Leveraging Program Analysis for Type Inference

by

Zvonimir Pavlinovic

A dissertation submitted in partial fulfillment of the requirements for the degree of Doctor of Philosophy Department of Computer Science New York University May 2019

Professor Thomas Wies

Acknowledgements

I would like to thank my mother, father, sister, and late grandparents for their undeniable support in pursuing my professional goals. Without their unconditional love, I would have never reached this point in my career that initially seemed as nothing more than a childish dream. This dissertation could have not been written without the guidance of my adviser, Thomas Wies. I am sincerely thankful for an unparalleled patience and countless hours spent helping me understand the phenomena of computation and the very core of being a scientist. To my friends, I thank you for being there for me despite my frequent absence and unavailability. To all the other people that helped me in any capacity during the course of my studies, I am sincerely grateful.

Abstract

Type inference is a popular feature of programming languages used to automatically guarantee the absence of certain execution errors in programs at compile time. The convenience of type inference, unfortunately, comes with a cost. Developing type inference algorithms is a challenging task that currently lacks a systematic approach. Moreover, programmers often have problems interpreting error reports produced by type inference. The overarching goal of this thesis is to provide a mathematically rigorous framework for the systematic development of sophisticated type inference algorithms that are convenient to use by the programmers. To this end, we focus on two specific problems in this thesis: (1) how to constructively design type inference algorithms that improve over the state-of-the-art and (2) how to automatically debug type errors that arise during inference. We base our approach on the observation that, similar to type inference, program analysis algorithms automatically discover various program properties that can be used to show program correctness. Type inference and program analysis techniques, although similar, have traditionally been developed independently of each other. In contrast, this thesis further explores the recent path of leveraging program analysis for type inference.

As our first contribution, we use abstract interpretation to constructively design type inference algorithms. We specifically focus on Liquid types, an advanced family of algorithms that combine classical typing disciplines and known static analyses to prove various safety properties of functional programs. By using abstract interpretation, we make the design space of Liquid type inference explicit. We also unveil the general type inference framework underlying Liquid types. By properly instantiating this general framework, one obtains novel type inference algorithms that are sound by construction.

Our second contribution is a framework for automatically debugging type errors for languages that deploy type inference in the style of Hindley-Milner, such as OCaml and Haskell. Such languages are notorious for producing cryptic type error reports that are often not helpful in fixing the actual bug. We formulate the problem of finding the root cause of type errors as an optimization problem expressed in a formal logic. We then show how to solve this problem using automated theorem provers. We experimentally illustrate how our framework can efficiently produce type error reports that outperform the state-of-the-art solutions in identifying the true cause of type errors.

In summary, this thesis introduces a mathematical framework for the systematic design of sophisticated type inference algorithms that are sound by construction. Our results further enable automatic generation of more meaningful type error diagnostics, ultimately making type inference more usable by the programmers.

Contents

	Ackn	owledgements	iii
	Abst	ract	iv
	List	of Figures	viii
	List	of Tables	х
	List	of Appendices	xi
1	Intro	oduction	1
	1.1	Data Flow Refinement Type Inference	4
	1.2	Type Error Localization	8
	1.3	Road Map	11
2	Prel	iminaries	13
	2.1	Notation	13
	2.2	Language	14
	2.3	Types	16
	2.4	Abstract Interpretation	17
3	Data	a Flow Refinement Type Inference	20
	3.1	Overview	20
	3.2	Data Flow Semantics	29

	3.3	Relational Data Flow Semantics		
	3.4	Collapsed Relational Data Flow Semantics	57	
3.5 Data Flow Refinement Type Semantics		Data Flow Refinement Type Semantics	67	
	3.6	Liquid Type Semantics	75	
	3.7	Related Work	80	
4	Typ	be Error Localization	85	
	4.1	Overview	85	
	4.2	Problem	93	
	4.3	Algorithm	95	
	4.4	Implementation and Evaluation	104	
	4.5	Taming The Exponential Explosion	109	
	4.6	Related Work	130	
5	Conclusions 1		.34	
Appendices			.36	
Bibliography			64	

List of Figures

Series of abstract semantics leading up to Liquid type inference		
Subsequent maps computed by data flow semantics for program		
from Example 1. Numbers above maps are used to identify the maps.	30	
Concrete transformer for data flow semantics	39	
Value propagation in the concrete data flow semantics		
Operations on relational abstract domains		
Value propagation in the relational data flow semantics	54	
Abstract transformer for relational data flow semantics	55	
Abstract transformer for collapsed semantics (recursive function def-		
inition rule)	64	
Propagation between values in the collapsed semantics	65	
) A structural definition of strengthening operations on refinement		
types	70	
Abstract transformer for data flow refinement semantics	72	
Adapted Liquid typing rules	78	
High-level overview of constraint-based type error localization. Thick arrows represent a looping interaction between a compiler and the		
SMT solver.	86	
	Series of abstract semantics leading up to Liquid type inference Subsequent maps computed by data flow semantics for program from Example 1. Numbers above maps are used to identify the maps. Concrete transformer for data flow semantics Value propagation in the concrete data flow semantics Operations on relational abstract domains Value propagation in the relational data flow semantics Abstract transformer for relational data flow semantics Abstract transformer for collapsed semantics (recursive function def- inition rule)	

4.2	Typing rules for λ^{\perp}
4.3	Rules defining the constraint typing relation for λ^{\perp} 100
4.4	Labeled abstract syntax tree for the program $p \dots $
4.5	Quality of AST ranking criterion compared to OCaml's type checker
	in pinpointing the true error source in 20 benchmark programs 107
4.6	Maximum, average, and minimum execution times for computing a
	minimum error source
4.7	Maximum, average, and minimum number of generated assertions
	(in thousands) for computing a minimum error source 108
4.8	Rules defining the refined constraint typing relation for λ^\perp 112
4.9	Maximum, average, and minimum number of typing assertions for
	computing a minimum error source by naive and iterative approach 125
4.10	Maximum, average, and minimum execution times for computing a
	minimum error source by naive and iterative approach $\ldots \ldots \ldots 126$
4.11	Maximum, average, and minimum number of typing assertions for
	computing a minimum error source by naive and iterative approach
	for larger programs
4.12	Maximum, average, and minimum execution times for computing
	a minimum error source by naive and iterative approach for larger
	programs

List of Tables

4.1	Statistics for the number of expansions and iterations when com-	
	puting a single minimum error source	127
4.2	Statistics for the number of expansions and iterations when com-	

puting a single minimum error source for larger programs 128 $\,$

List of Appendices

A	Data Flow Refinement Type Inference		
	A.1	Example Concrete Map	136
	A.2	Concrete Semantics Proofs	138
	A.3	Path Data Flow Semantics	158
	A.4	Relational Semantics Proofs	202
	A.5	Collapsed Semantics Proofs	227
	A.6	Data Flow Refinement Type Semantics Proofs	240
	A.7	Liquid Types Semantics Proofs	246
В	Тур	e Error Localization	255
	B.1	Proof of Lemma 25	255

Chapter 1

Introduction

Types are an integral part of modern programming languages. The role of types in a programming language is to facilitate the usage of powerful programming abstractions and, perhaps more importantly, enhance the reliability of programs. Types can be used to express information about program executions capturing the form of values being used by the program and the behavior of functions that operate over these values. For instance, types can capture program invariants stating that certain program variables always hold integer values during the execution or that functions implementing the integer division operation accept two integer arguments and produce an integer value. Programmers use types to communicate such information about their programs to the compiler that in turn uses this information to ensure the absence of certain runtime errors. Some programming languages, such as OCaml [70] and Haskell [38], automatically discover types in programs by employing type inference algorithms [74]. By using type inference, programming languages are therefore able to provide certain correctness guarantees for programs at compile time without requiring major programmer intervention.

Type inference algorithms can be seen as a form of static analysis of programs [15, 17, 16]. Both type inference and static analyses automatically deduce information about programs with the goal of showing program correctness. In type inference, programs are deemed correct if they do not exhibit a target class of execution errors, such as calling a Boolean negation operation with a string argument [74]. A deeper connection between type inference and static analyses can be observed in the way they both deal with the inherent infeasibility of automatically deciding program correctness. Static analyses approach this problem by sacrificing precision for soundness, i.e., by over-approximating how programs behave [17, 16]. That is, the behaviors of a program inferred by a static analysis always include the actual program behaviors, but they also might include same spurious behaviors. Similarly, type inference algorithms sacrifice precision by flagging some correct programs as potentially erroneous, but they ensure soundness by never classifying an erroneous program as correct. Despite many similarities [15], type inference and static analysis techniques are typically developed independently of each other. Type inference algorithms rarely reap many of the benefits provided by the advances in static analysis and program analysis in general [17, 68]. This is unfortunate, particularly since type inference algorithms and the process of their creation taken by both industry and academia have several notable shortcomings.

The presentations of many type inference algorithms found in the literature [40, 59, 78] and practical implementations [70, 38, 92] only argue the algorithm's soundness. Such works do not formally discuss the precision of the type inference being presented. The particular points where these algorithms can be made more precise, generalized, or how they can be compared against each other are not made formally explicit. Overall, the design space of existing type inference algorithms is mathematically unclear. Developing type inference algorithms is consequently a challenging problem lacking a more systematic approach. This problem is even more severe in the context of modern programming languages that are facing demands for providing type inference with an ever increasing precision and for increasingly complex classes of errors whose absence they are supposed to guarantee. The other shortcoming associated with type inference concerns its usability. When type inference flags a given program as possibly erroneous by discovering a type error, the compiler-generated messages explaining the error are often not helpful in fixing the actual bug. The produced error explanation frequently identifies the source of the error incorrectly, thus increasing the debugging times for the programmer [54]. Even worse, cryptic type error reports can significantly hinder the process of learning languages with type inference for novice programmers [54]. The goal of this thesis is to reduce the cost of developing and using sophisticated type inference algorithms in modern programming languages by addressing the problems of (1) how to constructively design type inference algorithms that improve over the state-of-the-art and (2) how to automatically debug type errors that arise during inference. We base our approach on the close connection between static analysis and type inference described earlier.

We make two key technical contributions in this thesis. First, we use abstract interpretation [17, 16] to systematically construct Liquid type inference [92], a recently developed algorithm based on classical type inference [40, 59] and predicate abstraction [4, 27] that automatically infers rich properties of higher-order programs. Abstract interpretation is a mathematically rigorous framework for constructively designing static analyses as abstract semantics of programs, i.e., as over-approximations of the actual program behaviors. By formulating Liquid type inference as a static analysis using abstract interpretation, we unveil the design space of such inference algorithms. Our formulation makes explicit the points where Liquid type inference loses precision and can hence be improved. Moreover, our results allow generalization of Liquid type inference and facilitate the development of powerful type inference algorithms that are sound by construction.

