume 6601 of *Lecture Notes in Computer Science*, pages 104–123. Springer, 2011. doi:10.1007/978-3-642-19861-8_7.
- [30] Marianna Rapoport, Ifaz Kabir, Paul He, and Ondřej Lhoták. A simple soundness proof for dependent object types. *Proc. ACM Program. Lang.*, 1(OOPSLA):46:1–46:27, 2017. doi:10.1145/3133870.
- [31] Marianna Rapoport and Ondřej Lhoták. A path to DOT: formalizing fully path-dependent types. *Proc. ACM Program. Lang.*, 3(OOPSLA):145:1–145:29, 2019. doi:10.1145/3360571.
- [32] Marianna Rapoport and Ondřej Lhoták. Mutable WadlerFest DOT. In *Proceedings of the 19th Workshop on Formal Techniques for Java-like Programs, Barcelona, Spain, June 20, 2017*, pages 7:1–7:6. ACM Press, 2017. doi:10.1145/3103111.3104036.
- [33] Tiark Rompf and Nada Amin. Type soundness for dependent object types (DOT). In Eelco Visser and Yannis Smaragdakis, editors, *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, OOPSLA '16*, pages 624–641. ACM, 2016. doi:10.1145/2983990.2984008.
- [34] Atanas Rountev. Precise identification of side-effect-free methods in Java. In *20th International Conference on Software Maintenance (ICSM 2004), 11-17 September 2004, Chicago, IL, USA*, pages 82–91. IEEE Computer Society, 2004. doi:10.1109/ICSM.2004.1357793.
- [35] Lukas Rytz, Nada Amin, and Martin Odersky. A flow-insensitive, modular effect system for purity. In Werner Dietl, editor, *Proceedings of the 15th Workshop on Formal Techniques for Java-like Programs, FTfJP 2013, Montpellier, France, July 1, 2013, FTfJP '13*, pages 4:1–4:7. ACM, 2013. doi:10.1145/2489804.2489808.
- [36] Amr Sabry and Matthias Felleisen. Reasoning about programs in continuation-passing style. In *Proceedings of the 1992 ACM Conference on LISP and Functional Programming, LFP '92*, page 288–298. New York, NY, USA, 1992. Association for Computing Machinery. doi:10.1145/141471.141563.
- [37] Alexandru Salcianu and Martin Rinard. A combined pointer and purity analysis for Java programs. Technical report, Massachusetts Institute of Technology Computer Science and Artificial Intelligence Laboratory, 2004. URL: <https://dspace.mit.edu/handle/1721.1/30470>.
- [38] Alexandru Salcianu and Martin C. Rinard. Purity and side effect analysis for Java programs. In Radhia Cousot, editor, *Verification, Model Checking, and Abstract Interpretation, 6th International Conference, VMCAI 2005, Paris, France, January 17-19, 2005, Proceedings*, volume 3385 of *Lecture Notes in Computer Science*, pages 199–215. Springer, 2005. doi:10.1007/978-3-540-30579-8_14.
- [39] Guannan Wei, Oliver Bračevac, Songlin Jia, Yuyan Bao, and Tiark Rompf. Polymorphic reachability types: Tracking freshness, aliasing, and separation in higher-order generic programs. *Proc. ACM Program. Lang.*, 8(POPL), jan 2024. doi:10.1145/3632856.
- [40] Haiying Xu, Christopher J. F. Pickett, and Clark Verbrugge. Dynamic purity analysis for Java programs. In Manuvir Das and Dan Grossman, editors, *Proceedings of the 7th ACM SIGPLAN-SIGSOFT Workshop on Program Analysis for Software Tools and Engineering, PASTE'07, San Diego, California, USA, June 13-14, 2007*, pages 75–82. ACM, 2007. doi:10.1145/1251535.1251548.

A Appendix – roDOT definitions

This appendix contains full versions of definitions, which are partially presented in the paper.

The appendix is organized as follows:

- In Appendix A.1, we include full definitions of general typing and subtyping in roDOT.
- Appendix A.2 contains full definitions of layered typing.
- We show several additional properties of layered typing, which are used in the soundness proof in Appendix A.3.
- In Appendix A.4 we show different side-effect-freedom conditions as alternatives to Definition 11 in Section 4.1.
- A more detailed discussion and formal definition of similarity (informally discussed in Section 5.3) is given in Appendix A.5.
- Appendix A.6 contains proofs of lemmas from Section 5.
- Appendix A.7 contains more detailed discussion of the transformation framework.
- The proof of the transformation guarantee (Theorem 27) is given in Appendix A.8.

A.1 Typing and reduction definitions

This section contains full definitions of typing and operational semantics of roDOT, which are partially presented in Figure 2 in the main paper. Changes to definitions described in the paper are highlighted.

- The typing rules for terms are in Figure 8.
- The typing rules for variables are in Figure 9.
- The subtyping rules are in Figure 10.
- The type splitting rules are in Figure 11.
- The typing rules for typing definitions are in Figure 12.
- The reduction rules (operational semantics) are in Figure 13.
- The rules for typing configurations and heap correspondence are in Figure 14.

$$\begin{array}{c}
\frac{\Gamma; \rho \vdash x : T \quad \Gamma \mathbf{vis} x}{\Gamma; \rho \vdash \mathbf{v}x : T} \text{(TT-Var)} \\
\frac{\Gamma; \rho \vdash t : T_1 \quad \Gamma; \rho \vdash T_1 <: T_2}{\Gamma; \rho \vdash t : T_2} \text{(TT-Sub)} \\
\frac{\Gamma; \rho \vdash t_1 : T_1 \quad z \notin \text{fv } T_2 \quad \Gamma, z : T_1; \rho \vdash t_2 : T_2}{\Gamma; \rho \vdash \text{let } z = t_1 \text{ in } t_2 : T_2} \text{(TT-Let)} \\
\frac{\Gamma; \rho \vdash x_1 : \{m(z : T_1, r : T_3) : T_2\} \quad \Gamma; \rho \vdash x_2 : T_1 \quad \Gamma \mathbf{vis} x_2}{\Gamma; \rho \vdash x_1 : [x_2/z]T_3 \quad \Gamma \mathbf{vis} x_1} \text{(TT-Call)} \\
\frac{\Gamma; \rho \vdash x_1 : T_1 \quad \Gamma \mathbf{vis} x_1 \quad \Gamma; \rho \vdash x : \{a : T_1..T_2\} \quad \Gamma \mathbf{vis} x}{\Gamma; \rho \vdash x : \{\mathbf{M}(r) : \perp.. \perp\}} \text{(TT-Write)} \\
\frac{\Gamma, s : T_1; \rho \vdash d : T_1 \quad z \notin \text{fv } T_2 \quad T_1 \mathbf{indep} s \quad \Gamma, z : \mu(s : T_1) \wedge \{\mathbf{M}(r) : \perp.. \perp\}; \rho \vdash t : T_2}{\Gamma; \rho \vdash \text{let } z = \nu(s : T_1)d \text{ in } t : T_2} \text{(TT-New)} \\
\frac{\Gamma; \rho \vdash x : \{a : T_1..T_2\} \quad \Gamma \mathbf{vis} x \quad \Gamma; \rho \vdash T_2 \mathbf{ro} T_3 \quad \Gamma; \rho \vdash T_2 \mathbf{mu}(r) T_4}{\Gamma; \rho \vdash x.a : T_3 \wedge \{\mathbf{M}(r) : \perp..(T_4 \vee x.\mathbf{M}(r))\}} \text{(TT-Read)}
\end{array}$$

Figure 8: roDOT term typing rules

$$\begin{array}{c}
\frac{! \notin \Gamma_2}{\Gamma_1, x : T, \Gamma_2 \mathbf{vis} x} \text{(Vis-Var)} \\
\frac{\Gamma; \rho \vdash x : \mu(s : T) \quad T \mathbf{indep} s}{\Gamma; \rho \vdash x : [x/s]T} \text{(VT-RecE)} \\
\frac{}{(\Gamma_1, x : T, \Gamma_2); \rho \vdash x : T} \text{(VT-Var)} \\
\frac{\Gamma; \rho \vdash x : T_1 \quad \Gamma; \rho \vdash x : T_2}{\Gamma; \rho \vdash x : T_1 \wedge T_2} \text{(VT-AndI)} \\
\frac{\Gamma; \rho \vdash x : [x/s]T \quad T \mathbf{indep} s}{\Gamma; \rho \vdash [x/s]T \mathbf{ro} [x/s]T} \text{(VT-RecI)} \\
\frac{\Gamma; \rho \vdash x : T}{\Gamma; \rho \vdash x : \{\mathbf{M}(r_0) : \perp.. \top\}} \text{(VT-M)} \\
\frac{\Gamma; \rho \vdash x : T_1 \quad \Gamma; \rho \vdash T_1 <: T_2}{\Gamma; \rho \vdash x : T_2} \text{(VT-Sub)}
\end{array}$$

Figure 9: roDOT variable typing rules

$$\begin{array}{c}
\overline{\Gamma; \rho \vdash T <: T} \text{ (ST-Refl)} \\
\overline{\Gamma; \rho \vdash T <: \top} \text{ (ST-Top)} \\
\overline{\Gamma; \rho \vdash \perp <: T} \text{ (ST-Bot)} \\
\frac{\rho \vdash T_1 \approx T_2}{\Gamma; \rho \vdash T_1 <: T_2} \text{ (ST-Eq)} \\
\overline{\Gamma; \rho \vdash T_1 <: T_1 \vee T_2} \text{ (ST-Or1)} \\
\overline{\Gamma; \rho \vdash T_2 <: T_1 \vee T_2} \text{ (ST-Or2)} \\
\overline{\Gamma; \rho \vdash T_1 \wedge T_2 <: T_1} \text{ (ST-And1)} \\
\overline{\Gamma; \rho \vdash T_1 \wedge T_2 <: T_2} \text{ (ST-And2)} \\
\overline{\Gamma; \rho \vdash \mathbf{N} <: \mu(s : T)} \text{ (ST-N-Rec)} \\
\frac{\Gamma; \rho \vdash x : \{B(r) : T_1..T_2\}}{\Gamma; \rho \vdash [x_2/r]T_1 <: x.B(x_2)} \text{ (ST-SelL)} \\
\frac{\Gamma; \rho \vdash x : \{B(r) : T_1..T_2\}}{\Gamma; \rho \vdash x.B(x_2) <: [x_2/r]T_2} \text{ (ST-SelU)} \\
\frac{\Gamma; \rho \vdash T_3 <: T_1 \quad \Gamma; \rho \vdash T_2 <: T_3}{\Gamma; \rho \vdash T_1 <: T_3} \text{ (ST-Trans)} \\
\frac{\Gamma; \rho \vdash T_1 <: T_2 \quad \Gamma; \rho \vdash T_1 <: T_3}{\Gamma; \rho \vdash T_1 <: T_2 \wedge T_3} \text{ (ST-And)} \\
\frac{\Gamma; \rho \vdash T_1 <: T_3 \quad \Gamma; \rho \vdash T_2 <: T_3}{\Gamma; \rho \vdash T_1 \vee T_2 <: T_3} \text{ (ST-Or)} \\
\frac{\Gamma; \rho \vdash T_3 <: T_1 \quad \Gamma; \rho \vdash T_2 <: T_4}{\Gamma; \rho \vdash \{B(r) : T_1..T_2\} <: \{B(r) : T_3..T_4\}} \text{ (ST-Typ)} \\
\frac{\Gamma; \rho \vdash T_3 <: T_1 \quad \Gamma; \rho \vdash T_2 <: T_4}{\Gamma; \rho \vdash \{a : T_1..T_2\} <: \{a : T_3..T_4\}} \text{ (ST-Fld)} \\
\overline{\Gamma; \rho \vdash \mathbf{N} <: \{A(r) : T_1..T_2\}} \text{ (ST-N-Typ)} \\
\overline{\Gamma; \rho \vdash \mathbf{N} <: \{m(z : T_1, r : T_3) : T_2\}} \text{ (ST-N-Met)} \\
\overline{\Gamma; \rho \vdash \mathbf{N} <: \{a : T_1..T_2\}} \text{ (ST-N-Fld)} \\
\overline{\Gamma; \rho \vdash \mathbf{N} \wedge \{M(r) : \perp.. \perp\} <: \perp} \text{ (ST-N-M)} \\
\overline{\Gamma; \rho \vdash \top <: \mathbf{N} \vee \{M(r) : T_1..T_2\}} \text{ (ST-NM)} \\
\frac{\Gamma; \rho \vdash T_3 <: T_1 \quad \Gamma, z : T_3, r : T_6; \rho \vdash T_2 <: T_4}{\Gamma, z : T_3; \rho \vdash T_6 <: T_5 \quad T_6 \text{ indep } z \Rightarrow T_5 \text{ indep } z} \text{ (ST-Met)} \\
\overline{\Gamma; \rho \vdash \{B(r) : T_1..T_2\} \wedge \{B(r) : T_3..T_4\} <: \{B(r) : T_1 \vee T_3..T_2 \wedge T_4\}} \text{ (ST-TypAnd)} \\
\overline{\Gamma; \rho \vdash T_1 \wedge (T_2 \vee T_3) <: (T_1 \wedge T_2) \vee (T_1 \wedge T_3)} \text{ (ST-Dist)}
\end{array}$$

Figure 10: roDOT subtyping rules

$$\begin{array}{c}
\frac{}{\Gamma; \rho \vdash \top \mathbf{ro} \top} \text{(TS-Top)} \\
\Gamma; \rho \vdash \top \mathbf{mu}(r) \top \\
\frac{}{\Gamma; \rho \vdash \perp \mathbf{ro} \mathbf{N}} \text{(TS-Bot)} \\
\Gamma; \rho \vdash \perp \mathbf{mu}(r) \perp \\
\frac{T = \{A(r) : T_1..T_2\}}{\Gamma; \rho \vdash T \mathbf{ro} T} \text{(TS-Typ)} \\
\Gamma; \rho \vdash T \mathbf{mu}(r_0) \top \\
\frac{T = \{m(z : T_1, r : T_3) : T_2\}}{\Gamma; \rho \vdash T \mathbf{ro} T} \text{(TS-Met)} \\
\Gamma; \rho \vdash T \mathbf{mu}(r_0) \top \\
\frac{T = \{a : T_1..T_2\}}{\Gamma; \rho \vdash T \mathbf{ro} T} \text{(TS-Fld)} \\
\Gamma; \rho \vdash T \mathbf{mu}(r) \top \\
\frac{T = \mu(s : T_1)}{\Gamma; \rho \vdash T \mathbf{ro} T} \text{(TS-Rec)} \\
\Gamma; \rho \vdash T \mathbf{mu}(r) \top \\
\frac{\Gamma; \rho \vdash \mathbf{N} <: T}{\Gamma; \rho \vdash T \mathbf{ro} T} \text{(TS-N)}
\end{array}
\qquad
\begin{array}{c}
\frac{\Gamma; \rho \vdash T_1 \mathbf{ro} T_2 \quad \Gamma; \rho \vdash T_3 \mathbf{ro} T_4}{\Gamma; \rho \vdash T_1 \wedge T_3 \mathbf{ro} T_2 \wedge T_4} \text{(TS-AndR)} \\
\frac{\Gamma; \rho \vdash T_1 \mathbf{mu}(r) T_2 \quad \Gamma; \rho \vdash T_3 \mathbf{mu}(r) T_4}{\Gamma; \rho \vdash T_1 \wedge T_3 \mathbf{mu}(r) T_2 \wedge T_4} \text{(TS-AndM)} \\
\frac{\Gamma; \rho \vdash T_1 \mathbf{ro} T_2 \quad \Gamma; \rho \vdash T_3 \mathbf{ro} T_4}{\Gamma; \rho \vdash T_1 \vee T_3 \mathbf{ro} T_2 \vee T_4} \text{(TS-OrR)} \\
\frac{\Gamma; \rho \vdash T_1 \mathbf{mu}(r) T_2 \quad \Gamma; \rho \vdash T_3 \mathbf{mu}(r) T_4}{\Gamma; \rho \vdash T_1 \vee T_3 \mathbf{mu}(r) T_2 \vee T_4} \text{(TS-OrM)} \\
\frac{\Gamma; \rho \vdash x : \{B(r) : T_1..T_2\} \quad \Gamma; \rho \vdash [x_2/r]T_2 \mathbf{ro} T_3}{\Gamma; \rho \vdash [x_2/r]T_2 \mathbf{mu}(r_0) T_4} \text{(TS-Sel)} \\
\Gamma; \rho \vdash x.B(x_2) \mathbf{ro} T_3 \\
\Gamma; \rho \vdash x.B(x_2) \mathbf{mu}(r_0) T_4 \\
\frac{\Gamma; \rho \vdash \{M(r) : T_1..T_2\} \mathbf{ro} \top}{\Gamma; \rho \vdash \{M(r) : T_1..T_2\} \mathbf{mu}(r) T_2} \text{(TS-M)}
\end{array}$$

Figure 11: roDOT type split rules

$$\begin{array}{c}
\frac{}{\Gamma, s : T_4; \rho \vdash \{A(r) = T\} : \{A(r) : T..T\}} \text{(DT-Typ)} \\
\frac{}{\Gamma, s : T_4; \rho \vdash \{A(r) = T\} : \{A(r) : \perp..T\}} \text{(DT-TypB)} \\
\frac{\Gamma, s : T_4; \rho \vdash x : T \quad \Gamma, s : T_4 \mathbf{vis} x}{\Gamma, s : T_4; \rho \vdash \{a = x\} : \{a : T..T\}} \text{(DT-Fld)} \\
\frac{\Gamma, s : T_4; \rho \vdash d_1 : T_1 \quad \Gamma, s : T_4; \rho \vdash d_2 : T_2 \quad d_1 \text{ and } d_2 \text{ have distinct member names}}{\Gamma, s : T_4; \rho \vdash d_1 \wedge d_2 : T_1 \wedge T_2} \text{(DT-And)} \\
\frac{z \notin \text{fv } T_1 \cup \text{fv } T_4, r \notin \text{fv } T_3 \cup \text{fv } T_1 \cup \text{fv } T_4 \quad \Gamma, s : T_4, !, z : T_1, r : T_4 \wedge [r/s]T_4 \wedge T_3; \rho \vdash t : T_2}{\Gamma, s : T_4; \rho \vdash \{m(z, r) = t\} : \{m(z : T_1, r : T_3) : T_2\}} \text{(DT-Met)}
\end{array}$$

Figure 12: roDOT definition typing rules

$$\begin{array}{c}
\frac{y_1 \rightarrow \dots_1 \{a = y_2\} \dots_2 \in \Sigma \quad w_1 \rightarrow y_1 \in \rho_1 \quad \rho_2 = \rho_1, w_2 \rightarrow y_2 \quad (w_2 \text{ fresh})}{\langle w_1.a; \sigma; \rho_1; \Sigma \rangle \mapsto \langle vw_2; \sigma; \rho_2; \Sigma \rangle} \text{(R-Read)} \\
\\
\frac{w_1 \rightarrow y_1 \in \rho \quad y_1 \rightarrow \dots_1 \{a = y_2\} \dots_2 \in \Sigma_1 \quad w_3 \rightarrow y_3 \in \rho \quad \Sigma_2 = \Sigma_1[y_1 \rightarrow \dots_1 \{a = y_3\} \dots_2]}{\langle w_1.a := w_3; \sigma; \rho; \Sigma_1 \rangle \mapsto \langle vw_3; \sigma; \rho; \Sigma_2 \rangle} \text{(R-Write)} \\
\\
\frac{w_1 \rightarrow y_1 \in \rho \quad y_1 \rightarrow \dots_1 \{m(z, r) = t\} \dots_2 \in \Sigma}{\langle w_1.m w_2; \sigma; \rho; \Sigma \rangle \mapsto \langle [w_1/r][w_2/z]t; \sigma; \rho; \Sigma \rangle} \text{(R-Call)} \\
\\
\frac{\rho_2 = \rho_1, w \rightarrow y \quad \Sigma_2 = \Sigma_1, y \rightarrow [y/s][\rho_1]d \quad (y, w \text{ fresh})}{\langle \text{let } z = \nu(s : T)d \text{ in } t; \sigma; \rho_1; \Sigma_1 \rangle \mapsto \langle [w/z]t; \sigma; \rho_2; \Sigma_2 \rangle} \text{(R-LetNew)} \\
\\
\overline{\langle \text{let } z = t_1 \text{ in } t_2; \sigma; \rho; \Sigma \rangle \mapsto \langle t_1; \text{let } z = \square \text{ in } t_2 :: \sigma; \rho; \Sigma \rangle} \text{(R-LetPush)} \\
\\
\overline{\langle vw; \text{let } z = \square \text{ in } t :: \sigma; \rho; \Sigma \rangle \mapsto \langle [w/z]t; \sigma; \rho; \Sigma \rangle} \text{(R-LetLoc)}
\end{array}$$

Figure 13: roDOT reduction rules

$$\begin{array}{c}
\frac{\Gamma; \rho \vdash T_1 <: T_2}{\Gamma; \rho \vdash \cdot : T_1, T_2} \text{(CT-EmptyS)} \quad \frac{\Gamma, z : T_1; \rho \vdash t : T_2 \quad z \notin \text{fv } T_2 \quad \Gamma; \rho \vdash \sigma : T_2, T_3}{\Gamma; \rho \vdash \text{let } z = \square \text{ in } t :: \sigma : T_1, T_3} \text{(CT-LetS)} \\
\\
\overline{\Gamma; \rho \vdash \cdot} \text{(CT-EmptyH)} \quad \overline{\Gamma \sim \cdot} \text{(CT-EmptyE)} \quad \frac{\Gamma_1; \rho \vdash \Gamma_2 \sim \Sigma}{\Gamma_1; \rho \vdash \Gamma_2, w : T \sim \Sigma} \text{(CT-RefH)} \\
\\
\frac{\Gamma_1, y/s : R; \rho \vdash d : [y/s]R \quad \Gamma_1; \rho \vdash \Gamma_2 \sim \Sigma \quad R \text{ indep } s}{\Gamma_1; \rho \vdash \Gamma_2, y : \mu(s : R) \wedge \{M(r) : \perp.. \perp\} \sim \Sigma, y \rightarrow d} \text{(CT-ObjH)} \\
\\
\frac{\Gamma = \Gamma_1, w : \mu(s : R) \wedge \{M(r) : \perp.. T\}, \Gamma_2 \quad \Gamma_1 \sim \rho \quad \Gamma_1 = \Gamma_3, y : \mu(s : R) \wedge \{M(r) : \perp.. \perp\}, \Gamma_4}{\Gamma \sim \rho, w \rightarrow y} \text{(CT-RefE)} \\
\\
\frac{\Gamma; \rho \vdash \Gamma \sim \Sigma \quad \text{all fields in } \Sigma \text{ are locations}}{\Gamma; \rho \sim \Sigma} \text{(CT-CorrH)} \quad \frac{\Gamma \sim \rho \quad \Gamma; \rho \sim \Sigma \quad \Gamma; \rho \vdash t : T_1 \quad \Gamma; \rho \vdash \sigma : T_1, T_2 \quad \text{no locations in } t \text{ and } \sigma}{\Gamma \vdash \langle t; \sigma; \rho; \Sigma \rangle : T_2} \text{(CT-Corr)} \\
\\
\frac{z \notin \text{fv } T_1 \cup \text{fv } T_4, r \notin \text{fv } T_3 \cup \text{fv } T_1 \cup \text{fv } T_4 \quad \Gamma, !, z : T_1, r : [y/s]T_4 \wedge [r/s]T_4 \wedge T_3; \rho \vdash t : T_2}{\Gamma, y/s : T_4; \rho \vdash \{m(z, r) = t\} : \{m(z : T_1, r : T_3) : T_2\}} \text{(HT-Met)}
\end{array}$$

Figure 14: roDOT configuration typing rules

A.2 Layered typing definitions

In this section, we show the typing rules of layered typing, which was introduced in Section 4.3.