As a second contribution, we formally illustrate how the problem of type error debugging for languages such as OCaml and Haskell can be cast as an optimization problem that neatly separates two constituent subproblems: devising a heuristic for pinpointing the root source of a type error and searching for the root source given such a heuristic. We express the optimization problem of type error debugging using formal logic and then solve it using automated theorem provers [21, 5]. By relying on recent advances in automated theorem proving, we obtain a general, effective, and efficient framework for type error debugging.

We now discuss each of our two technical contributions in more detail.

1.1 Data Flow Refinement Type Inference

Classical type inference algorithms, such as those found in OCaml and Haskell, guarantee the absence of a certain class of runtime errors. For instance, such algorithms can capture that an array read, read a i, will always operate on the base pointer a to some memory region that holds an array. However, to show that the array read is always within the array bounds, the types need to relate the index i at which the read occurs and the length of the array a. Automatically inferring such types is out of the scope of classical type systems. Refinement types take an exciting direction to designing static type inference algorithms that can infer fine-grained program invariants [28, 95, 78, 92, 99, 93]. Here, the types of a classical type system are augmented with auxiliary *type refinements* that further constrain the members of the basic type.

Liquid types [78, 91, 92, 90] form an important family of refinement type systems. Here, basic types are augmented with *refinement predicates* that are drawn from some decidable logic. These predicates can express relational dependencies between inputs and outputs of functions. For example, the contract of the array read operator, **read**, can be expressed using the Liquid type

read ::
$$(a : \alpha \text{ array}) \rightarrow (i : \{\nu : \text{int} \mid 0 \leq \nu < \text{length } a\}) \rightarrow \alpha$$

This type indicates that read is a function that takes an array a over some element type α and an index i as input and returns a value of type α . The type $\{\nu : \text{int } | 0 \leq \nu < \text{length } a\}$ of the parameter i refines the base type int to indicate that imust be an index within the bounds of the array a. Liquid type inference uses this type to check that array accesses in a given program are within the array bounds by automatically inferring refinement types for all program expressions. For instance, consider the following illustrative OCaml program.

```
let f a i = read a (i - 1) in
let b = f [| 1; 2 |] 2 in
let c = f [| 3; 4; 5 |] 3 in
b + c
```

The above program defines a function f that accepts an integer array a, an integer i, and returns the element of a at position i - 1. This function is then called two times with integer arrays and integers whose value equals the length of the respective arrays. Liquid type inference first calls the OCaml type inference to

get the basic types of program expressions. For every expression, the Liquid type inference then infers a conjunction of auxiliary refinement predicates drawn from a given finite set of candidate predicates. The inferred conjunction refines the basic type of the expression and is consistent with the way the expression is used in the program. For instance, one possible inferred Liquid type for f is

$$f :: (a : int array) \rightarrow (i : \{\nu : int \mid \nu = \text{length } a \land 2 \leq \nu\}) \rightarrow int$$

The above type correctly states that f is always called with an integer array a whose length is equal to the value of the argument i in the above program. The value of i is also correctly inferred to be greater or equal than 2. Hence, the actual access index i - 1 is always provably within the bounds of a for the above program. This illustrative example showcases an interesting feature of Liquid type inference. The type of f is not inferred based on the body of f in isolation, as is the case with more classical type inference [40, 59], but also on the remainder of the program: the types for parameters a and i are deduced based on how f is called in the program. The algorithm discovers that the values *flowing* to the two arguments of f are always an integer array a and an integer i such that i is in the interval $[2,\infty]$ and i = length a. Intuitively, this type expresses an invariant of the program capturing how data flows to and from f.

Approach. In this thesis, we formally model the distinguishing features of Liquid type inference with the goal of generalizing it and unveiling the design space of such analyses. We shed new light on Liquid type inference in terms of abstract interpretation, a framework that enables the systematic construction of static analyses as abstractions of a concrete program semantics [16, 17]. First, we propose a new semantics of higher-order functional programs that precisely captures the program properties abstracted by Liquid types. Our new semantics explicitly models the data flow in higher-order programs that is over-approximated by Liquid types. To the best of our knowledge, this is the first semantics where the notion of data flow is made explicit in a higher-order setting. We then systematically construct a *data flow refinement type analysis* via a sequence of Galois abstractions [16] of this concrete semantics. The resulting analysis infers valid refinement types for all subexpressions of a program and is parametric in the abstract domain used to express the type refinements. Finally, we show that Liquid type inference is a particular instance of this parametric analysis where the abstract domain of type refinements is defined by the particular choice of candidate refinement predicates made by the Liquid type inference.

Contributions. Our technical development brings several benefits. First, the points where Liquid type inference loses precision are made explicit in (1) the Galois connections used in the abstractions building up to our data flow refinement type analysis and (2) the way in which the convergence of the analysis is enforced, which we formally capture as widening operators [16, 17]. These operators are a well-known mechanism used in abstract interpretation to ensure the sound termination of static analyses. Next, we generalize from the specific choice of logical predicates as type refinements used in Liquid types by allowing for the use of arbitrary abstract domains (including e.g. polyhedra [83, 19] and automatabased domains [3, 50]). Instantiations of our parametric type analysis are sound by construction. Finally, the type inference analyses that build on our results can be more easily compared against each other and used in combination with other static analyses (e.g., via reduced products [17]).

1.2 Type Error Localization

Another important aspect of type inference besides its systematic construction is usability. For a type inference to be usable, it is important that the programmers understand the results the inference produces. This particularly applies for the cases where type inference rejects a program and reports a type error. The quality of the error report has a direct impact on the programmer's ability to fix the bug [54]. In this thesis, we focus on the problem of type error debugging for Hindley-Milner type inference [40, 59] found in languages such as OCaml and Haskell. Our focus on the Hindley-Milner algorithm is motivated by its widespread usage [70, 38] and often unsatisfactory error reports that increase debugging times for the programmers [54]. Moreover, addressing the problem of type error debugging for the Hindley-Milner type inference is an essential step in improving type error reports for the Liquid type inference, as the latter algorithm relies on the results produced by the Hindley-Milner algorithm [78].

Hindley-Milner type inference reports type errors on the fly. If the inferred type of the current program expression conflicts the inferred type of its context, the inference algorithm immediately stops and reports an error at the current program location. Although fast in practice, this approach also produces poor error diagnostics. In particular, it might be the case that the programmer made a mistake with the previous usages of the offending expression, or with some other related expressions. For example, consider the following simple OCaml program taken from the student benchmarks in [54]:

```
1 type 'a lst = Null | Cons of 'a * 'a lst
2 let x = Cons(3, Null)
```

3 let _ = print_string x

The standard OCaml compiler [70] reports a type mismatch error for expression **x** on line 3, as the code before that expression is well typed. However, perhaps the programmer defined **x** incorrectly on line 2 or misused the **print_string** function. As **x** is defined just before the error line, it seems more likely that the error is caused by a misuse of **print_string**. In fact, the student author of this code confirmed that this is the real source of the error. This simple example suggests that in order to generate useful error reports, compilers should consider several possible error causes and rank them by their relevance. Hence, there is a need for an infrastructure that can supply compilers with error sources that best match their relevance criteria. This thesis proposes a general algorithm based on constraint solving that provides this functionality.

Approach. Unlike typical type inference algorithms, we do not simply report the location of the first observed type inconsistency. Instead, we compute all minimum sets of expressions each of which, once corrected, yields a type correct program. Compilers can then use these computed sets for generating more meaningful error reports or even for providing automatic error correction. The considered optimality criterion is controlled by the compiler. For example, the compiler may only be interested in those error causes that require the fewest changes to fix the program.

The crux of our approach is to reduce type error localization to the weighted maximum satisfiability modulo theory (MaxSMT) problem. Specifically, our algorithm builds on existing work that rephrases type inference in terms of constraint solving [87, 1, 71]. Each program expression is assigned a type variable and typing information is captured in terms of constraints over those variables. If an input program has a type error, then the corresponding set of typing constraints is unsatisfiable. We encode the compiler-specific ranking criterion by assigning weights to the generated typing constraints. A weighted MaxSMT solver then computes the satisfiable subsets of the constraints that have maximum cumulative weight. As constraints directly map to program expressions, the complements of these maximum sets represent minimum sets of program expressions that may have caused the type error.

One issue with our constraint-based approach to the type error debugging problem is that the constraints for polymorphic functions need to be freshly copied for each usage of such functions, potentially resulting in an exponential number of constraints [74]. We hence also propose an improved algorithm that deals with the inherent complexity of polymorphic typing [57]. Our new algorithm makes the optimistic assumption that the relevant type error sources only involve few polymorphic functions, even for large programs. The new algorithm abstracts polymorphic functions by principal types [40] in such a way that all potential error sources involving the definition of an abstracted function are represented by a single error source whose cost is smaller or equal to the cost of all these potential error sources. If a computed error source does not involve a usage of such an abstracted polymorphic function, the minimum error source has been found. Otherwise, we expand the instantiations of the principal type of the function involved in the discovered error source to the corresponding typing constraints and repeat the whole process of finding the minimum error source. In the worst case scenario, the algorithm terminates by solving the constraint set where all usages of polymorphic functions are expanded, ending up with the same constraint set as the naive algorithm.

Contributions. The benefits of our approach are multifold. Our approach clearly separates the problems of devising a heuristic for pinpointing the actual source of a type error, expressed as a ranking criterion, and the problem of searching for the best error source given the criterion, solved by MaxSMT solvers. The use of SMT solvers has several additional advantages. First, it allows support for a variety of type systems by instantiating the MaxSMT solver with an adequate reasoning theory. In particular, the typing constraints for the Hindley-Milner type systems [40, 59] can be encoded in the theory of inductive data types [7]. Second, the framework does not introduce a substantial implementation overhead since the SMT solver can be used as a black box. Finally, we experimentally illustrate how our type error debugging framework is able to incorporate sophisticated ranking criteria that outperform standard type inference implementations in finding the true source of type errors while being efficient even for programs having several thousands of lines of code. To the best of our knowledge, no other type error debugging technique that formally guarantees the optimality of the computed error sources has comparable performance.

1.3 Road Map

In summary, the results presented in this thesis enable a systematic development of powerful and programmer-friendly type inference algorithms. The presentation of these results in this thesis is structured as follows. We start by introducing the programming language, notation, and definitions we use throughout the thesis in Chapter 2. We present our formulation of Liquid type inference as abstract interpretation in Chapter 3. After providing an overview of the technical contribution in § 3.1, we introduce our novel concrete data flow semantics of untyped higher-order languages in § 3.2 and build several subsequent abstract semantics in § 3.3, § 3.4, and § 3.5. We finally present Liquid type semantics of functional programs in § 3.6. Chapter 4 describes our solution to the problem of type error debugging aimed at type inference algorithms in the style of Hindley-Milner. We first overview our approach in § 4.1. Then, we formalize the problem of type error localization in § 4.2, present the solution in § 4.3, and discuss our implementation and evaluation in § 4.4. § 4.5 presents an improved type error localization algorithm that effectively deals with the inherent complexity of polymorphic typing. We conclude the thesis in Chapter 5.

Chapter 2

Preliminaries

We first introduce notation and definitions that we use throughout the thesis.

2.1 Notation

We often use meta-level let ... in ... and conditional if ... then ... else ... constructs in mathematical definitions. The context and the notation should make it clear when these constructs are used on the meta-level and as actual constructs of the programming language of investigation. We will often compress consecutive let statements let x = ... in let y = ... in ... as let x = ...; y = ... in

We use capital lambda notation $(\Lambda x. ...)$ for defining mathematical functions. For a function $f: X \to Y, x \in X$, and $y \in Y$, we use f.x: y and $f[x \mapsto y]$ to denote a function that maps x to y and otherwise agrees with f on every element of $X \setminus \{x\}$. For a set X we denote its powerset by $\wp(X)$. For a relation $R \subseteq X \times Y$ over sets X, Y and a natural number n > 0, we use \dot{R}_n to refer to the pointwise lifting of R to a relation on n-tuples $X^n \times Y^n$. That is $\langle (x_1, \ldots, x_n), (y_1, \ldots, y_n) \rangle \in \dot{R}_n$ iff $\bigwedge_{1 \le i \le n} (x_i, y_i) \in R$. Similarly, for any nonempty set Z we denote by \dot{R}_Z the pointwise lifting of R to a relation over $(Z \to X) \times (Z \to Y)$. More precisely, if $f_1 : Z \to X$ and $f_2 : Z \to Y$, then $(f_1, f_2) \in \dot{R}_Z$ iff $\forall z \in Z. (f_1(z), f_2(z)) \in R$. Typically, we drop the subscripts from these lifted relations when they are clear from the context or use an alternative "dot" notation. For instance, the pointwise lifting \dot{R}_2 of a relation $R : X \times Y$, for some sets X and Y, is also denoted by \ddot{R} .

For sets X, Y, and a function $d: X \to \wp(Y)$, we use the notation $\Pi x \in X.d(x)$ to refer to the set $\{f: X \to Y \mid \forall x \in X. f(x) \in d(x)\}$ of all dependent functions with respect to d. Similarly, for given sets X and Y we use the notation $\Sigma x \in$ X.d(x) to refer to the set $\{\langle x, y \rangle : X \times Y \mid y \in d(x)\}$ of all dependent pairs with respect to d. We use the operators π_1 and π_2 that project to the first, respectively, second component of a pair. For set comprehensions $\{t_1 \mid t_2\}$ we assume that free variables of the term t_2 are existentially quantified unless they are free in t_1 .

2.2 Language

For exposition purposes, our formal presentations are built around an idealized language: a lambda calculus with ML-style let polymorphism. The syntax of our language is defined as follows:

Expressions	e ::= x	variables
	c	constants
	$\mid \mu f. \lambda x. e$	(recursive) lambda abstractions
	$ e_1 e_2$	applications
	$ e_0 ? e_1 : e_2$	conditionals
	$ $ let $x = e_1$ in e_2	let bindings

The language supports constants $c \in Cons$ and recursive lambda abstractions. We assume the set of constants includes Booleans $b \in Bool \stackrel{\text{def}}{=} \{true, false\}$, which we will use to assign meaning to conditional expressions. In addition to conditional expressions and the usual constructs of the lambda calculus, our language also supports let bindings. Specifically, we use let bindings to define polymorphic functions.

Variables $x \in Vars$ and $f \in Vars$ are drawn from an infinite set that is disjoint from all other syntactic constructs. An expression e is *closed* if all using occurrences of variables within e are bound in definitions of recursive functions $\mu f \lambda x.e$.

Let e be an expression. Each subexpression of e is uniquely annotated with a location $\ell \in Loc$. We use Loc(e) to denote the set of all locations of e. A location $\ell \in Loc(e)$ hence uniquely identifies a subexpression of e. We denote this subexpression by $e(\ell)$. We use superscript notation to indicate these locations in (sub)expressions as in $(e_1^{\ell_1} e_2^{\ell_2})^{\ell}$ and $(\mu f.\lambda x. e_1^{\ell_1})^{\ell}$. Sometimes, we also use $e' \in e$ notation to denote the fact that e' is a subexpression of e. The label annotations are omitted whenever possible to avoid notational clutter. Variables are also locations $Vars \subseteq Loc$. A program is a closed expression where all location labels are unique.

Next, we define a mapping Uloc for the usage locations of a variable. Formally, Uloc is a partial function such that given a location ℓ of a let variable definition and a program p returns the set $Uloc_p(\ell)$ of all locations where this variable is used in p. Note that a location of a let variable definition is a location corresponding to the root of the defining expression. We also make use of a function for the definition location dloc. The function dloc reverses the mapping of Uloc for a variable usage. More precisely, $dloc(p, \ell)$ returns the location where the variable appearing at ℓ was defined in p. Also, for a set of locations L we define $Vloc(\ell)$ to be the set of all locations in $Loc(\ell)$ that correspond to usages of let variables. We assume the above sets are precomputed for a given program. We do not provide detailed algorithms for computing them as they are straightforward.

2.3 Types

The types of our language are as follows:

$$\begin{array}{ll} \mathbf{Monotypes} & \tau ::= \mathsf{bool} \mid \mathsf{int} \mid \alpha \mid \tau \to \tau \\ \mathbf{Polytypes} & \sigma ::= \tau \mid \forall \alpha. \sigma \end{array}$$

Monotypes τ include the base types **bool** and **int**, type variables α , which are drawn from an infinite set disjoint from the other types (as well as locations and program expressions), and function types $\tau \to \tau$. A monotype in which no type variables occur is called ground. A polytype $\forall \alpha.\sigma$ represents the intersection of all types obtained by instantiating the type variable α in σ by a ground monotype. That is, $\forall \alpha.\sigma$ binds α . We write $\forall \vec{\alpha}.\tau$ as a shorthand for $\forall \alpha_1....\forall \alpha_n.\tau$ where $\vec{\alpha} =$ $\alpha_1, \ldots, \alpha_n$. We further denote by $fv(\sigma)$ the set of free type variables in type σ . Finally, we write $\sigma[\vec{\beta}/\vec{\alpha}]$ for capture-avoiding substitution of free occurrences of the type variables $\vec{\alpha}$ in σ by the type variables $\vec{\beta}$.

We define typing environments Γ as maps from variables to types. We write \emptyset for the empty typing environment and extend the function fv to typing environments in the expected way.

2.4 Abstract Interpretation

A partially ordered set (poset) is a set L equipped with a relation \sqsubseteq that is (1) reflexive, (2) transitive, and (3) antisymmetric. We use the notation (L, \sqsubseteq) for such a poset.

Galois connections. Let (L_1, \sqsubseteq_1) and (L_2, \sqsubseteq_2) be two posets. We say a pair of functions $\langle \alpha, \gamma \rangle$ with $\alpha \in L_1 \to L_2$ and $\gamma \in L_2 \to L_1$ forms a *Galois connection* iff

$$\forall x \in L_1, \forall y \in L_2. \ \alpha(x) \sqsubseteq_2 y \iff x \sqsubseteq_1 \gamma(y) .$$

This fact is denoted by $(L_1, \sqsubseteq_1) \xleftarrow{\gamma}{\alpha} (L_2, \sqsubseteq_2)$. We call L_1 the concrete domain and L_2 the abstract domain of the Galois connection. Similarly, α is called abstraction function (or left adjoint) and γ concretization function (or right adjoint). Intuitively, $\alpha(x)$ is the most precise approximation of $x \in L_1$ in L_2 while $\gamma(y)$ is the least precise element of L_1 that can be approximated by $y \in L_2$.

A complete lattice is a tuple $\langle L, \sqsubseteq, \bot, \top, \sqcup, \sqcap \rangle$ where (L, \sqsubseteq) is a poset such that for any $X \subseteq L$, the least upper bound $\sqcup X$ (join) and greatest lower bound $\sqcap X$ (meet) with respect to \sqsubseteq exist. In particular, we have $\bot = \sqcap L$ and $\top = \sqcup L$. We will often identify a complete lattice with its carrier set L. Some of our definitions for a join (meet) operator will be defined as a binary operator $\sqcup \subseteq L \times L$ ($\sqcap \subseteq L \times L$) that is lifted to arbitrary subset of L when possible.

Let $(L_1, \sqsubseteq_1, \bot_1, \top_1, \sqcup_1, \sqcap_1)$ and $(L_2, \sqsubseteq_2, \bot_2, \top_2, \sqcup_2, \sqcap_2)$ be two complete lattices. Let the pair $\langle \alpha, \gamma \rangle$, where $\alpha \in L_1 \to L_2$ and $\gamma \in L_2 \to L_1$, form a Galois connections. Each function in the pair uniquely determines the other:

$$\alpha(x) = \sqcap_2 \{ y \in L_2 \mid x \sqsubseteq_1 \gamma(y) \}$$

$$\gamma(y) = \sqcup_1 \{ x \in L_1 \mid \alpha(x) \sqsubseteq_2 y \}$$

Also, α is a complete join-morphism

$$\forall S \subseteq L_1. \ \alpha(\sqcup_1 S) = \sqcup_2 \{ \alpha(x) \mid x \in S \}, \alpha(\bot_1) = \bot_2$$

and γ is a complete meet-morphism

$$\forall S \subseteq L_2. \ \gamma(\sqcap_2 S) = \sqcap_1 \{ \gamma(x) \mid x \in S \}, \gamma(\top_2) = \top_1.$$

Proposition 1 ([17]). The following statements are equivalent:

- 1. $\langle \alpha, \gamma \rangle$ is a Galois connection
- 2. α and γ are monotone, $\alpha \circ \gamma$ is reductive: $\forall y \in L_2$. $\alpha(\gamma(y)) \sqsubseteq_2 y$, and $\gamma \circ \alpha$ is extensive: $\forall x \in L_1$. $x \sqsubseteq_1 \gamma(\alpha(x))$
- 3. γ is a complete meet-morphism and $\alpha = \Lambda x \in L_1$. $\prod_2 \{ y \in L_2 \mid x \sqsubseteq_1 \gamma(y) \}$
- 4. α is a complete join-morphism and $\gamma = \Lambda y \in L_2$. $\bigsqcup_1 \{ x \in L_1 \mid \alpha(x) \sqsubseteq_2 y \}$.