- The first, atomic layer is in Appendix A.2.1 and Figure 15.
- The second, basic layer is in Appendix A.2.2 and Figure 16.
- The third, union layer is in Appendix A.2.3 and Figure 17.
- The fourth, logic layer is in Appendix A.2.4 and Figure 18.
- The fifth, main layer is in Appendix A.2.5 and Figure 19.
- Additional definitions required for layered typing are in Appendix A.2.6.

A.2.1 Atomic layer typing

The atomic layer is between precise typing and the basic layer. It handles declaration subtyping and equivalence. The subtyping here is limited to simple subtyping between atomic member declarations. Rules of atomic layer typing are shown in Figure 15.

$$\frac{\Gamma \vdash_1 x : T_1 \quad \rho \vdash T_1 \approx T_2}{\Gamma; \rho \vdash_a x : T_2} \text{(LTA-Prec)} \quad \frac{\Gamma; \rho \vdash_a x : T_1}{\Gamma; \rho \vdash_a T_1 <: T_2} \text{(LTA-Sub)} \quad \frac{}{\Gamma; \rho \vdash_a x : T_2} \text{(LTA-Sub)}$$

Figure 15: Rules of atomic layer typing

A.2.2 Basic layer typing

The basic layer is between the atomic layer and the union layer. It handles the top-like types $N \vee \{M(r_0) : T_1..T_2\}$ and $\{M(r_0) : T_1..T_2\} \vee N$. Rules of basic layer typing are shown in Figure 16.

$$\frac{\Gamma; \rho \vdash_a x : T}{\Gamma; \rho \vdash_b x : T} \text{(LTB-Atom)} \quad \frac{\vdash_c N \triangleleft T_1 \quad \vdash_c M \triangleleft T_2}{\Gamma; \rho \vdash_b x : T_1 \vee T_2} \text{(LTB-NM)} \quad \frac{\vdash_c M \triangleleft T_1 \quad \vdash_c N \triangleleft T_2}{\Gamma; \rho \vdash_b x : T_1 \vee T_2} \text{(LTB-MN)}$$

Figure 16: Rules of basic layer typing

A.2.3 Union layer typing

The union layer is between the basic layer and logic layer. It handles union types and the top type. Rules of union layer typing are shown in Figure 17.

$$\frac{\Gamma; \rho \vdash_b x : T}{\Gamma; \rho \vdash_c x : T} \text{(LTC-Basic)} \quad \frac{}{\Gamma; \rho \vdash_c x : \top} \text{(LTC-Top)}$$

$$\frac{\Gamma; \rho \vdash_c x : T_1}{\Gamma; \rho \vdash_c x : T_1 \vee T_2} \text{(LTC-Or1)} \quad \frac{\Gamma; \rho \vdash_c x : T_2}{\Gamma; \rho \vdash_c x : T_1 \vee T_2} \text{(LTC-Or2)}$$

Figure 17: Rules of union layer typing

A.2.4 Logic layer typing

The logic layer is between the union layer and main layer. It handles intersection types, by treating the unions derived by the union layer as CNF, and pushing intersections down. Rules of logic layer typing are shown in Figure 18.

The (LTL-And) rule takes derivations of two types which may have some parts in common, but differ in one place. For example, we can write the two types $\{a_1 : T_1..T_1\} \vee \{a_2 : T_2..T_2\}$ and $\{a_1 : T_1..T_1\} \vee \{M(r_0) : \perp..\perp\}$ as $C^\vee[\{a_2 : T_2..T_2\}]$ and $C^\vee[\{M(r_0) : \perp..\perp\}]$, where C^\vee is the common part – a syntactical context which combines the argument into a union with other types. The rule takes two such types, preserves their common part, and combines the differing parts using an intersection type – pushing the intersection down from the top to its target place.

$$\frac{\Gamma;\rho \vdash_c x : T \quad x \in \text{dom } \Gamma}{\Gamma;\rho \vdash_1 x : T} \text{(LTL-Or)} \quad \frac{\Gamma;\rho \vdash_1 x : C^\vee[T_1] \quad \Gamma;\rho \vdash_1 x : C^\vee[T_2]}{\Gamma;\rho \vdash_1 x : C^\vee[T_1 \wedge T_2]} \text{(LTL-And)}$$

Figure 18: Rules of logic typing

A.2.5 Main layer typing

The main layer is above the logic layer. It handles selections and recursive types. Rules of main layer typing are shown in Figure 19.

Here, the syntactic context C^{\wedge} can consist of a mixture of unions and intersections. The (LTM-Sel) and (LTM-Rec) have additional conditions which correspond to the conditions in the corresponding rules in tight typing and subtyping. The (LTM-N) rule handles the fact that $\Gamma;\rho \vdash N <: \mu(s : T)$.

$$\frac{\Gamma;\rho \vdash_1 x : T}{\Gamma;\rho \vdash_m x : T} \text{(LTM-Logic)} \quad \frac{\Gamma;\rho \vdash_m x : C^{\wedge}[[x/s]T] \quad T \text{ indep } s}{\Gamma;\rho \vdash_m x : C^{\wedge}[\mu(s : T)]} \text{(LTM-Rec)}$$

$$\frac{\Gamma;\rho \vdash_m x : C^{\wedge}[[v_3/r]T_1] \quad \Gamma \vdash_! v_2 : \{B(r) : T_1..T_2\}}{\Gamma;\rho \vdash_m x : C^{\wedge}[v_2.B(v_3)]} \text{(LTM-Sel)} \quad \frac{\Gamma;\rho \vdash_m x : C^{\wedge}[N]}{\Gamma;\rho \vdash_m x : C^{\wedge}[\mu(s : T)]} \text{(LTM-N)}$$

Figure 19: Rules of main typing

A.2.6 Additional definitions

The layered typing uses additional typing judgements:

- **Atomic subtyping** (Figure 20) is tight subtyping restricted to simple declarations.
- The definition of **N- and M-supertypes** (used in the basic typing layer) is also layered. The basic layer (Figure 21) recognizes declarations and recursive types, and the union layer (Figure 22) adds top and union types.

$$\begin{array}{c}
\frac{\Gamma; \rho \vdash_{\#} T_3 <: T_1 \quad \Gamma; \rho \vdash_{\#} T_2 <: T_4}{\Gamma; \rho \vdash_a \{B(r) : T_1..T_2\} <: \{B(r) : T_3..T_4\}} \text{(LSA-Typ)} \\
\\
\frac{\Gamma; \rho \vdash_{\#} T_3 <: T_1 \quad \Gamma, z : T_3; \rho \vdash T_6 <: T_5 \quad \Gamma, z : T_3, r : T_6; \rho \vdash T_2 <: T_4}{\Gamma; \rho \vdash_a \{m(z : T_1, r : T_5) : T_2\} <: \{m(z : T_3, r : T_6) : T_4\}} \text{(LSA-Met)} \\
\\
\frac{\Gamma; \rho \vdash_{\#} T_3 <: T_1 \quad \Gamma; \rho \vdash_{\#} T_2 <: T_4}{\Gamma; \rho \vdash_a \{a : T_1..T_2\} <: \{a : T_3..T_4\}} \text{(LSA-Fld)}
\end{array}$$

Figure 20: Rules of atomic subtyping

$$\begin{array}{c}
\frac{}{\vdash_b M \triangleleft \{M(r_0) : T_1..T_2\}} \text{(LNB-M)} \\
\\
\frac{}{\vdash_b N \triangleleft N} \text{(LNB-N)} \quad \frac{}{\vdash_b N \triangleleft \{m(z : T_1, r : T_5) : T_2\}} \text{(LNB-Met)} \\
\\
\frac{}{\vdash_b N \triangleleft \{a : T_1..T_2\}} \text{(LNB-Fld)} \quad \frac{}{\vdash_b N \triangleleft \{A(r) : T_1..T_2\}} \text{(LNB-Typ)}
\end{array}$$

Figure 21: N and M supertypes

$$\begin{array}{c}
\frac{\vdash_b N \triangleleft T}{\vdash_c N \triangleleft T} \text{(LNC-Basic)} \quad \frac{}{\vdash_c N \triangleleft \top} \text{(LNC-Top)} \\
\\
\frac{\vdash_c N \triangleleft T_1}{\vdash_c N \triangleleft T_1 \vee T_2} \text{(LNC-Or1)} \quad \frac{\vdash_c N \triangleleft T_2}{\vdash_c N \triangleleft T_1 \vee T_2} \text{(LNC-Or2)} \\
\\
\frac{\vdash_b \mathcal{M} \triangleleft T}{\vdash_c \mathcal{M} \triangleleft T} \text{(LMC-Basic)} \quad \frac{}{\vdash_c \mathcal{M} \triangleleft \top} \text{(LMC-Top)} \\
\\
\frac{\vdash_c \mathcal{M} \triangleleft T_1}{\vdash_c \mathcal{M} \triangleleft T_1 \vee T_2} \text{(LMC-Or1)} \quad \frac{\vdash_c \mathcal{M} \triangleleft T_2}{\vdash_c \mathcal{M} \triangleleft T_1 \vee T_2} \text{(LMC-Or2)}
\end{array}$$

Figure 22: N and M supertypes union closure

A.3 Layered typing properties

The layered typing introduced in Section 4.3 is used in the proof of soundness of the updated roDOT calculus. The proof relies on several properties of layered typing. We stated the most important in Section 4.3.

In this section, we show additional properties of layered typing, which are used in proofs of the main properties and show the general way of working with layered typing.