Proposition 2 ([17]). α is onto iff γ is one-to-one iff $\alpha \circ \gamma = \lambda y. y.$

Widening operators. A widening operator for a complete lattice $\langle L, \sqsubseteq, \bot, \top, \sqcup, \sqcap \rangle$ is a function $\nabla : L \times L \to L$ such that: (1) ∇ is an upper bound operator, i.e., for all $x, y \in L, x \sqcup y \sqsubseteq x \lor y$, and (2) for all infinite ascending chains $x_0 \sqsubseteq x_1 \sqsubseteq \ldots$ in L, the chain $y_0 \sqsubseteq y_1 \sqsubseteq \ldots$ eventually stabilizes, where $y_0 \stackrel{\text{def}}{=} x_0$ and $y_i \stackrel{\text{def}}{=} y_{i-1} \lor x_i$ for all i > 0.

Proposition 3 ([16]). Given a complete lattice $\langle L, \sqsubseteq, \bot, \top, \sqcup, \sqcap \rangle$, let $F : L \to L$ be a monotone function and $\nabla : L \times L \to L$ a widening operator for L. Let \overline{X} be an (upward) iteration sequence defined over L, as follows

Let X_f be the least fixpoint $\mathbf{lfp}_{\perp}^{\sqsubseteq}F$ of F, which exists due to Tarski [88]. It follows then that \overline{X} is stationary with the limit X such that $X_f \sqsubseteq X$.

Chapter 3

Data Flow Refinement Type Inference

We motivate our work on casting Liquid type inference as abstract interpretation by providing a more detailed introduction to the actual inference algorithm, its distinguishing features, and challenges of modeling the algorithm as an abstract semantics of programs. Proofs for all the theorems, lemmas, and corollaries stated in this chapter can be found in Appendix A, unless already presented in the chapter.

3.1 Overview

To better understand how the Liquid Type inference works, consider the following OCaml program that is slightly more involved than the example from § 1.

```
1 let dec y = y - 1 in

2 let f x g = if x >= 0 then g x else x in

3 f (f 1 dec) dec
```

The program first defines a function dec that decrements the input value y by one and then calls the higher-order function f twice, passing it an integer value and dec in each call. In turn, f calls the given function g on its parameter x if x is not negative, which is true for both calls to f in the program.

3.1.1 Liquid Type Inference

The Liquid type inference algorithm works as follows. First, the analysis performs a standard Hindley-Milner type inference to infer the basic shape of the refinement type for every subexpression of the program. For instance, the inferred type for the function **f** is int \rightarrow (int \rightarrow int) \rightarrow int. For every function type $\tau_1 \rightarrow \tau_2$, where τ_1 is a base type such as int, the analysis next introduces a fresh dependency variable x which stands for the function parameter, $x : \tau_1 \rightarrow \tau_2$. The scope of x is the result type τ_2 , i.e., refinement predicates inferred for τ_2 can express dependencies on the input τ_1 by referring to x. Further, every base type τ is replaced by a refinement type, { $\nu : \tau \mid \phi(\nu, \vec{x})$ }, with a placeholder refinement predicate ϕ that expresses a relation between the members ν of τ and the other variables \vec{x} in scope of the refinement type. For example, the augmented type for function f is

$$\begin{aligned} x: \{\nu: \mathsf{int} \mid \phi_1(\nu)\} &\to (y: \{\nu: \mathsf{int} \mid \phi_2(\nu, x)\} \to \\ \{\nu: \mathsf{int} \mid \phi_3(\nu, x, y)\}) &\to \{\nu: \mathsf{int} \mid \phi_4(\nu, x)\}. \end{aligned}$$

The algorithm then derives a system of Horn clauses modeling the subtyping constraints imposed on the refinement predicates by the program *data flow*. For example, the body of the function **f** induces the following Horn clauses over the refinement predicates in f's type:

$$x \ge 0 \land \phi_1(x) \Rightarrow \phi_2(x, x)$$
$$x \ge 0 \land \phi_1(x) \land \phi_2(x, x) \land \phi_3(\nu, x, y) \Rightarrow \phi_4(\nu, x)$$
$$x < 0 \land \phi_1(x) \land \nu = x \Rightarrow \phi_4(\nu, x)$$

The first two clauses model the data flow from the parameter \mathbf{x} to the input of the function parameter \mathbf{g} of \mathbf{f} and, in turn, the result value of \mathbf{g} to the result value of \mathbf{f} in the *then* branch of the conditional. The third clause captures the flow of the parameter x to \mathbf{f} 's result value in the *else* branch. Similar entailments are derived from the other subexpressions.

The algorithm finally solves the Horn clauses using monomial predicate abstraction [52] to derive the refinement predicates ϕ_i . That is, the analysis assumes a given set of atomic predicates $Q = \{p_1(\nu, \vec{x}), \dots, p_n(\nu, \vec{x})\}$, which are either provided by the programmer or derived from the program using heuristics, and then infers an assignment for each ϕ_i to a conjunction over Q such that all Horn clauses are valid. This can be done effectively and efficiently using the Houdini algorithm [26, 52]. For example, if we choose $Q = \{0 \leq \nu, \nu < 2, \nu = x, \nu = x - 1, \nu = y - 1, \nu = 1\}$, then the final type inferred for function **f** will be:

$$\begin{aligned} x \colon \{\nu \colon \text{int} \mid 0 \le \nu < 2\} \to &(y \colon \{\nu \colon \text{int} \mid 0 \le \nu < 2 \land \nu = x\} \to \\ \{\nu \colon \text{int} \mid \nu < 2 \land \nu = y - 1\}) \to &\{\nu \colon \text{int} \mid \nu < 2 \land \nu = x - 1\} \end{aligned}$$

The meaning of this type is tied to the current program. For instance, the above inferred type guarantees that the program only calls \mathbf{f} with an integer in the interval [0, 2) as the first argument. In other words, the inference algorithm is able to infer that only the integer values in the interval [0, 2) flow to \mathbf{f} as inputs and

that the output values flowing out of f are integers smaller than 2 and are equal to the decrement of the respective inputs. Note that for a different program, the analysis would potentially infer a different refinement type for f. For instance, if we replaced dec by the identity function, the analysis would infer $\{\nu: \text{int} \mid \nu = 1\}$ for the first input x of f. Further, the type inferred for the input g of f captures how functions passed as g to f are used by f. To see this in more detail, consider the following two types inferred for the two usages of f in the program:

$$\begin{aligned} x \colon \{\nu : \text{int} \mid \nu = 1\} &\to (y \colon \{\nu : \text{int} \mid \nu = x = 1\} \to \\ \{\nu : \text{int} \mid \nu = x = y = 1\}) \to \{\nu : \text{int} \mid \nu = x = 1\} \\ x \colon \{\nu : \text{int} \mid \nu = 0\} \to (y \colon \{\nu : \text{int} \mid \nu = x = 0\} \to \\ \{\nu : \text{int} \mid \nu = x = y = 0\}) \to \{\nu : \text{int} \mid \nu = x = 0\} \end{aligned}$$

Intuitively, the function f takes two different execution paths whose endpoints correspond to the two usages of f. The types inferred for these two usages capture the behavior of f for the inputs generated at these two call sites. The first type above, corresponding to the nested call to f, states that f is here called with integer 1 as the first input and a function accepting and producing integers as the other input. This function, once passed to f at this call site, is subsequently called with integer 1, also producing integer 1 as the result. Finally, the output of this invocation to f is also integer 1. The similar explanation can be given for the type corresponding to the outer call to f. Both of these two types are encompassed by the type inferred for the definition of f. In general, Liquid type inference infers data flow invariants of higher-order programs that capture not only the kind of values to which a program expression evaluates, but also how these values are used by the program from that point onward in the execution paths taken by the values.

3.1.2 Problem

Rather than indirectly describing the inference as a constraint solving problem induced by a type system with a subtyping relation, we model Liquid type inference directly as an abstraction, formalized by Galois connections, of higher-order data flow that induces the subtyping relation. One of the key technical challenges in modeling Liquid types in this way concerns the choice of the concrete semantics. While soundness of the Liquid typing rules can be proved against denotational [92] and operational semantics [78], neither is well suited for calculationally constructing Liquid type inference using abstract interpretation. Intuitively, these semantics do not explicitly capture the data flow properties of programs.

The problem with denotational semantics is that it is inherently compositional; functions are analyzed irrespective of the context in which they appear. Formally, Liquid types and their denotational interpretations do not form a Galois connection. Consider a function let $\mathbf{f} \mathbf{x} = \mathbf{x} + \mathbf{1}$ and the two types $\tau_1 = \{\nu : \operatorname{int} \mid \nu \geq 0\}$ $0\} \rightarrow \{\nu : \operatorname{int} \mid \nu \geq 0\}$ and $\tau_2 = \{\nu : \operatorname{int} \mid \nu < 0\} \rightarrow \{\nu : \operatorname{int} \mid \nu \leq 0\}$. The denotational semantics of \mathbf{f} is a function that maps all input values that are integers to their successors and all other inputs to the error value. The type τ_1 describes all functions that map a positive integer to a positive integer and all other values to arbitrary values. The meaning of τ_2 is similar. The denotation of \mathbf{f} is contained in the concretization of both τ_1 and τ_2 and hence in their intersection. However, the meet, i.e., the greatest lower bound, of τ_1 and τ_2 subject to the subtyping ordering is $\tau = \{\nu : \operatorname{int} \mid true\} \rightarrow \{\nu : \operatorname{int} \mid \nu = 0\}$. Clearly, the denotation of \mathbf{f} is not described by τ , which means that the denotational concretization of Liquid types cannot be the right adjoint of a Galois connection [17]. One way of forming a Galois connection between types and a denotational domain is by using intersection types [15]. However, intersection types are not supported by Liquid type inference and lead to a fundamentally different analysis [28].

Operational semantics, on the other hand, do not make explicit how a function value flows from the point of its creation to its call sites. A formal connection to Liquid type inference would then involve complex reasoning that relates steps in the program's execution trace where function values are created with subsequent steps where they are called.

3.1.3 Approach

We therefore start our technical development with a new concrete *data flow* semantics for untyped higher-order functional programs. This new semantics is inspired by minimal function graphs [46] and the exact semantics for flow analysis [41].

3.1.3.1 Data Flow Semantics

The key idea behind our data flow semantics is as follows. We first formally model unique evaluation points in a program execution. Intuitively, these points are identified by an execution environment and the unique location of the expression currently being evaluated and can be used to describe execution paths that values take in a higher-order program. As opposed to more classical semantics, such as big-step operational semantics, the meaning of programs computed by our data flow semantics is a map from execution points to *data flow* values.

Consider the illustrative program in Example 1 (left). We explain the meaning of the content to the right shortly. Each expression in the program is annotated with a unique location but we show only a portion of these locations.

Example 1. A simple OCaml program (to the left) and its meaning (to the right)
as given by the data flow semantics for higher-order untyped languages.

1	let id $x = x^g$ in	${\rm id} \to [a \mapsto$	$\langle 1,1\rangle,\ d\mapsto \langle 2,2\rangle]$
2	let $u = (id^a \ 1^b)^c$ in $(id^d \ 2^e)^f$	$\mathbf{u}, b, c \to 1,$	$a \to [a \mapsto \langle 1, 1 \rangle]$
5		$e, f \rightarrow 2,$	$d \to [d \mapsto \langle 2, 2 \rangle]$
		$g^a \to 1,$	$g^d \rightarrow 2$

The data flow map for our running example is shown in Example 1 to the right. Here, we identify execution points with a notation based on locations only, to avoid clutter. For instance, for the entries **b**, **c**, **u**, **e**, and **f**, the map shows the values to which the expressions annotated with those locations evaluated in the above program. As an example, the $e, f \rightarrow 2$ mapping indicates that the result of executing expressions associated with locations e and f is integer 2. The more interesting case is with execution points whose location is g. We have two such points, g^a and g^d , that store the results of evaluating the body of **id** stemming from the two invocations of **id** at call site points identified with locations a and d, respectively. In other words, we use a and d designations on g to informally describe different stack information present when the program evaluation reaches the expression with the location g two times in the above program. The valuation for points described so far are constants with the expected value, i.e., the same value one obtains using, say, a big-step operational semantics.

The more interesting cases occur at execution points that store function values. We model function values as *tables* that keep track of the call sites points at which the associated functions are called and the pair of input and output values produced. For instance, the function *id* is assigned a meaning as a table $[a \mapsto$ $\langle 1, 1 \rangle$, $d \mapsto \langle 2, 2 \rangle$] that stores input-output values for both call site points of *id*, *a* and *d*, where an input to *id* has been observed, i.e., where the function was called. Interestingly, the execution points corresponding to the two usages of id are also tables but with only a fraction of the content computed for the definition of id. Intuitively, id takes two separate execution paths in the program, reaching locations a and d. For each execution point on these two paths, the associated tables capture how id has been used from that point onward in the corresponding paths. Tables stored at execution points a and d thus contain information about inputs and outputs of id for the call sites a and d, respectively.

As we show later, the computation in our data flow semantics boils down to propagation of information between data flow values stored at the successive execution points. On a high level, function inputs are propagated from call site points to the function definitions and the outputs are propagated back in the other direction. For instance, once it observes the input 1 generated at the call site a, the table assigned to the point a propagates this input to the table associated with the point id. After this input is evaluated, the produced output is stored in the table mapped to id and then propagated back to a. Note how this change of direction in which the inputs and outputs of functions are propagated intuitively corresponds to the input contravariance and output covariance of subtyping for function types.

3.1.3.2 Liquid Types as Abstract Interpretation

We continue our technical development by creating a series of semantics that abstract, in the sense of Galois connections, the concrete data flow semantics. This series of abstractions is shown in Figure 3.1. We use the concrete data flow semantics to define the properties of higher-order programs that are abstracted by the Liquid types. These properties are captured by the *collecting semantics*. We then abstract the collecting semantics by a *relational data flow semantics*. The basic



Figure 3.1: Series of abstract semantics leading up to Liquid type inference

idea behind the relational semantics is to abstract concrete data flow values by relations since Liquid types intuitively represent relations expressed in a certain decidable logic. In the next abstraction step, we introduce an abstract semantics that does not maintain the precise information on the call sites at which the inputs to functions are being generated. More precisely, the tables modeling functions in this new semantics, called the *collapsed semantics*, do not distinguish inputs by their call sites, as is the case with Liquid types and in contrast to the concrete and relational data flow semantics. The call site information in tables is collapsed and, in effect, functions are modeled essentially as single relations capturing the input-output behavior of the functions in the program. The key idea behind the general refinement type semantics, the next semantics in our abstraction chain, is to abstract relations with refinement types that consist of a basic type and an element of an arbitrary relational abstract domain representing type refinements. That is, the general refinement type inference is parametric in the abstract domain of type refinements. This semantics also ensures the sound termination of the type inference via widening operators [16, 17]. As we show, these operators are defined on the structure of refinement types using the available widening operators of the chosen abstract domain of type refinements. Finally, Liquid type inference is a particular instance of the general data flow refinement type inference.

Consequences. The benefits of modeling Liquid type inference as abstract interpretation are multifold. First, the design space, i.e., the precision loss points, of type inference algorithms developed using our framework are explicitly captured by Galois connections used in the abstraction chain and widening operators used to enforce the finite convergence of the analysis. Second, we obtain a generalization of the Liquid type inference in the form of general refinement type semantics. Any instantiation of this general system yields a data flow refinement type inference that is sound by construction. Lastly, type inference analyses that build on our results can be more easily compared against each other and used in combination with other static analyses (e.g., via reduced products [17]).

3.2 Data Flow Semantics

In this section, we formally introduce our data flow semantics of untyped functional programs.

Language. Our formal presentation in this chapter assumes that our language does not support let expressions. This simplification is made as we are not investigating polymorphism in this chapter. We name the resulting language λ^d . However, we will informally use let constructs in our examples to make them more readable. In general, our examples will use the constructs of more familiar programming languages, such as OCaml, that can be easily compiled to λ^d . For instance, the let constructs can be expanded into function applications of λ^d as expected [74]. The same holds for the conditional expressions.

3.2.1 Examples

Before we formally introduce the data flow semantics, we first provide more informal intuition through several additional examples.

3.2.1.1 Data Flow Computation

First, we explain how our semantics computes data flow maps. In Figure 3.2, we show the data flow map computed in each step by our data flow semantics for the program in Example 1. We only show the map entries for execution points storing tables, for simplicity. Initially, no useful information for execution points id, a,

Figure 3.2: Subsequent maps computed by data flow semantics for program from Example 1. Numbers above maps are used to identify the maps.

and d is known, denoted with \perp in the map identified with (1). After evaluating line 1, id, being a function definition, is assigned the empty table [] (2). This

table indicates that the meaning of id is a function, with currently no information on how this function is used in the program. In the next step (3), execution point a is reached and is also assigned [] based on the value of id. However, in the next step (4) our semantics observes that id is called with 1 at execution point a; the table associated with a hence contains this information. This information is then propagated back to the table associated with id in the next step (5). As we shall see later, this exchange of information is formally defined as propagation of data between the table associated with the point id and the table stored at a. Now that id has seen input 1, our semantics evaluates the body of id with 1, producing 1 as the result (6). In the next step (7), this output is propagated back in the other direction to the table stored at a. This change of direction in which inputs and outputs are propagated between tables intuitively corresponds to the contravariance of subtyping for function types [74]. After this step, the execution point d becomes reachable and is assigned the empty table [] (8). The data flow computation then continues in the similar fashion as described earlier. Note how a data flow value stored at some execution point, such as id, can be updated multiple times in different execution steps of our semantics.

3.2.1.2 Recursion

We now exemplify how our semantics models recursion through a program that can be found below in Example 2.

Example 2. The OCaml program below first defines a function f that returns 0 when given 1 as input, and it otherwise just recurses with a decrement of the input. This function is then called two times with inputs 2 and 0, respectively. The latter call never terminates. To avoid clutter, we only annotate program expressions of

particular interest with their locations.

let rec f n = if n=1 then 0 else
$$f^a$$
 (n-1)
in $(f^b \ 2^c) + (f^d \ 0^e)$

A portion of the map computed by our data flow semantics for the above program is shown below. We again identify execution points with their locations.

$$\begin{array}{ll} b \to & [b \mapsto \langle 2, 0 \rangle] & d \to [d \mapsto \langle 0, \bot \rangle] \\ f \to & [b \mapsto \langle 2, 0 \rangle, \ a^2 \mapsto \langle 1, 0 \rangle, \\ & d \mapsto \langle 0, \bot \rangle, \ a^0 \mapsto \langle -1, \bot \rangle, \ a^{-1} \mapsto \langle -2, \bot \rangle, \ \ldots] \end{array}$$

Let us first discuss the execution point identified by b, which corresponds to the call to f with integer 2. The meaning assigned to b is a table indicating that the corresponding function is invoked at the call site b with integer 2, returning 0 as the output. The computation of this output clearly requires a recursive call to f with 1 as the input. The existence and effect of this call can be observed by the entry $a^2 \mapsto \langle 1, 0 \rangle$ in the table associated with f. Intuitively, this entry states that f was called at the point identified by a, which is inside the body of f indicating recursion, while analyzing f with the previous input 2. Here, we use the integer value 2 designation on a to informally describe the stack information present when the execution reaches a for this particular case. We note that our choices of describing stack information in the examples are made for presentation purposes only; we formally explain how our data flow semantics encodes stack information in § 3.2.2.

Consider now the execution point identified by d, which corresponds to the call to f with integer 0. The meaning assigned to d is a table indicating that the corresponding function is invoked at the call site d with integer 0 but with no

computed output \perp . Clearly, there is no useful output computed since **f** does not terminate with 0 as the input. The recursive calls to **f** observed after calling **f** with 0 and their effects can be again observed in the table computed for **f**. For instance, the table entry $a^0 \mapsto \langle -1, \perp \rangle$ states that the recursive call to **f** at the call site point *a* has seen -1 as the input, but no output has been computed. The designation 0 on *a* again informally describes the stack information stating that this call to **f** is induced by the preceding call to **f** with 0 as the input. Similarly, the entry $a^{-1} \mapsto \langle -2, \perp \rangle$ indicates that **f** was recursively called at *a* with -2 as input and with no useful output observed. This recursive call was induced by the previous call to **f** with -1 as the input.

Observe how for every recursive call to \mathbf{f} we remember the stack information describing the preceding call to \mathbf{f} that led to this recursive call. That is, the table computed for \mathbf{f} has entries of the form $a^i \mapsto \langle i - 1, \perp \rangle$ for every integer $i \leq 0$. Each call site point a^i explicitly stores the information about the previous call to \mathbf{f} , here designated with the input integer i, that led to the current execution of the body of \mathbf{f} . This way, we uniquely distinguish between different recursive calls to \mathbf{f} in the program execution. We note again that the examples use a simplified notion of stack information for simplicity. As we shall see later in this section, each execution point a^i is modeled as pair of a particular execution environment and the location a. Since \mathbf{f} is a recursive function, this environment will contain a certain binding for \mathbf{f} itself. That binding will explicitly contain the information about the call site point, say x, of \mathbf{f} that led to the current recursive execution of \mathbf{f} . The call site x will in turn remember the information about the preceding call to \mathbf{f} that ultimately led to x, and so on. Modeling stack information in this way allows us to uniquely identify execution points even in the presence of recursive functions. For instance, suppose there are two calls with the same two argument values passed to a recursive function f in a program. These two top-level calls will induce recursive calls to f with the same argument values. Our semantics, however, is able to distinguish between these seemingly same recursive calls since they were induced by the two top-level calls to f happening at the two different call sites.