Lemma 28 (Type declaration inversion). *If $\Gamma; \rho \vdash_m x : \{B(r) : T_1..T_2\}$, then there exist T_3, T_4 such that $\Gamma; \rho \vdash_l x : \{B(r) : T_3..T_4\}$, $\Gamma; \rho \vdash_{\#} T_1 <: T_3$ and $\Gamma; \rho \vdash_{\#} T_4 <: T_2$.*

Proof. Declaration types are only affected by the atomic layer, so we can simply invert all rules on the layers above, and the rules on the atomic layer give us the desired result. \square

Lemma 29 (Field declaration inversion). *If $\Gamma; \rho \vdash_m x : \{a : T_1..T_2\}$, then there exist T_3, T_4 such that $\Gamma; \rho \vdash_1 x : \{a : T_3..T_4\}$, $\Gamma; \rho \vdash_{\#} T_1 <: T_3$ and $\Gamma; \rho \vdash_{\#} T_4 <: T_2$.*

Proof. Declaration types are only affected by the atomic layer, so we can simply invert all rules on the layers above, and the rules on the atomic layer give us the desired result. \square

Lemma 30 (Method declaration inversion). *If $\Gamma; \rho \vdash_m x : \{m(z : T_1, r : T_5) : T_2\}$, then there exist T_3, T_4, T_6 , such that $\Gamma; \rho \vdash_1 x : \{m(z : T_3, r : T_6) : T_3\}T_4$, $\Gamma; \rho \vdash_{\#} T_3 <: T_1$, $\Gamma, z : T_3; \rho \vdash_{\#} T_6 <: T_5$ and $\Gamma, z : T_3, r : T_6; \rho \vdash_{\#} T_2 <: T_4$.*

Proof. Declaration types are only affected by the atomic layer, so we can simply invert all rules on the layers above, and the rules on the atomic layer give us the desired result. \square

Lemma 31 (Tight to layered typing). *If $\Gamma; \rho \vdash_{\#} x : T$, then $\Gamma; \rho \vdash_m x : T$.*

The implication from tight typing to invertible typing requires us to show that invertible typing is closed under the rules of tight typing and subtyping. Because the main layer deals with parts of a type under some syntactic context C^{\wedge} , it is useful to formulate the property in this way:

Lemma 32 (Layered typing closed under subtyping). *If $\Gamma; \rho \vdash_m x : C^{\wedge}[T_1]$, and $\Gamma; \rho \vdash_{\#} T_1 <: T_2$, then $\Gamma; \rho \vdash_m x : C^{\wedge}[T_2]$.*

Similarly to how it was done for invertible typing in the original roDOT, the proof involves inverting the rules of layered typing. Here though, the inversion has to go through multiple layers of typing, until it reaches the layer where the type system feature involved in the subtyping rule is handled.

Below, we show examples of properties of typing on the main layer that are analogous to subtyping rules.

The subtyping rules for intersection types (ST-And1, ST-And2 and ST-And) must be handled at the logic layer by the rule LTL-And. On the main layer, we show that the features of the main layer do not interfere with intersection types, so we can derive and invert them in the same manner:

Lemma 33 (Intersection at the main layer). *If $\Gamma; \rho \vdash_m x : C^{\wedge}[T_1]$, and $\Gamma; \rho \vdash_m x : C^{\wedge}[T_2]$, then $\Gamma; \rho \vdash_m x : C^{\wedge}[T_1 \wedge T_2]$.*

Lemma 34 (Inversion of intersection at the main layer). *If $\Gamma; \rho \vdash_m x : C^{\wedge}[T_1 \wedge T_2]$, then $\Gamma; \rho \vdash_m x : C^{\wedge}[T_1]$ and $\Gamma; \rho \vdash_m x : C^{\wedge}[T_2]$.*

The rule ST-Or can be handled, thanks to induction on the tight subtyping and to the syntactic context C^{\wedge} , by a simpler property:

Lemma 35 (Inversion of union types at the main layer). *If $\Gamma; \rho \vdash_m x : C^{\wedge}[T \vee T]$, then $\Gamma; \rho \vdash_m x : C^{\wedge}[T]$.*

On the main layer, the N type can be replaced by any read-only type:

Lemma 36 (Replacement of N at the main layer). *If $\Gamma; \rho \vdash_m x : C^{\wedge}[N]$, and $\Gamma \vdash_m N \triangleleft T$, then $\Gamma; \rho \vdash_m x : C^{\wedge}[T]$.*

Finally, the other implication (invertible typing implies tight typing) is expressed by Lemma 37.

Lemma 37 (Layered to tight typing). *If $\Gamma; \rho \vdash_m x : T$, then $\Gamma; \rho \vdash_{\#} x : T$.*

This was very easy in the original DOT, because there, each rule of invertible typing can be mirrored by application of just a few rules of tight typing and subtyping.

In layered typing though, the same approach only works up to the logic layer. On the main layer, there is a problem that the rule LTM-Rec allows closing a recursive type anywhere under a syntactic context C^{\wedge} .

However, general and tight typing do not have subtyping rules for recursive types. The rule that closes recursive types in tight typing, which has the same form as VT-Recl, can only close the whole type, and not just a part of it.

That is because this rule is part of a typing judgment, not subtyping. Therefore, it cannot be combined with other subtyping rules (subtyping of unions, intersections, transitivity).

To solve this problem, we can make use of the requirement that the inside of a recursive type must be read-only. Then, we can look at what importance this part of the type has in the typing derivation. We can actually show that any read-only part of a type is either derived from the type of the variable, or it can be replaced by N:

Lemma 38 (Replacement of N supertypes at the main layer). *If $\Gamma; \rho \vdash_m x : C^{\wedge\vee}[T]$, and $\Gamma \vdash_m N \triangleleft T$, then either $\Gamma; \rho \vdash_m x : T$ or $\Gamma; \rho \vdash_m x : C^{\wedge\vee}[N]$.*

This lemma allows us, for a type $C^{\wedge\vee}[T]$, to either bring the type T to the top level and use VT-RecI to derive the recursive type, or to replace T by N and use subtyping with $\Gamma; \rho \vdash_{\#} N <: \mu(s : T)$.

This approach has one issue, however: we cannot prove Lemma 37 by induction on the main-layer typing derivation, because when we use Lemma 38, it gives a derivation of just T , which is not a sub-derivation of the derivation of $C^{\wedge\vee}[T]$, so we cannot use the induction hypothesis on it.

Instead of doing the proof directly by induction on the derivation, we do the proof by first using induction on the number of main-layer rules applied in the derivation, and second induction on the number of union and intersection constructors within the type.

We handle the main-layer rule applications one by one, so this metric is decreasing, except the case when Lemma 38 needs to be used. It is only needed if $C^{\wedge\vee}$ has at least one union/intersection constructor. In that case, we can find a derivation of T that has the same number of main-layer rule application as $C^{\wedge\vee}[T]$ and the number of union/intersection constructors is lowered.

A.4 Alternative SEF conditions

We stated the SEF condition (Definition 15) informally in Section 5.1 and formally in Section 3.1, also mentioning the possibility to stating side-effect-freedom conditions in other variants with weaker or stronger requirements. Here, we formally state several other variants of how a SEF condition could be defined in roDOT.

The conditions Definition 39, Definition 40 and Definition 41 have a stronger requirement of not creating any new objects. The conditions Definition 39 and Definition 42 have a stronger requirement of not writing to fields of existing objects, even if the write does not actually change the object. The conditions Definition 41 and Definition 43 have a weaker requirement of object being modifiable if they are restored back to their original state.

The Sef-N0 property (Definition 39) states that a method does execute any heap-modifying steps.

Definition 39 (Sef-N0). *A method execution $\langle w_1.m w_2; \sigma; \rho_1; \Sigma_1 \rangle \mapsto^k \langle v w_3; \sigma; \rho_2; \Sigma_2 \rangle$ is Sef-N0 when for every $j \leq k$ and $\langle w_1.m w_2; \sigma; \rho_1; \Sigma_1 \rangle \mapsto^j \langle t_3; \sigma_3; \rho_3; \Sigma_3 \rangle$ the term t_3 is not an object creation or write term.*

The Sef-I0 property (Definition 40) states that the heap stays the same during the execution of the method.

Definition 40 (Sef-I0). *A method execution $\langle w_1.m w_2; \sigma; \rho_1; \Sigma_1 \rangle \mapsto^k \langle v w_3; \sigma; \rho_2; \Sigma_2 \rangle$ is Sef-I0 when for every $j \leq k$ and $\langle w_1.m w_2; \sigma; \rho_1; \Sigma_1 \rangle \mapsto^j \langle t_3; \sigma_3; \rho_3; \Sigma_3 \rangle$ we have $\Sigma_3 = \Sigma_1$.*

The Sef-O0 property, defined in Definition 41, allows the method to modify the heap as long as it restores the heap to its original state before it returns.

Definition 41 (Sef-O0). *A method execution $\langle w_1.m w_2; \sigma; \rho_1; \Sigma_1 \rangle \mapsto^k \langle v w_3; \sigma; \rho_2; \Sigma_2 \rangle$ is Sef-O0 when $\Sigma_2 = \Sigma_1$.*

The Sef-N property (Definition 42) states that a method does execute any steps that write to objects on the heap. It is similar to Sef-N0, but allows creation of new objects.

Definition 42 (Sef-N). *A method execution $\langle w_1.m w_2; \sigma; \rho_1; \Sigma_1 \rangle \mapsto^k \langle v w_3; \sigma; \rho_2; \Sigma_2 \rangle$ is Sef-N when for every $j \leq k$ and $\langle w_1.m w_2; \sigma; \rho_1; \Sigma_1 \rangle \mapsto^j \langle t_3; \sigma_3; \rho_3; \Sigma_3 \rangle$ the term t_3 is not a write term.*

The Sef-O property (Definition 43) states that at the end of the method, the new heap has the original heap as a prefix. It is similar to Sef-O0, but allows creation of new objects.

Definition 43 (Sef-O). *A method execution $\langle w_1.m w_2; \sigma; \rho_1; \Sigma_1 \rangle \mapsto^k \langle v w_3; \sigma; \rho_2; \Sigma_2 \rangle$ is Sef-O when Σ_1 is a prefix of Σ_2 .*

From the three Definitions 15, 42 and 43, Definition 42 is unnecessarily restrictive, because it forbids modifications of existing objects. The difference between Definition 15 and Definition 43 is that Definition 43 allows modifying an object, as long as it is changed back before returning from the method. Making use of this possibility would require tracking the values assigned to the fields, which is out of scope of roDOT's type system. Also, although we only consider single-threaded execution, Definition 43 would not adapt well to possible extensions to multi-threaded environments, because any changes to the object within the method would be observable from other threads. For these reasons, we use the Sef-I property (Definition 15) as the definition of a SEF method and the basis for our SEF guarantee.

A.5 Similarity

In Section 5.4, we introduced the similarity relation. In this section, we provide the formal definition of similarity and prove a few important properties.

Similarity formalizes structural equivalence of syntactic elements such as terms, objects or whole configurations, which differ only in names of variables. The correspondence of variable names on two sides of the equivalence is specified by a renaming list.

Definition 44 (Renaming). *A list of pairs of variables $X := (x_{1,\ell}, x_{2,\ell})_{\ell=1,\dots,L}$ is a renaming, which renames each variable $x_{1,\ell}$ on the left to $x_{2,\ell}$ on the right.*

For simple syntactic elements, similarity is defined by the ability to get both sides of the equivalence by substitution starting from a common element.