3.2.1.3 Higher-order Functions

We also illustrate how our semantics models higher-order functions of λ^d . Intuitively, a higher-order function will be modeled as a table some of whose inputs and outputs are also tables. We will work with the program shown in Example 3.

Example 3. The OCaml program below first defines a constant function g1 always returning 0, identity function g2, and a higher-order function h that accepts a function and simply calls it with integer 1. Then, h is called with g1 and g2 after which the respective results of these two calls are divided. This program is safe, returning 0 as the result. For simplicity, we only annotate few program expressions of interest with their locations.

let $h g = g 1^{a}$ in let g1 x = 0 in let g2 x = x in $(h^{b} g1^{c})/(h^{d} g2^{e})$ The corresponding data flow map is shown below.

$$g1, c \to [a^b \mapsto \langle 1, 0 \rangle] \quad g2, e \to [a^d \mapsto \langle 1, 1 \rangle]$$
$$b \to [b \mapsto \langle [a^b \mapsto \langle 1, 0 \rangle], 0 \rangle] \quad d \to [d \mapsto \langle [a^d \mapsto \langle 1, 1 \rangle], 1 \rangle]$$
$$h \to [b \mapsto \langle [a^b \mapsto \langle 1, 0 \rangle], 0 \rangle, d \mapsto \langle [a^d \mapsto \langle 1, 1 \rangle], 1 \rangle]$$

Let us focus on the valuation for the function h. The meaning of h here is a table that has (non-bottom) entries for call site points identified with locations b and d, indicating that these are the two call site points from which the inputs to h are emitted. For the call site b, where h is applied to g1, the observed input-output pair consists of a table $T = [a^b \mapsto \langle 1, 0 \rangle]$ and integer 0. Indeed, h produces 0 with g1 as the input. The table T encodes the fact that h applies the input argument g, which is in this case g1, at the call site identified by a^b , i.e., inside the body of h at location a while analyzing h with the input seen at b. Note how we here again use the call site b to informally describe the stack information. Finally, the contents of T indicate that the input g of h was here called with integer 1 as the input, producing 0 as the output. Similar analysis applies for the second entry dof table for h. The explanation for other nodes in the map is similar.

We now proceed to the formal presentation of our data flow semantics for untyped higher-order programs.

3.2.2 Semantic Domains

We start our formal exposition of concrete data flow semantics by introducing semantic domains used for giving meaning to expressions of λ^d .

Nodes and environments

Every intermediate point of a program's execution is uniquely identified by an execution node (or simply node), $n \in \mathcal{N}$. We distinguish expression nodes \mathcal{N}_e and variable nodes \mathcal{N}_x . An expression node $\langle E, \ell \rangle$, denoted $E \diamond \ell$, captures the execution point where a subexpression e^{ℓ} is evaluated in the environment E.

nodes	$n \in \mathcal{N} \stackrel{\scriptscriptstyle\mathrm{def}}{=} \mathcal{N}_e \cup \mathcal{N}_x$
expression nodes	$\mathcal{N}_e \stackrel{\text{\tiny def}}{=} \mathcal{E} imes Loc$
variable nodes	$\mathcal{N}_x \stackrel{\text{\tiny def}}{=} \mathcal{E} imes \mathcal{N}_e imes Vars$
environments	$E \in \mathcal{E} \stackrel{\text{\tiny def}}{=} Vars \rightharpoonup_{fin} \mathcal{N}_x$
(data flow) values	$v \in \mathcal{V} \ ::= \ \bot \mid \omega \mid c \mid T$
tables	$T \in \mathcal{T} \stackrel{\text{\tiny def}}{=} \mathcal{N}_e \to \mathcal{V} \times \mathcal{V}$
execution maps	$M \in \mathcal{M} \stackrel{\text{\tiny def}}{=} \mathcal{N} \to \mathcal{V}$

An environment E is a (finite) partial map binding variables to variable nodes. A variable node $\langle E, n_{cs}, x \rangle$, denoted $E \diamond n_{cs} \diamond x$, models the execution point where an argument value is bound to a formal parameter of a function. Here, E is the environment of the expression node where the function was defined. The expression node n_{cs} is referred to as the *call site* of the variable node. Intuitively, this is the node of the subexpression that emits the argument value of the function application associated with the binding. That is, if the call came from the expression $(e_1^{\ell_1} e_2^{\ell_2})$, then the location of n_{cs} is ℓ_2 and the environment of n_{cs} is the environment used when evaluating the function application. The call site is included in the variable node to uniquely identify the binding. This way, we implicitly model the call stack of a program in our semantics. Finally, x is the formal parameter that is being bound to the variable node. We explain the role of expression and variable nodes in more detail later. For any node n, we denote by env(n) the environment component of n and by loc(n) the location component.

A pair $\langle e, E \rangle$ is called well-formed if E(x) is defined for every variable x that occurs free in e.

Values and execution maps

There are four kinds of (data flow) values $v \in \mathcal{V}$. First, every constant c is also a value. The value \perp stands for nontermination or unreachability of a node, and the value ω models execution errors. Functions are represented by *tables* T. A table maintains an input/output value pair for each call site. We denote by T_{\perp} the empty table that maps every call site to the pair $\langle \perp, \perp \rangle$. We say that a table T has been called at a call site n_{cs} when the associated input value is non-bottom $n_{cs} \in T \Leftrightarrow \pi_1(T(n_{cs})) \neq \perp$. We write $[n_{cs} \mapsto \langle v_1, v_2 \rangle]$ as shorthand for the table $T_{\perp}[n_{cs} \mapsto \langle v_1, v_2 \rangle]$.

The data flow semantics computes execution maps $M \in \mathcal{M}$, which map nodes to values. We denote by M_{\perp} (M_{ω}) the execution maps that assign \perp (ω) to every node. For a set of nodes N, we write $M_1 =_{\backslash N} M_2$ to express that the two maps M_1 and M_2 are pointwise equal on all nodes in $\mathcal{N} \setminus N$.

As a precursor to defining the data flow semantics, we define a *computational* ordering \sqsubseteq on values. In this ordering, \perp is the smallest element, ω is the largest element, and tables are ordered recursively on the pairs of values for each call site:

$$v_1 \sqsubseteq v_2 \iff v_1 = \bot \lor v_2 = \omega \lor (v_1, v_2 \in Cons \land v_1 = v_2) \lor (v_1, v_2 \in \mathcal{T} \land \forall n_{cs}. v_1(n_{cs}) \doteq v_2(n_{cs}))$$

Defining the error value as the largest element of the ordering is nonessential but simplifies the presentation. In particular, this definition ensures that the least upper bounds (lub) of arbitrary sets of values exist, which we denote by the join operator \sqcup .

$$\perp \sqcup v \stackrel{\text{def}}{=} v \qquad v \sqcup \perp \stackrel{\text{def}}{=} v \qquad c \sqcup c \stackrel{\text{def}}{=} c$$
$$T_1 \sqcup T_2 \stackrel{\text{def}}{=} \lambda n_{cs}. T_1(n_{cs}) \mathrel{\dot{\sqcup}} T_2(n_{cs}) \qquad v_1 \sqcup v_2 \stackrel{\text{def}}{=} \omega \quad \text{(otherwise)}$$

The meet of values is defined in the expected way. These two operators together with the above ordering relation on values form a complete lattice.

Lemma 1. Let
$$V \in \wp(\mathcal{V})$$
. Then, $\sqcup V = lub(V)$ and $\sqcap V = glb(V)$ subject to \sqsubseteq .

As expected, execution maps also form a complete lattice with the ordering relation $\dot{\sqsubseteq}$, join $\dot{\sqcup}$, and meet $\dot{\sqcap}$ operators that are simply pointwise lifts of the corresponding ordering relation and operators on values.

3.2.3 Concrete Transformer

The data flow semantics of an expression e is defined as the least fixpoint of a concrete transformer, **step**, on execution maps. The idea is that we start with the map M_{\perp} and then use **step** to consecutively update the map with new values as more and more execution nodes are reached. The signature of this transformer is as follows:

$$\mathsf{step}: \mathbb{N} \to \lambda^d \to \mathcal{E} \to \mathcal{M} \to \mathcal{M}$$

Essentially, step takes $k \in \mathbb{N}$, an expression e, and an environment E and returns a transformer $\operatorname{step}_k[\![e]\!](E) : \mathcal{M} \to \mathcal{M}$ on execution maps. The definition of the transformer is given in Fig. 3.3 using induction over both k and the structure of e. Note that in the definition we implicitly assume that $\langle e, E \rangle$ is well-formed. Let $\operatorname{step}_0[\![e]\!](E)(M) \stackrel{\text{def}}{=} M$ $\mathsf{step}_{k+1}\llbracket c^{\ell} \rrbracket(E)(M) \stackrel{\text{def}}{=} M[E \diamond \ell \to M(E \diamond \ell) \sqcup c]$ $\mathsf{step}_{k+1}\llbracket x^{\ell} \rrbracket(E)(M) \stackrel{\text{def}}{=}$ let $\langle v'_x, v' \rangle = \operatorname{prop}(M(E(x)), M(E \diamond \ell))$ in $M[E(x) \mapsto v'_x, E \diamond \ell \mapsto v']$ $step_{k+1} [\![(e_1^{\ell_1} e_2^{\ell_2})^{\ell}]\!](E)(M) \stackrel{\text{def}}{=}$ let $M_1 = \text{step}_k[e_1](E)(M); v_1 =_{M_1} M_1(E \diamond \ell_1)$ in if $v_1 \notin \mathcal{T}$ then M_{ω} else let $M_2 = \text{step}_k [\![e_2]\!](E)(M_1); v_2 =_{M_2} M_2(E \diamond \ell_2)$ in let $\langle v_1', T' \rangle = \operatorname{prop}(v_1, [E \diamond \ell_1 \mapsto \langle v_2, M_2(E \diamond \ell) \rangle])$ $\langle v_2', v' \rangle = T'(E \diamond \ell_1)$ in $M_2[E \diamond \ell_1 \mapsto v'_1, E \diamond \ell_2 \mapsto v'_2, E \diamond \ell \mapsto v']$ $\operatorname{step}_{k+1} \left[(\mu f \cdot \lambda x \cdot e_1^{\ell_1})^{\ell} \right] (E)(M) \stackrel{\text{def}}{=}$ let $T =_M T_{\perp} \sqcup M(E \diamond \ell)$ in let $\overline{M} = \Lambda n_{cs} \in T$. if $\pi_1(T(n_{cs})) = \omega$ then M_{ω} else let $n_x = E \diamond n_{cs} \diamond x$; $n_f = E \diamond n_{cs} \diamond f$; $E_1 = E[x \mapsto n_x, f \mapsto n_f]$ $\langle T'_{x}, T' \rangle = \operatorname{prop}([n_{cs} \mapsto \langle M(n_{x}), M(E_{1} \diamond \ell_{1}) \rangle], T)$ $\langle T'', T'_f \rangle = \operatorname{prop}(T, M(n_f)); \langle v'_r, v'_1 \rangle = T'_r(n_{cs})$ $M_1 = M[E \diamond \ell \mapsto T' \sqcup T'', n_f \mapsto T'_f, n_x \mapsto v'_x, E_1 \diamond \ell_1 \mapsto v'_1]$ in step_k $[e_1](E_1)(M_1)$ in $M[E \diamond \ell \mapsto T] \sqcup \dot{|}_{n \in T} \overline{M}(n_{cs})$ $\operatorname{step}_{k+1} \llbracket (e_0^{\ell_0} ? e_1^{\ell_1} : e_2^{\ell_2})^{\ell} \rrbracket (E)(M) \stackrel{\text{def}}{=}$ let $M_0 = \text{step}_k \llbracket e_0 \rrbracket (E)(M); v_0 =_{M_0} M_0(E \diamond \ell_0)$ in if $v_0 \notin Bool$ then M_{ω} else let $b = if v_0 = true$ then 1 else 2 in let $M_b = \operatorname{step}_k \llbracket e_b \rrbracket(E)(M_0)$; $v_b =_{M_b} M_b(E \diamond \ell_b)$ in let $\langle v'_b, v' \rangle = \operatorname{prop}(v_b, M_b(E \diamond \ell))$ in $M_b[E \diamond \ell \mapsto v', E \diamond \ell_b \mapsto v'_b]$