Definition 45 (Similar variables, objects, terms and stacks). *Two terms t_1 and t_2 (or variables, objects and stacks) are similar by renaming $X := (x_{1,\ell}, x_{2,\ell})_{\ell=1,\dots,L}$, written $t_1 \overset{X}{\approx} t_2$, if there exists a term t_3 with list of variables $(x_{3,\ell})_{\ell=1,\dots,L}$, such that $\text{fv } t_3 \subseteq \{x_{3,\ell}\}_{\ell=1,\dots,L}$ and $[x_{1,\ell}/x_{3,\ell}]_{\ell} t_3 = t_1$ and $[x_{2,\ell}/x_{3,\ell}]_{\ell} t_3 = t_2$.*

The definition of similarity for configurations is complicated by variables in the domains Σ and environment ρ . Because the domains cannot contain duplicates, but we want to have duplicates in the renaming used for the term, in order to be able to make y_i correspond to both y_i and w_i in Lemma 20. To allow this, are two lists of variables – one used for domains of Σ and ρ and the other for their codomain, for the term and the stack.

Definition 46 (Similar configurations). *Two configurations $\langle t_1; \sigma_1; \rho_1; \Sigma_1 \rangle$ and $\langle t_2; \sigma_2; \rho_2; \Sigma_2 \rangle$ are similar up to renaming $X := (x_{1,\ell}, x_{2,\ell})_{\ell=1,\dots,L}$ in the input and $Y := (y_{1,\ell}, y_{2,\ell})_{\ell=1,\dots,L}$ in the output exactly when*

- Terms and stacks are similar by Y : $t_1 \overset{Y}{\approx} t_2$ and $\sigma_1 \overset{Y}{\approx} \sigma_2$.
- Heaps and environments are similar with respect to renaming inputs. For each $(x_{1,k}, x_{2,k})$ from X , either $\Sigma_1(x_{1,\ell}) \overset{Y}{\approx} \Sigma_2(x_{2,\ell})$ or $\rho_1(x_{1,\ell}) \overset{Y}{\approx} \rho_2(x_{2,\ell})$.
- Heap correspondence in similarity. If z_i are heap locations in Σ_i such that $(z_1, z_2) \in Y$ then $(z_1, z_2) \in X$.

Notation for similarity of configurations is $-\overset{Y}{\underset{X}{\approx}}-$, with X, Y omitted for the one that induces equality.

Lemma 47 (Similarity is preserved by reduction). *Let c_1, c_2 be well-typed configurations, $c_1 \overset{Y}{\underset{X}{\approx}} c_2$ and $c_i \mapsto^n c'_i$ for some n . Then there exists a renaming $L := (y_{1,k}, y_{2,k})_{k=1,\dots,K}$, such that each $y_{i,k}$ is fresh in c_i , and $c'_1 \overset{Y \cup L}{\underset{X \cup L}{\approx}} c'_2$.*

This is illustrated by the following diagram:

$$\begin{array}{ccc} c_1 & \overset{Y}{\underset{X}{\approx}} & c_2 \\ \downarrow^n & & \downarrow^n \\ c'_1 & \overset{Y \cup L}{\underset{X \cup L}{\approx}} & c'_2 \end{array}$$

Proof. It suffices to consider the case of a single-step reduction (i.e. when $n = 1$) because one can just use subject reduction and induction for the multi-step case. In the single-step case, we solve each case of $c_i \mapsto c'_i$ (defined in Figure 13 in the appendix), taking advantage of the fact that $t_1 \overset{Y}{\approx} t_2$, where t_i is the term of c_i , to ensure that the same reduction rule is used to reduce both $c_i \mapsto c'_i$.

The freshness of the variables $y_{i,k}$ follows from the fact that new in all reduction rules, the domains of Σ and ρ are either unchanged or extended with fresh variables. The cases of field read and writes are justified by the second and third conditions of Definition 46 along with the fact that configuration typing means heap correspondence. \square

Lemma 48 (Similarity ensures reduction). *If $c_1 \overset{Y}{\underset{X}{\approx}} c_2$ for c_i well-typed and $c_1 \mapsto^n c'_1$, then there exists c'_2 for which $c_2 \mapsto^n c'_2$. This is illustrated by the following diagram:*

$$\begin{array}{ccc} c_1 & \overset{Y}{\underset{X}{\approx}} & c_2 \\ \downarrow n & & \downarrow n \\ c'_1 & & c'_2 \end{array}$$

Proof. If c_2 is an answer configuration, then similarity means c_1 must be an answer configuration too, which is impossible unless $n = 0$. In the inductive case, use the subject reduction property and the similarity part of Lemma 47. \square

Proof. By definitions of c_2 , c'_1 and Definition 46. \square

A.6 Proofs of SEF guarantee lemmata

In this section, we give proofs to selected lemmata from Section 5.

Lemma (Lemma 20, Similarity for eliminated references). *Let c'_1 satisfy the conditions of Lemma 17 and c''_1 the conditions of Lemma 19. Then $c'_1 \overset{(w_1, y_1), (w_2, y_2)}{\approx} c''_1$.*

Lemma (Lemma 19, Read-only weakening). *Let c_1 and Γ satisfy the conditions of Theorem 16, and $\rho_1[w_i] = y_i$. Then there is a context Γ_2 binding w_i to read-only types such that the configuration $c'_1 := \langle w_1.m w_2; \cdot; \rho_1; [y_i/w_i]_i \Sigma_1 \rangle$, is well-typed in Γ_2 (formally $\Gamma_2 \vdash c'_1 : \top$).*

Proof. Take Γ' and ρ' as from Lemma 18. Let T_i be the read-only version of $\Gamma(y_i)$ so that T_i differs from $\Gamma(w_i)$ only by mutability. The goal is to take $\Gamma_2 := \Gamma'(w_i : T_i)_i$.

The hardest part of showing $\Gamma_2 \vdash c'_1 : \top$ is the part of typing $w_1.m w_2$. Note that the only difference between T_i and $\Gamma(w_i)$ is that T_i is by construction read-only while $\Gamma(w_i)$ is possibly mutable. So the idea is to use the same derivation of $\Gamma' \vdash y_1.m y_2 : \top$, in the process taking advantage of the fact that $\Gamma \vdash w_1 : \{m(z : N)(r : N) : \top\}$ means nowhere in the derivation did one need to use the fact that $\Gamma(w_i)$ is potentially mutable.

In the derivation of $\Gamma' \vdash y_1.m y_2 : \top$, one sees there are sub-derivations giving read-only types T_i to y_i . Lemma 9 says that for any context Γ and environment ρ that are in correspondence, if a location y_i can be derived to have a read-only type T_i then one may give the same type T_i to any references w_i to the location y_i . The result of Lemma 19 follows. \square

Lemma (Lemma 22, Reduction only changes fields). *If $\langle -; -; -; \Sigma \rangle \mapsto^n \langle -; -; -; \Sigma' \rangle$ and y is a location in Σ , then $\Sigma(y) \overset{\text{fld}}{\approx} \Sigma'(y)$.*

Proof. The only reduction rule that changes existing objects is the field-write reduction rule, which changes just the fields. \square

Lemma (Lemma 17, SEF guarantee without stack). *For c_1 satisfying the conditions of Theorem 16, let $c'_1 := \langle w_1.m w_2; \cdot; \rho_1; \Sigma_1 \rangle$. If $c'_1 \mapsto^n c'_2$ for some n and c'_2 , then the heap of c'_2 contains all objects of c'_1 without modification.*

Proof. We use Lemmas 20, 47 and 48 to get

$$\begin{array}{ccc} c'_1 & \overset{(w_i, y_i)_i}{\approx} & c'_1 \\ \downarrow n & & \downarrow n \\ c'_2 & \overset{(w_i, y_i)_i \cup L}{\underset{L}{\approx}} & c'_2 \end{array}$$

where L is fresh in both c'_1 and c'_2 . Let Σ_2 be the heap of c'_2 and Σ''_2 be the heap of c'_2 . Let y be a heap location of Σ_1 , so that Lemma 22 applied to the left edge means that $\Sigma_1(y) \overset{\text{fld}}{\approx} \Sigma_2(y)$. The immutability guarantee applies to the right edge because of Lemma 19 and so we know $\Sigma''_2(y) = \Sigma_1(y)$. The freshness of L and the bottom edge then implies $\Sigma_2(y) \overset{(w_i, y_i)_i}{\approx} \Sigma''_2(y) \overset{(w_i, y_i)_i}{\approx} \Sigma_1(y)$. But $\Sigma_1(y)$ can only differ from $\Sigma_2(y)$ in field values and w_i are references, so they never occur as values of fields, therefore $\Sigma_2(y) = \Sigma_1(y)$. \square

And with this, Theorem 16 follows straightforwardly.

Theorem (Theorem 16, SEF guarantee). *Let the configuration $c_1 := \langle w_1.m \ w_2; \sigma_1; \rho_1; \Sigma_1 \rangle$ be well-typed in a context Γ . Further assume that $\Gamma \vdash w_1 : \{m(z : \mathbb{N})(r : \mathbb{N}) : \top\}$. Then for any k steps of execution:*

1. Either the method call has finished executing. *There is $j < k$ for which $c_1 \mapsto^j \langle \nu w_3; \sigma_1; -; - \rangle$.*
2. Or, the method call has not finished executing and in this period existing objects in the heap are unchanged. *For each $c_1 \mapsto^k c_2$, all heap locations in c_1 also exist in c_2 and moreover they are unchanged in c_2 .*

Proof. By classical reasoning, assume the condition 1 is false so that the goal is to prove the condition 2. That is, assume that there is no $j < k$ such that the top-most frame of c_1 is popped after execution by j steps: $c_1 \mapsto^j \langle \nu w_3; -; -; \Sigma_2 \rangle$. Then, the sequence of reductions $c_1 \mapsto \dots \mapsto c_4$ by Lemma 48 corresponds to a sequence of reductions $c_2 \mapsto \dots \mapsto c_3$ because even though c_2 has no awaiting frames, there are no frame pops in this execution sequence by the current assumption. By Lemma 17, the condition 2 follows. \square

A.7 Transformation framework

In Section 6.1, we described the general transformation framework in short. The framework consists of definitions and lemmas that enable reasoning about safe transformations of programs in roDOT. In this section, we present a more detailed discussion and basic definitions of the framework.

A.7.1 Framework design

The framework design is influenced by the following considerations:

- Programs in DOT are *terms* in λ -calculus style and A-normal form. They are not formed by sequences of statements, but rather by nesting let-in terms.
- The program can contain object literals and method definitions, so a transformation such as swapping two calls can also be located inside a body of a method.
- The transformation can be conditioned on local typing information, such as the method calls being recognized as SEF. In roDOT, this necessitates looking at the typing context Γ that is used for typing that part of the program.
- The execution of a program is defined using small-step operational semantics, starting from an initial configuration and applying reduction steps until reaching an answer.

For discussion of what a safe transformation is, we will consider two programs, an original program and a transformed program, which only differ in ways allowed by the transformation. Such a transformation is safe if executions of those two programs reach the “same” answer. We cannot require the answers of the two programs to be identical, because of the following points.

- During execution of an roDOT program, objects are created on the heap and these objects contain methods, containing code where the transformation may be located.

Therefore, if we look at execution of the two programs in parallel, the transformation affects not only the program being executed, but also the bodies of these methods in objects on the heap.

In DOT, the code (terms) and values (objects) are mixed with each other. In order to allow reasoning about the transformation, we must consider that during execution, the transformation may be located at multiple places in the focus of execution, stack and heap.

This also applies to the answer of the program. To define a safe transformation, it would be too restrictive to require the two programs to produce the exact same output value since the output value may be an object that may contain the transformed code.