Figure 3.3: Concrete transformer for data flow semantics

$$\begin{aligned} \operatorname{prop}(T_1, T_2) \stackrel{\text{def}}{=} \\ & \operatorname{let} T' = \Lambda n_{\operatorname{cs}}. \\ & \operatorname{if} n_{\operatorname{cs}} \notin T_2 \operatorname{then} \langle T_1(n_{\operatorname{cs}}), T_2(n_{\operatorname{cs}}) \rangle \operatorname{else} \\ & \operatorname{let} \langle v_{1i}, v_{1o} \rangle = T_1(n_{\operatorname{cs}}); \langle v_{2i}, v_{2o} \rangle = T_2(n_{\operatorname{cs}}) \\ & \langle v'_{2i}, v'_{1i} \rangle = \operatorname{prop}(v_{2i}, v_{1i}); \langle v'_{1o}, v'_{2o} \rangle = \operatorname{prop}(v_{1o}, v_{2o}) \\ & \operatorname{in} (\langle v'_{1i}, v'_{1o} \rangle, \langle v'_{2i}, v'_{2o} \rangle) \\ & \operatorname{in} \langle \Lambda n_{\operatorname{cs}}. \pi_1(T'(n_{\operatorname{cs}})), \Lambda n_{\operatorname{cs}}. \pi_2(T'(n_{\operatorname{cs}})) \rangle \\ & \operatorname{prop}(T, \bot) \stackrel{\text{def}}{=} \langle T, T_{\bot} \rangle \qquad \operatorname{prop}(T, \omega) \stackrel{\text{def}}{=} \langle \omega, \omega \rangle \\ & \operatorname{prop}(v_1, v_2) \stackrel{\text{def}}{=} \langle v_1, v_1 \sqcup v_2 \rangle \qquad (\operatorname{otherwise}) \end{aligned}$$

Figure 3.4: Value propagation in the concrete data flow semantics

us ignore the parameter k for a moment and focus on the equations that case split on the structure of e. We discuss the cases for e one at a time.

Constant $e = c^{\ell}$. Here, we simply set the current node to the join of the current value stored for $E \diamond \ell$ and the value c. The use of join is non-essential, but it simplifies our presentation; in practice, the result of the join is always c.

Variable $e = x^{\ell}$. This case is already more interesting. It implements the data flow propagation between the variable node E(x) where x is bound and the expression node $E \diamond \ell$ where x is used. This is realized using the helper function prop defined in Fig. 3.4. Let $v_x = M(E(x))$ and $v = M(E \diamond \ell)$ be the current values stored at the two nodes in M. The function prop takes these values as input and propagates information between them, updating them to the new values v'_x and v' which are then stored back into M. The propagation works as follows. If v_x is a constant or the error value and v is still \perp , then we simply propagate v_x forward, replacing v and leaving v_x unchanged. The interesting cases are when we propagate information between tables. The idea is that inputs in a table v flow backward to v_x whereas outputs for these inputs flow forward from v_x to v. Thus, if v_x is a table but v is still \perp , we initialize v to the empty table T_{\perp} and leave v_x unchanged (because we have not yet accumulated any inputs in v). If both v_x and v are tables, we propagate inputs and outputs as described above by calling prop recursively for every call site n_{cs} that is called in v. Note how the recursive call for the propagation of the inputs prop (v_{2i}, v_{1i}) inverts the direction of the propagation. As we shall see, this will give rise to contravariant subtyping of function types.

Function application $e = (e_1^{\ell_1} e_2^{\ell_2})^{\ell}$. We first evaluate e_1 to obtain the updated map M_1 and extract the new value v_1 stored at the corresponding expression node. Here, we use a special meta-level let binding let $x =_M v$. That is, let $x =_M v$ in tshortcircuits to M if $v = \bot$, M_{ω} if $v = \omega$, and otherwise yields t where x is bound to v. Hence, if the $E \diamond \ell_1$ node has not yet been reached $(v_1 = \bot)$, we return M_1 . If v_1 is not a table, then we must be attempting an unsafe call, in which case we return the error map M_{ω} . If v_1 is a table, we continue evaluation of e_2 obtaining the new map M_2 and value v_2 at the associated expression node $E \diamond \ell_2$. If v_2 is neither \bot nor ω , then we need to propagate the information between this call site and v_1 . To this end, we first extract the return value $v = M(E \diamond \ell)$ for the node of ecomputed thus far and create a singleton table storing the input-output $\langle v_2, v \rangle$ for this call site $E \diamond \ell_1$. We propagate between v_1 and this table to obtain v'_1 and a table T'. Note that this propagation essentially boils to down to (1) the propagation between v_2 and the input of v_1 at $E \diamond \ell_1$ and (2) the propagation between the output of v_1 at $E \diamond \ell_1$ and the return value v. The updated table T' hence contains the updated input and output values v'_2 and v'. All of these values are stored back into M_2 . Intuitively, v'_1 contains the information that the corresponding function received an input coming from the call site $E \diamond \ell_1$. This information will ultimately be propagated back to the function definition that will in turn evaluate this input, as follows.

Recursive function $e = (\mu f.\lambda x. e_1^{e_1})^{\ell}$. We first extract the table T computed for the function thus far. Then, for every call site n_{cs} for which an input has already been back-propagated to T (i.e. $\pi_1(T(n_{cs})) \neq \bot$), we do the following. First, we create a variable node n_x that will store the input value that was intuitively emitted at the call site n_{cs} . The contents of this value can potentially be updated later on during the computation if the actual input is a table that, for instance, has yet to be called. We hence define n_x using n_{cs} as a unique variable node $E \diamond n_{cs} \diamond x$, thus ensuring the execution of the lambda body for other inputs does not interfere with the execution for this input. The similar reasoning underlies the creation of n_f . The unique nature of n_f allows our semantics to precisely track recursive calls to f that occurred while evaluating the input to f coming from n_{cs} . Also, note that by remembering the call site node n_{cs} in the environment nodes n_x and n_f we effectively model the program stack. That is, each call site node n_{cs} contains an environment that in turn contains environment nodes, which again keep track of the call site information, and so on.

We then propagate information between the values stored at the node for the body e_1 , the nodes n_x and n_f for the bound variables, and the table T. That is, we propagate information from (1) the input of T at n_{cs} and the corresponding environment node n_x , and (2) the value assigned to the function body under the update environment E_1 and the output of T at n_{cs} . Further, the value associated with n_f is a table that will intuitively contain information about the recursive calls to the function observed during the analysis of the body for the current input. The propagation between T and the table stored at n_f allows T to pick up such observed inputs and eventually propagate the corresponding outputs back to $M(n_f)$.

The new values obtained by propagation are stored back into the map to obtain a new map M_1 and then we evaluate the body e_1 . The updated maps obtained for all call sites are then merged into a single map via a join. Note that this join does not lose precision by introducing ω . This is because the extended environment E_1 is unique for each call site, which means that the evaluation of the body e_1 for different call sites will update different expression nodes. Moreover, a concurrent update to the same variable node n in two maps $\overline{M}(n_{cs})$ and $\overline{M}(n'_{cs})$ can only occur if M(n) stores a table and the updates are the result of back-propagation of inputs and outputs. In this case, these inputs and outputs must agree if they come from the same call site and the updated tables in $\overline{M}(n_{cs})(n)$ and $\overline{M}(n'_{cs})(n)$ can be safely merged via join. This reflects the deterministic nature of the concrete semantics.

Conditionals. The discussion for conditional expressions is similar to the previous expression cases.

Lemma 2. The function prop is monotone and increasing.

Lemma 3. For every $k \in \mathbb{N}$, $e \in \lambda^d$ and $E \in \mathcal{E}$ such that $\langle e, E \rangle$ is well-formed, then $\operatorname{step}_k[\![e]\!](E)$ is monotone and increasing. Step-indexing. Let us now return to the discussion of the parameter k of step. Intuitively, we would like to just do away with k and define the semantics of a program e as the least fixpoint of step over the complete lattice of execution maps. However, if e does not terminate, then the tables stored in the fixpoint iterates may grow arbitrarily deep. Since tables are defined inductively, they can only have finite nesting depth, which means that the fixpoint would approximate these unbounded tables with ω . Possible solutions include to define tables coinductively or to use a different representation of tables based on Scott domains [80]. However, we here use a simpler solution inspired by step-indexing [2]. We introduce the step-index parameter k, informally called *fuel*, that bounds the number of function applications executed by $\text{step}_k[\![e]\!](E)(M)$ and, in turn, the nesting depth of the tables in the resulting execution map. The fixpoint then simply computes a separate map for every *fuel* k.

We define the semantics $\mathbf{S}[\![e]\!]$ of a program e as the least fixpoint of step over the complete lattice of step-indexed execution maps $\mathbb{N} \to \mathcal{M}$, which is ordered by the pointwise lifting of the computational ordering on values:

$$\mathbf{S}\llbracket e \rrbracket \stackrel{\text{def}}{=} \mathbf{lfp}_{\Lambda k. M_{\perp}}^{\sqsubseteq} \Lambda \mathbf{M}. \Lambda k. \operatorname{step}_{k}\llbracket e \rrbracket(\epsilon)(\mathbf{M}(k))$$

Here, ϵ denotes the empty environment. Lemma 3 guarantees that $\mathbf{S}[\![e]\!]$ is well-defined.