- The locations of objects on the heap are chosen non-deterministically, therefore the framework must deal with the fact that in the answer and in intermediate states of two executions, location names can differ.

- Otherwise the execution is deterministic. Programs in DOT and roDOT do not read any input and do not make non-deterministic choices. This makes the situation easier, as we can assume that all possible executions of one program will reach answers that are similar to each other. Therefore, the differences in executions of the original and the transformed program are only caused by the transformation of the initial program.

The transformation framework consists of the following parts:

- Definition of transformations of syntactic elements such as terms and configurations
- Definition of lifting transformations from local term transformations to transformations of whole programs and run-time machine configurations
- Definition of a *similarity transformation*, which uses similarity defined in Appendix A.5 to deal with differences in variable names.
- Definition of important properties of local transformations, which are relevant to showing that those transformations are safe.
- Lemmas about lifting such properties from local transformations to whole programs.
- Definition of a safe transformation – that does not change the result of execution of a program.
- Theorem 25, about safety of transformations which have the required properties and are lifted to whole programs.

A.7.2 Transformations

A *transformation* in general is defined in Definition 23.

- A *term transformation* is a transformation $\langle \Gamma_1; \rho_1 \vdash t_1 : T_1 \rangle \rightarrow_\tau \langle \Gamma_2; \rho_2 \vdash t_2 : T_2 \rangle$, where $t_{1,2}$ are terms, $T_{1,2}$ are types, $\Gamma_{1,2}$ are typing contexts (mapping variables to types) and $\rho_{1,2}$ are environments mapping object references to heap locations.
- A *stack transformation* is a transformation $\langle \Gamma_1; \rho_1 \vdash \sigma_1 : T_1, T_3 \rangle \rightarrow_\tau \langle \Gamma_2; \rho_2 \vdash \sigma_2 : T_2, T_4 \rangle$, where $\sigma_{1,2}$ are stacks, $T_{1,2}$ are types of the holes in the top frame of $\sigma_{1,2}$ respectively, $T_{3,4}$ are types of the bottom frame of $\sigma_{1,2}$, $\Gamma_{1,2}$ are typing contexts (mapping variables to types) and $\rho_{1,2}$ are environments mapping object references to heap locations.
- A *heap transformation* is a transformation $\langle \rho_1 \vdash \sigma_1 : \Gamma_1 \rangle \rightarrow_\tau \langle \rho_2 \vdash \sigma_2 : \Gamma_2 \rangle$, where $\Sigma_{1,2}$ are heaps, $\Gamma_{1,2}$ are typing contexts (mapping variables to types) and $\rho_{1,2}$ are environments mapping object references to heap locations. In this definition, the typing context take the position of a type, because the types of variables in $\Gamma_{1,2}$ correspond to types in $\Sigma_{1,2}$.
- A *configuration transformation* is a transformation $\langle \Gamma_1 \vdash c_1 : T_1 \rangle \rightarrow_\tau \langle \Gamma_2 \vdash c_2 : T_2 \rangle$, where $c_{1,2}$ are configurations, $T_{1,2}$ are types, and $\Gamma_{1,2}$ are typing contexts (mapping variables to types).

Definition 49 (Type-identical transformation). *A transformation τ is type-identical, if $\langle \Gamma_1 \vdash X_1 : T_1 \rangle \rightarrow_\tau \langle \Gamma_2 \vdash X_2 : T_2 \rangle$ implies $\Gamma_1 = \Gamma_2$ and $T_1 = T_2$.*

Definition 50 (Type-safe transformation). *A transformation τ is type-safe, if $\langle \Gamma_1 \vdash X_1 : T_1 \rangle \rightarrow_\tau \langle \Gamma_2 \vdash X_2 : T_2 \rangle$ implies $\Gamma_1 \vdash X_1 : T_1$ and $\Gamma_2 \vdash X_2 : T_2$.*

A.7.3 Transformation lifting

The lifting to configurations is used to relate intermediate states of execution of the original program with execution of the transformed program. During execution, which is defined using small-step semantics, the place where a transformation is located will move around as the machine state evolves. For example, when an object is instantiated, if the transformation was applied within a method of that literal, then after instantiation, the transformation will apply in the heap.

Example 51. Let's consider two configurations related by a lifted transformation τ : $\langle \Gamma \vdash \langle \text{let } x = \nu(r : R)\{m_1(r, z) = t_1\} \text{ in } t_3; \sigma; \Sigma; \rho \rangle : T \rangle \rightarrow_{\text{cfg } \tau} \langle \Gamma \vdash \langle \text{let } x = \nu(r : R)\{m_1(r, z) = t_2\} \text{ in } t_3; \sigma; \Sigma; \rho \rangle : T \rangle$, where the bodies of the method are related by τ : $\langle \Gamma, r : R, z : T_z \vdash t_1 : T_t \rangle \rightarrow_{\tau} \langle \Gamma, r : R, z : T_z \vdash t_2 : T_t \rangle$. After one reduction step, the next configurations will still be related by the lifted transformation. $\langle \Gamma \vdash \langle [w/x_2]t_3; \sigma; \Sigma, y \sim \nu(r : R)\{m(r, z) = t_1\}; \rho, w \sim y \rangle : T \rangle \rightarrow_{\text{cfg } \tau} \langle \Gamma \vdash \langle [w/x_2]t_3; \sigma; \Sigma, y \sim \nu(r : R)\{m(r, z) = t_2\}; \rho, w \sim y \rangle : T \rangle$

Term, stack, heap and configuration transformations can be defined by lifting a local term transformation, as defined in Figure 23 and Figure 24. This lifting contain premises that ensure that both sides of the transformation are well typed. The rules in these definition mimic the typing rules and heap correspondence rules, but applying to two terms, stacks, heap or configurations at the same time.

$$\begin{array}{c}
\frac{\langle \Gamma; \rho \vdash t_1 : T \rangle \xrightarrow{\tau} \langle \Gamma; \rho \vdash t_2 : T \rangle}{\Gamma; \rho \vdash T <: S} \text{(TRFL-Local)} \\
\langle \Gamma; \rho \vdash t_1 : S \rangle \xrightarrow{\text{lift } \tau} \langle \Gamma; \rho \vdash t_2 : S \rangle \\
\\
\frac{\langle \Gamma; \rho \vdash t_1 : T_1 \rangle \xrightarrow{\text{lift } \tau} \langle \Gamma; \rho \vdash t_2 : T_1 \rangle \quad \Gamma, z : T_1; \rho \vdash t_3 : T_2 \quad \Gamma; \rho \vdash T_2 <: S}{\langle \Gamma; \rho \vdash \text{let } z = t_1 \text{ in } t_3 : S \rangle \xrightarrow{\text{lift } \tau} \langle \Gamma; \rho \vdash \text{let } z = t_2 \text{ in } t_3 : S \rangle} \text{(TRFL-Let1)} \\
\\
\frac{\langle \Gamma, z : T_1; \rho \vdash t_1 : T_2 \rangle \xrightarrow{\text{lift } \tau} \langle \Gamma, z : T_1; \rho \vdash t_2 : T_2 \rangle \quad \Gamma; \rho \vdash t_3 : T_1 \quad \Gamma; \rho \vdash T_2 <: S}{\langle \Gamma; \rho \vdash \text{let } z = t_3 \text{ in } t_1 : S \rangle \xrightarrow{\text{lift } \tau} \langle \Gamma; \rho \vdash \text{let } z = t_3 \text{ in } t_2 : S \rangle} \text{(TRFL-Let2)} \\
\\
\frac{\langle \Gamma, s : R; \rho \vdash d_1 : R \rangle \xrightarrow{\text{lift } \tau} \langle \Gamma, s : R; \rho \vdash d_2 : R \rangle \quad \Gamma, z : \mu(s : R) \wedge \mathcal{M}_\perp; \rho \vdash t : T \quad \Gamma; \rho \vdash T <: S \quad R \text{ indep } s}{\langle \Gamma; \rho \vdash \text{let } z = \nu(s : R)d_1 \text{ in } t : S \rangle \xrightarrow{\text{lift } \tau} \langle \Gamma; \rho \vdash \text{let } z = \nu(s : R)d_2 \text{ in } t : S \rangle} \text{(TRFL-Lit1)} \\
\\
\frac{\langle \Gamma, z : \mu(s : R) \wedge \mathcal{M}_\perp; \rho \vdash t_1 : T \rangle \xrightarrow{\text{lift } \tau} \langle \Gamma, z : \mu(s : R) \wedge \mathcal{M}_\perp; \rho \vdash t_2 : T \rangle \quad \Gamma, s : R; \rho \vdash d : R \quad \Gamma; \rho \vdash T <: S \quad R \text{ indep } s}{\langle \Gamma; \rho \vdash \text{let } z = \nu(s : R)d \text{ in } t_1 : S \rangle \xrightarrow{\text{lift } \tau} \langle \Gamma; \rho \vdash \text{let } z = \nu(s : R)d \text{ in } t_2 : S \rangle} \text{(TRFL-Lit2)} \\
\\
\frac{\langle \Gamma; \rho \vdash d_1 : T_1 \rangle \xrightarrow{\text{lift } \tau} \langle \Gamma; \rho \vdash d_2 : T_1 \rangle \quad \Gamma; \rho \vdash d_3 : T_2 \quad d_1 \text{ and } d_3 \text{ have distinct member names} \quad d_2 \text{ and } d_3 \text{ have distinct member names}}{\langle \Gamma; \rho \vdash d_1 \wedge d_3 : T_1 \wedge T_2 \rangle \xrightarrow{\text{lift } \tau} \langle \Gamma; \rho \vdash d_2 \wedge d_3 : T_1 \wedge T_2 \rangle} \text{(TRFL-And1)} \\
\\
\frac{\langle \Gamma; \rho \vdash d_1 : T_1 \rangle \xrightarrow{\text{lift } \tau} \langle \Gamma; \rho \vdash d_2 : T_1 \rangle \quad \Gamma; \rho \vdash d_3 : T_2 \quad d_1 \text{ and } d_3 \text{ have distinct member names} \quad d_2 \text{ and } d_3 \text{ have distinct member names}}{\langle \Gamma; \rho \vdash d_3 \wedge d_1 : T_2 \wedge T_1 \rangle \xrightarrow{\text{lift } \tau} \langle \Gamma; \rho \vdash d_3 \wedge d_2 : T_2 \wedge T_1 \rangle} \text{(TRFL-And2)} \\
\\
\frac{\Gamma' = \Gamma, s : T_4, !, z : T_1, r : T_4 \wedge [r/s]T_4 \wedge T_3 \quad R = \{m(z : T_1, r : T_3) : T_2\} \quad \langle \Gamma'; \rho \vdash t_1 : T_2 \rangle \xrightarrow{\text{lift } \tau} \langle \Gamma'; \rho \vdash t_2 : T_2 \rangle}{\langle \Gamma, s : T_4; \rho \vdash \{m(z, r) = t_1\} : R \rangle \xrightarrow{\text{lift } \tau} \langle \Gamma, s : T_4; \rho \vdash \{m(z, r) = t_2\} : R \rangle} \text{(TRFL-Met)}
\end{array}$$

Figure 23: Lifting local transformations to terms

$$\begin{array}{c}
\frac{\langle \Gamma; \rho \vdash t_1 : T_1 \rangle \xrightarrow{\tau} \langle \Gamma; \rho \vdash t_2 : T_1 \rangle}{\Gamma \sim \rho \quad \Gamma; \rho \sim \Sigma} \\
\frac{\Gamma; \rho \vdash T_1 <: T_2 \quad \Gamma; \rho \vdash \sigma : T_2, S}{\langle \Gamma \vdash \langle t_1; \sigma; \Sigma; \rho \rangle : S \rangle \xrightarrow{\text{foc } \tau} \langle \Gamma \vdash \langle t_2; \sigma; \Sigma; \rho \rangle : S \rangle} \text{(TRFF-Focus)}
\end{array}$$

$$\begin{array}{c}
\frac{\langle \Gamma; \rho \vdash t_1 : T_1 \rangle \xrightarrow{\text{lift } \tau} \langle \Gamma; \rho \vdash t_2 : T_1 \rangle}{\Gamma \sim \rho \quad \Gamma; \rho \sim \Sigma} \\
\frac{\Gamma; \rho \vdash T_1 <: T_2 \quad \Gamma; \rho \vdash \sigma : T_2, S}{\langle \Gamma \vdash \langle t_1; \sigma; \Sigma; \rho \rangle : S \rangle \xrightarrow{\text{cfg } \tau} \langle \Gamma \vdash \langle t_2; \sigma; \Sigma; \rho \rangle : S \rangle} \text{(TRFC-Focus)}
\end{array}$$

$$\begin{array}{c}
\frac{\langle \Gamma; \rho \vdash \sigma_1 : T_2, S \rangle \xrightarrow{\text{stack } \tau} \langle \Gamma; \rho \vdash \sigma_2 : T_2, S \rangle}{\Gamma \sim \rho \quad \Gamma; \rho \sim \Sigma} \\
\frac{\Gamma; \rho \vdash T_1 <: T_2 \quad \Gamma; \rho \vdash t : T_1}{\langle \Gamma \vdash \langle t; \sigma_1; \Sigma; \rho \rangle : S \rangle \xrightarrow{\text{cfg } \tau} \langle \Gamma \vdash \langle t; \sigma_2; \Sigma; \rho \rangle : S \rangle} \text{(TRFC-Stack)}
\end{array}$$

$$\begin{array}{c}
\frac{\langle \rho \vdash \Sigma_1 : \Gamma \rangle \xrightarrow{\text{heap } \tau} \langle \rho \vdash \Sigma_2 : \Gamma \rangle}{\Gamma \sim \rho \quad \Gamma; \rho \vdash t : T_1} \\
\frac{\Gamma; \rho \vdash T_1 <: T_2 \quad \Gamma; \rho \vdash \sigma : T_2, S}{\langle \Gamma \vdash \langle t; \sigma; \Sigma_1; \rho \rangle : S \rangle \xrightarrow{\text{cfg } \tau} \langle \Gamma \vdash \langle t; \sigma; \Sigma_2; \rho \rangle : S \rangle} \text{(TRFC-Heap)}
\end{array}$$

Figure 24: Lifting local transformations to configurations

A.7.4 General transformation safety

Such transformation guarantees can be proven by looking at executions of the two programs (original and transformed) in parallel, and showing that the intermediate states are still related by the iterated lifted transformation (plus similarity), until both executions reach answer states. The answers will therefore also be related by that transformation, and so the two programs will have the same result in the sense of Definition 24.

As the two programs execute, the transformation will be moved around in the configuration. The ability to apply the transformation anywhere is ensured by the lifting operator, but in order for the transformation to still apply, the typing conditions need to be preserved as well as the configuration and typing context changes. For that reason, the local transformation has to have additional properties: weakening, narrowing and substitution.

These properties are analogous to the weakening, narrowing and substitution lemmas, which are a part of the soundness proof for the roDOT calculus and were adapted from the soundness proof for DOT [30]. The properties state that the transformation will relate two terms even if the typing context is extended, when a type in the typing context is refined, or when a variable is substituted in the term and its type. As an example, the weakening property is stated in Definition 52.

Definition 52 (Transformation weakening). *A transformation τ is compatible with weakening, if $\langle \Gamma_1, \Gamma_2 \vdash t_1 : T \rangle \rightarrow_\tau \langle \Gamma_1, \Gamma_2 \vdash t_2 : T \rangle$ implies $\langle \Gamma_1, \Gamma_3, \Gamma_2 \vdash t_1 : T \rangle \rightarrow_\tau \langle \Gamma_1, \Gamma_3, \Gamma_2 \vdash t_2 : T \rangle$.*

If a local transformation has these properties, then so does the lifted transformation. The weakening property can also be lifted to stacks and heaps.

Lemma 53 (Transformation weakening lifting). *If a transformation τ is compatible with weakening, then lift τ is compatible with weakening.*

Another important property of a transformation is that it *preserves answers*:

Definition 54 (Transformation preserves answers). *A transformation τ is preserves answers if when $\langle \Gamma \vdash c_1 : T \rangle \rightarrow_\tau \langle \Gamma \vdash c_2 : T \rangle$ and c_1 is an answer, then c_2 is an answer too.*

meaning that it only relates answers with answers and non-answers with non-answers.

With these properties, we can show that as the programs both execute a single step, the transformation will relate the next states, with exceptions stated below. First, the transformation may occur at more than one location after the step. For example, suppose that before a method call, the transformation occurred in the body on the heap of the method to be called. After the call, the transformation occurs both in the heap and in the focus of execution. For that reason, the transformation is iterated in the conclusion of Lemma 55. Second, when the execution has reached an answer, there are no more steps. This case is handled separately. Third, when the transformation occurs in the root of the focus, the executions may diverge for a while, but converge eventually. This case is also handled separately, and will be resolved later in the proof of Theorem 25 using Definition 56.

Lemma 55 (Transformation execution step). *Let τ be a transformation that is type-identical, type-safe, compatible with weakening, narrowing and substitution, and preserves answers.*

If $\langle \Gamma \vdash c_1 : T \rangle \rightarrow_{\text{cfg } \tau} \langle \Gamma \vdash c_2 : T \rangle$, then one of these holds:

- Both c_1 and c_2 are answers.
- $\langle \Gamma \vdash c_1 : T \rangle \rightarrow_{\text{foc } \tau} \langle \Gamma \vdash c_2 : T \rangle$
- There exist c_3 and c_4 , such that $c_1 \longrightarrow c_3$, $c_2 \longrightarrow c_4$, and $\langle \Gamma \vdash c_3 : T \rangle \rightarrow_{(\text{cfg } \tau \cup \approx)^*} \langle \Gamma \vdash c_4 : T \rangle$.

Proof. By the progress lemma, which is a part of the safety proof [12], the configurations c_1 and c_2 can each be an answer or reduce to another configuration. Because τ preserves answers, either both are answers, and the first case applies, or both are not answers, so both can progress. In that case, we do case analysis for the reduction step and the location where the transformation occurs.

- If in the root of the focus, then the second case applies.
- If in a first part of a let-term, then it moves up as the let term is split. If it applies in a first part of a let-lit-term, then it moves to the heap as part of a newly instantiated object. If it applies in a second part of a let-term, then moves to the stack.

- If in the stack, then in case of steps that change the stack, it either moves up the stack, down the stack or into the focus. In other cases, it applies the same way in the same stack, but weakening and similarity is needed in steps that create new variables.
- If in the heap, then it always stays in the heap, where weakening may be applied. In the case of a method call, it will be additionally copied to the focus, while applying substitution.

□

The case when the transformation applies in the root of the focus represents the instant when the two programs start to go in different ways, but the transformation is safe because the programs will still eventually reach similar states.

In the case of call swapping, this means that in one program, the call to m_1 is executed first, then the call to m_2 . In the transformed program, it is vice versa. During this time, the intermediate states may differ arbitrarily, but when they return from both of the calls, the results are the same regardless of the order.

Definition 56 (Eventual similarity). *A transformation τ eventually reduces to similarity, if for $\langle \Gamma_1 \vdash c_1 : T \rangle \rightarrow_{\text{loc } \tau} \langle \Gamma_2 \vdash c_2 : T \rangle$ where c_1 terminates, there exists $n \geq m$ and $c_3, c_4, \Gamma_3, \Gamma_4$, such that $c_1 \xrightarrow{n} c_3, c_2 \xrightarrow{m} c_4, \Gamma_3 \vdash c_3 : T, \Gamma_4 \vdash c_4 : T$, and $\langle \Gamma_3 \vdash c_3 : T \rangle \rightarrow_{\approx} \langle \Gamma_4 \vdash c_4 : T \rangle$.*

With swapping calls, both executions reach similarity in the same number steps, but in general, the numbers of steps n and m may differ. The condition $n \geq m$ is needed in order to avoid the case where the execution of the program on the right side would be indefinitely prolonged by increasing the number of steps of the execution over and over.

Additionally, we will use the fact that the \approx transformation is also preserved as the program executes.

Lemma 57 (Similarity compatible with reduction). *If $\langle \Gamma_1 \vdash c_1 : T_1 \rangle \rightarrow_{\approx} \langle \Gamma_2 \vdash c_2 : T_2 \rangle$ and $c_1 \xrightarrow{} c_3, c_2 \xrightarrow{} c_4$, then $\langle \Gamma_1 \vdash c_3 : T_1 \rangle \rightarrow_{\approx} \langle \Gamma_2 \vdash c_4 : T_2 \rangle$.*

The property Definition 56 is only concerned with what happens if the transformation occurs once, in the root of the focus. But with Lemma 55 and Lemma 57, we can show that for a local transformation with the properties above, if this transformation occurs anywhere in the configuration, any number of times, the answers will also differ only by occurrences of this transformation and by similarity.

Theorem (Theorem 25, Transformation execution). *If τ is a transformation that is type-identical, type-safe, compatible with weakening, narrowing and substitution, preserves answers, and eventually reduces to similarity, then $(\text{cfg } \tau \cup \approx)^*$ is safe.*

Proof. First, do induction on j , the number of steps to reach the answer. For 0 steps, the conclusion trivially holds. For $n + 1$ steps, do induction on the iteration of the transformation $(\text{cfg } \tau \cup \approx)^*$. Each step of this iteration is either similarity, or application of $\text{cfg } \tau$.

For similarity, use Lemma 57 on all steps of the execution until the answers are reached. That way, show that the answers are similar. For application of $\text{cfg } \tau$, apply Lemma 55 and do case analysis:

- In the case where c_1 and c_2 are answers, there are no more steps, so just apply the transformation between c_1 and c_2 .
- In the case where the transformation applies in the focus, use the fact that τ eventually reduces to similarity. Use Lemma 57 to apply the similarity to the answer and continue using the second inductive hypothesis with a lesser or equal number of steps.
- In the case where a reduction step is made and the transformation is preserved, apply the first inductive hypothesis with one fewer step on the rest of the execution.

□

A.8 The call-swapping transformation guarantee

In this section, we show the proof of Theorem 27, the transformation guarantee stated in Section 6.2. To prove the transformation guarantee, we first apply Definition 24 to Definition 26:

Lemma 58 (Call swapping is safe). *The transformation $\text{cfg } \text{csw}$ is safe.*

Thanks to using a local transformation csw lifted to configurations, we can prove this lemma using Theorem 25.

First, we need to prove the premises of Theorem 25, which are local properties of the local transformation csw defined in Definition 26.

The local transformation csw is trivially symmetric and type-identical, and it preserves answers, because it only applies to let-terms, which are never in the focus of an answer.

It is also type-safe and compatible with weakening, narrowing and substitution, which can be proved by applying weakening, narrowing or substitution to each of the typing and subtyping conditions in the definition of csw .

Lemma 59 (Call-swap weakening). *The transformation csw is compatible with weakening.*

The crucial property that ensures safety of the transformation is that csw satisfies Definition 56.

Lemma 60 (Call swapping eventual similarity). *The transformation csw eventually reduces to similarity.*

Proof. The proof is based on looking at the execution of let-terms and the method calls in the two programs, using the SEF guarantee to relate intermediate states by similarity, and combining the information obtained from that to relate the states after the calls by similarity. This is captured in Figure 25.

We have two configurations $c_a = \langle \text{let } x_{a1} = w_{r1}.m_1 w_{z1} \text{ in let } x_{a2} = w_{r2}.m_2 w_{z2} \text{ in } t(x_{a1}, x_{a2}); \sigma; \rho; \Sigma \rangle$ and $c_b = \langle \text{let } x_{b2} = w_{r2}.m_2 w_{z2} \text{ in let } x_{b1} = w_{r1}.m_1 w_{z1} \text{ in } t(x_{b1}, x_{b2}); \sigma; \rho; \Sigma \rangle$. Assuming that they both terminate, we need to show that they reduce to similar configurations.

c_a reduces to method call begin $\langle w_{r1}.m_1 w_{z1}; \sigma_{a1}\sigma; \rho; \Sigma \rangle$, where $\sigma_{a1} = \text{let } x_{a1} = \square \text{ in } (\text{let } x_{a2} = w_{r2}.m_2 w_{z2} \text{ in } t(x_{a1}, x_{a2}))$. Assuming termination, there must exist a method call end $\langle vw_{a1}; \sigma_{a1}\sigma; \rho\rho_{a1}; \Sigma\Sigma_{a1} \rangle$. By the SEF guarantee, in the method call end, the heap has a prefix Σ . The stack does not affect the result, so we can also say that $\langle w_{r1}.m_1 w_{z1}; \cdot; \rho; \Sigma \rangle \mapsto^{k_1} \langle vw_{a1}; \cdot; \rho\rho_{a1}; \Sigma\Sigma_{a1} \rangle$.

The execution continues by reducing to $\langle \text{let } x_{a2} = w_{r2}.m_2 w_{z2} \text{ in } t(w_{a1}, x_{a2}); \sigma; \rho\rho_{a1}; \Sigma\Sigma_{a1} \rangle$, then to $c_{a2} = \langle w_{r2}.m_2 w_{z2}; \sigma_{a2}\sigma; \rho\rho_{a1}; \Sigma\Sigma_{a1} \rangle$, where $\sigma_{a2} = \text{let } x_{a2} = \square \text{ in } t(w_{a1}, x_{a2})$.

This is a method call begin which will reduce to a method call end $\langle vw_{a2}; \sigma_{a2}\sigma; \rho\rho_{a1}\rho_{a2}; \Sigma\Sigma_{a1}\Sigma_{a2} \rangle$. Again, by the SEF guarantee, the heap in the end has a prefix $\Sigma\Sigma_{a1}$. And again, the stack does not affect the execution, so we also have $\langle w_{r2}.m_2 w_{z2}; \cdot; \rho\rho_{a1}; \Sigma\Sigma_{a1} \rangle \mapsto^{k_2} \langle vw_{a2}; \cdot; \rho\rho_{a1}\rho_{a2}; \Sigma\Sigma_{a1}\Sigma_{a2} \rangle$. With the stack removed, the method call begin does not have any references to the objects created in the previous call, so it is similar to $\langle w_{r2}.m_2 w_{z2}; \cdot; \rho; \Sigma \rangle$ and the method call end is similar to $\langle vw_{a2}; \cdot; \rho\rho_{a2}; \Sigma\Sigma_{a2} \rangle$.

The execution continues by reducing to $c_{a3} = \langle t(w_{a1}, w_{a2}); \sigma; \rho\rho_{a1}\rho_{a2}; \Sigma\Sigma_{a1}\Sigma_{a2} \rangle$. Now, we will show that c_b reduces to a similar configuration.

c_b reduces to method call begin $\langle w_{r2}.m_2 w_{z2}; \sigma_{b2}\sigma; \rho; \Sigma \rangle$, where $\sigma_{b2} = \text{let } x_{b2} = \square \text{ in } (\text{let } x_{b1} = w_{r1}.m_1 w_{z1} \text{ in } t(x_{b1}, x_{b2}))$.

If we remove the stack, we get a configuration $\langle w_{r2}.m_2 w_{z2}; \cdot; \rho; \Sigma \rangle$ for which we already know that it reduces in k_2 steps to something similar to $\langle vw_{a2}; \cdot; \rho\rho_{a2}; \Sigma\Sigma_{a2} \rangle$. That means there are $w_{b2}, \Sigma_{b2}, \rho_{b2}$ similar to $w_{a2}, \Sigma_{a2}, \rho_{a2}$, such that $\langle w_{r2}.m_2 w_{z2}; \cdot; \rho; \Sigma \rangle \mapsto^{k_2} \langle vw_{b2}; \cdot; \rho\rho_{b2}; \Sigma\Sigma_{b2} \rangle$.

Adding back the stack, we have $\langle w_{r2}.m_2 w_{z2}; \sigma_{b2}\sigma; \rho; \Sigma \rangle \mapsto^{k_2} \langle vw_{b2}; \sigma_{b2}\sigma; \rho\rho_{b2}; \Sigma\Sigma_{b2} \rangle$.

The execution continues by reducing to $\langle \text{let } x_{b1} = w_{r1}.m_1 w_{z1} \text{ in } t(x_{b1}, w_{b2}); \sigma; \rho\rho_{b2}; \Sigma\Sigma_{b2} \rangle$, then to $c_{b1} = \langle w_{r1}.m_1 w_{z1}; \sigma_{b1}\sigma; \rho\rho_{b2}; \Sigma\Sigma_{b2} \rangle$, where $\sigma_{b1} = \text{let } x_{b1} = \square \text{ in } t(x_{b1}, w_{b2})$.

If we remove the stack, we get a configuration $\langle w_{r1}.m_1 w_{z1}; \cdot; \rho\rho_{b2}; \Sigma\Sigma_{b2} \rangle$. With the stack removed, this method call begin does not have any references to the objects created in the previous call, so it is similar to $\langle w_{r1}.m_1 w_{z1}; \cdot; \rho; \Sigma \rangle$, for which we already know that it reduces in k_1 steps to something similar to $\langle vw_{a1}; \cdot; \rho\rho_{a1}; \Sigma\Sigma_{a1} \rangle$. That means there are $w_{b1}, \Sigma_{b1}, \rho_{b1}$ similar to $w_{a1}, \Sigma_{a1}, \rho_{a1}$, such that $\langle w_{r1}.m_1 w_{z1}; \cdot; \rho; \Sigma \rangle \mapsto^{k_1} \langle vw_{b1}; \cdot; \rho\rho_{b1}; \Sigma\Sigma_{b1} \rangle$.

Adding back the stack and heap, we have $\langle w_{r1}.m_1 w_{z1}; \sigma_{b1}\sigma; \rho\rho_{b2}; \Sigma\Sigma_{b2} \rangle \mapsto^{k_1} \langle vw_{b1}; \sigma_{b1}\sigma; \rho\rho_{b2}\rho_{b1}; \Sigma\Sigma_{b2}\Sigma_{b1} \rangle$.

The execution continues by reducing to $c_{b3} = \langle t(w_{b1}, w_{b2}); \sigma; \rho\rho_{b1}\rho_{b2}; \Sigma\Sigma_{b2}\Sigma_{b1} \rangle$. Because the variables and heaps in c_{b3} were chosen to be similar to the variables and heaps in c_{a3} , we get $c_{a3} \approx c_{b3}$. \square

Applying Theorem 25 on csw gives us that $(\text{cfg } \text{csw} \cup \approx)^*$ is safe. That means $\text{cfg } \text{csw}$ is safe as well, because $\text{cfg } \text{csw} \subset (\text{cfg } \text{csw} \cup \approx)^* = ((\text{cfg } \text{csw} \cup \approx)^* \cup \approx)^*$. This concludes the proof of Lemma 58.

Theorem (Theorem 27, Transformation guarantee). *If $\langle \vdash t_1 : T \rangle \rightarrow_{\text{lift } \text{csw}} \langle \vdash t_2 : T \rangle$ and $\langle t_1; \cdot; \cdot \rangle \rightarrow^k c_3$, where c_3 is an answer typed as $\Gamma_3 \vdash c_3 : T$, then there exists c_4, Γ_4 and j such that $\langle t_2; \cdot; \cdot \rangle \rightarrow^j c_4$, c_4 is an answer typed as $\Gamma_4 \vdash c_4 : T$ and $\langle \Gamma_3 \vdash c_3 : T \rangle \rightarrow_{(\text{cfg } \text{csw} \cup \approx)^*} \langle \Gamma_4 \vdash c_4 : T \rangle$.*

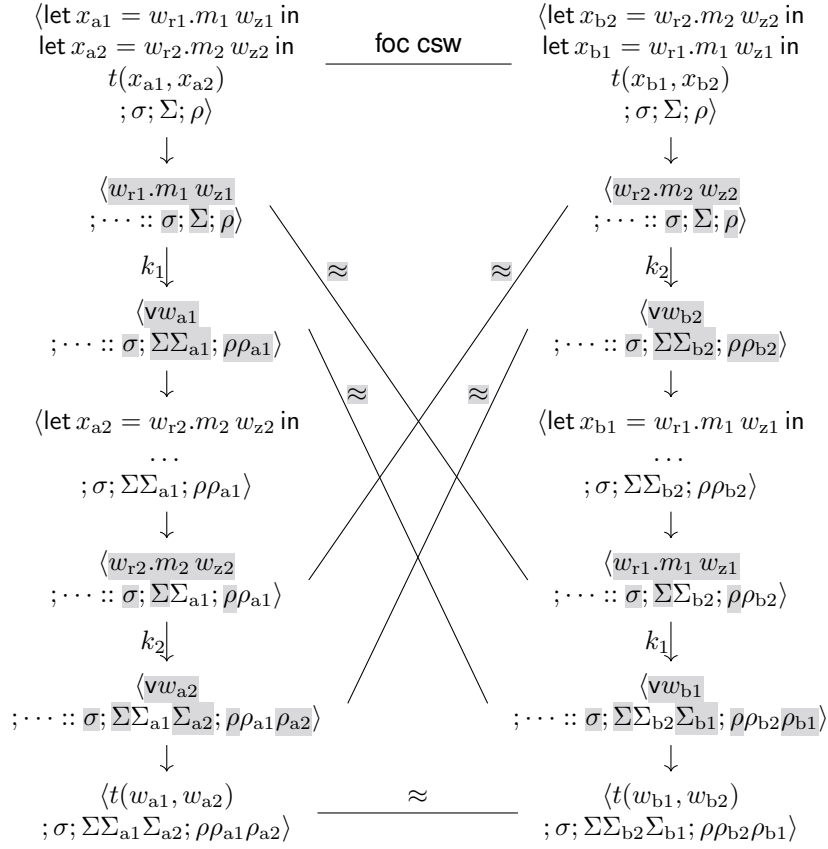


Figure 25: Reduction and transformation of call swapping in focus

Proof. Follows directly by unfolding definitions in Lemma 58 and specializing to initial configurations. \square