In addition to the example from the beginning of this section, we also show the final execution map for the example from \S 3.1 in Appendix A.

3.2.4 Properties and Collecting Semantics

As we shall see, Liquid type systems abstract programs by properties $P \in \mathcal{P}$, which are step-indexed sets of execution maps: $\mathcal{P} \stackrel{\text{def}}{=} \mathbb{N} \to \wp(\mathcal{M})$. We order properties by the pointwise lifting of subset inclusion. Properties form the concrete domain of our abstract interpretation. We therefore raise the data flow semantics **S** to a *collecting semantics* $\mathbf{C} : \lambda^d \to \mathcal{P}$ that maps programs to properties:

$$\mathbf{C}\llbracket e \rrbracket \stackrel{\text{def}}{=} \Lambda k. \{ \mathbf{S}\llbracket e \rrbracket(k) \}$$

An example of a property is *safety*: let $\mathcal{M}_{\mathsf{safe}}$ be the set of all execution maps that map no node to the error value ω and define $P_{\mathsf{safe}} = \Lambda k.\mathcal{M}_{\mathsf{safe}}$. Then a program eis called *safe* if $\mathbf{C}[\![e]\!] \subseteq P_{\mathsf{safe}}$.

3.3 Relational Data Flow Semantics

We now present the first abstraction of the concrete data flow semantics in the sequence of abstract semantics leading up to Liquid types. As we shall see later, a critical abstraction step performed by Liquid types is to conflate the information in tables that are propagated back from different call sites of that function. That is, a pair of input/output values that has been collected from one call site will also be considered as a possible input/output pair at other call sites to which that function flows. However, in order for Liquid type inference not to lose too much precision, output values observed for a particular call site explicitly remember the corresponding input value. We refer to this semantics as the *relational (data flow) semantics*.

3.3.1 Example

We first provide an intuition behind relational semantics through our running program from Example 1. Relational semantics abstracts the value of every node n in an execution map by a relation over the values observed at n and the values observed for the scope of n that is defined by env(n).

$$\begin{split} & \text{id} \to [a \to \langle \{(\nu : 1)\}, \{(a : 1, \nu : 1)\} \rangle, d \to \langle \{(\nu : 2)\}, \{(d : 2, \nu : 2)\} \rangle] \\ & \text{u}, b, c \to \{(\text{id} : \mathsf{F}, \nu : 1)\} \qquad e, f \to \{(\text{id} : \mathsf{F}, \mathbf{u} : 1, \nu : 2)\} \\ & a \to [a \mapsto \langle \{(\text{id} : \mathsf{F}, \nu : 1)\}, \{(\text{id} : \mathsf{F}, a : 1, \nu : 1)\} \rangle] \\ & d \to [d \mapsto \langle \{(\text{id} : \mathsf{F}, \mathbf{u} : 1, \nu : 2)\}, \{(\text{id} : \mathsf{F}, \mathbf{u} : 1, d : 2, \nu : 2)\} \rangle] \\ & g^a \to \{(x^a : 1, \nu : 1)\} \qquad g^d \to \{(x^d : 2, \nu : 2)\} \end{split}$$

In the above, we present the final map computed by our relational semantics for our running example. Let us explain the map by first focusing on nodes u, b, and c. The concrete semantics assigns integer 1 to these nodes, as explained in § 3.2, while relational semantics assigns the relational value $\{(id: F, \nu: 1)\}$. This value is a relation consisting of a single tuple $(id: F, \nu: 1)$. This tuple states that the value possibly observed for the nodes u, b, and c is integer 1, via mapping $\nu: 1$. Observe that, similar to Liquid types, we use the symbolic variable ν to refer to the values of the current node/expression. On the other hand, the mapping id: Fin the tuple indicates that the node associated with id is in the current scope and it contains some function/table value. Note how this mapping does not state what is that table exactly, yet only the fact that id is a function. This particular design choice in modeling relations is motivated by the fact that Liquid types do not track precise dependencies on function values. The relation $\{(id: F, u: 1, \nu: 2)\}$ computed for nodes e and f states that these nodes possibly evaluate to integer 2, where the values in the scope corresponding to id and u are a function and integer 1, respectively. Note how the relations described so far contain only a single tuple. As we show later, refinements of basic types will concretize to relations having multiple tuples of the form explained above.

Let us now turn to explaining relational tables by considering the entry for id node in the above map. Let us discuss the table entry $a \mapsto \langle \{(\nu : 1)\}, \{(a : 1, \nu : 1)\}\rangle$. As expected, this entry indicates that id is called at the node identified with a. Here, the observed input is a relational value $\{(\nu : 1)\}$ indicating that the actual input passed to id at b is integer 1. Note how there are no other mappings in the tuple of the input relation except $\nu : 1$. This is because the initial environment/scope, used to analyze id, is empty. The output relation $\{(a: 1, \nu: 1)\}$ is more interesting. This value indicates via the pair $\nu : 1$ that the actual output of id observed at the call site a is 1. Additionally, the output value also explicitly remembers the input value. This way, relational semantics tracks dependencies between inputs and outputs, which becomes important in the abstraction steps we introduce later on.

The maps computed by the relational data flow semantics for programs of Example 3 and Example 2 are obtained similarly. For instance, here is the relational map for our program from Example 3 exhibiting higher-order features of λ^d .

$$\begin{split} g1, b &\to [a^b \mapsto \langle \{(\nu \colon 1)\}, \{(a^b \colon 1, \nu \colon 0)\} \rangle] \quad g2, d \to [a^d \mapsto \langle \{(\nu \colon 1)\}, \{(a^d \colon 1, \nu \colon 1)\} \rangle] \\ b &\to [b \mapsto \langle [a^b \mapsto \langle \{(\nu \colon 1)\}, \{(a^b \colon 1, \nu \colon 0)\} \rangle], \{(b \colon \mathsf{F}, \nu \colon 0)\} \rangle] \\ d &\to [d \mapsto \langle [a^d \mapsto \langle \{(\nu \colon 1)\}, \{(a^d \colon 1, \nu \colon 1)\} \rangle], \{(d \colon \mathsf{F}, \nu \colon 1)\} \rangle] \\ h &\to [b \mapsto \langle [a^b \mapsto \langle \{(\nu \colon 1)\}, \{(a^b \colon 1, \nu \colon 0)\} \rangle], \{(b \colon \mathsf{F}, \nu \colon 0)\} \rangle, \\ d \mapsto \langle [a^d \mapsto \langle \{(\nu \colon 1)\}, \{(a^d \colon 1, \nu \colon 1)\} \rangle], \{(d \colon \mathsf{F}, \nu \colon 1)\} \rangle] \end{split}$$

Note again the use call site points in tables to express input-output dependencies. To avoid clutter, we also do not show all entries for scope nodes within rows of relations. For instance, input and output relations in the relational tables of b and d should also contain entries $h: \mathsf{F}, g1: \mathsf{F}$, and $g2: \mathsf{F}$.

3.3.2 Abstract Domains

We start our formal exposition of relational semantics by introducing the abstract domains. The definitions of nodes and environments are the same as in the concrete semantics. The relational abstractions of values, constants, and tables are defined with respect to *scopes*, which are finite sets of nodes $N \subseteq \mathcal{N}$. Given a node n, we denote its scope by $N_n \stackrel{\text{def}}{=} \operatorname{rng}(env(n))$. The new domains are defined as follows:

$$\begin{split} r_{N} \in \mathcal{V}_{N}^{\mathsf{r}} & ::= \ \bot^{\mathsf{r}} \mid \top^{\mathsf{r}} \mid R_{N}^{\mathsf{r}} \mid T_{N}^{\mathsf{r}} & \text{relational values} \\ R_{N}^{\mathsf{r}} \in \mathcal{R}_{N}^{\mathsf{r}} \stackrel{\text{def}}{=} \ \wp(\mathcal{D}_{N}^{\mathsf{r}}) & \text{relations} \\ D_{N}^{\mathsf{r}} \in \mathcal{D}_{N}^{\mathsf{r}} \stackrel{\text{def}}{=} \ N \cup \{\nu\} \to Cons \cup \{\mathsf{F}\} & \text{dependency vectors} \\ T_{N}^{\mathsf{r}} \in \mathcal{T}_{N}^{\mathsf{r}} \stackrel{\text{def}}{=} \ \Pi n_{\mathsf{cs}} \in \mathcal{N}_{e}. \ \mathcal{V}_{N \setminus \{n_{\mathsf{cs}}\}}^{\mathsf{r}} \times \mathcal{V}_{N \cup \{n_{\mathsf{cs}}\}}^{\mathsf{r}} & \text{relational tables} \\ M^{\mathsf{r}} \in \mathcal{M}^{\mathsf{r}} \stackrel{\text{def}}{=} \ \Pi n \in \mathcal{N}. \ \mathcal{V}_{N_{n}}^{\mathsf{r}} & \text{relational maps} \end{split}$$

Relational values r_N model how concrete values, stored at some node, depend on the concrete values of nodes in the current scope N. The relational value \perp^r again models nontermination or unreachability and imposes no constraints on the values in its scope. The relational value \top^r models every possible concrete value, including ω . Concrete constant values are abstracted by relations R_N^r , which are sets of *dependency vectors* D_N^r . A dependency vector tracks the dependency between a constant value stored at the current node, represented by the special symbol ν , and the values stored at the other nodes in the current scope N. Here, we assume that ν is not contained in \mathcal{N} .

We only track dependencies between constant values precisely: if a node in the scope stores a table, we abstract it by the symbol F which stands for an arbitrary concrete table. We assume F to be different from all other constants. We also require that for all $D_N^r \in r_N$, $D^r(\nu) \in Cons$. We will see later that the inferred refinement relations $\phi(\nu, \vec{x})$ of Liquid types are abstractions of dependency relations. That is, every satisfying assignment of a $\phi(\nu, \vec{x})$ concretizes to a set of dependency vectors.

Relational tables T^r are defined analogously to concrete tables except that they now map call sites n_{cs} to pairs of relational values $\langle r_i, r_o \rangle$ rather than concrete values. The call site node is not allowed to be in the scope of the input relational r_i but is explicitly added to the scope of r_o . The idea is that r_o tracks the dependency between the output value and the input values described by r_i via n_{cs} . We denote by $T_{\perp r}^r$ the empty relational table that maps every call site to the pair $\langle \perp^r, \perp^r \rangle$. We denote the fact that a relational table T^r has been invoked at a call site n_{cs} by $n_{cs} \in T^r$. Relational execution maps M^r assign each node n a relational value with scope N_n . We denote by $M_{\perp r}^r$ $(M_{\perp r}^r)$ the relational execution maps that assign \perp^r (\top^{r}) to every node.

We denote by \mathcal{V}^{r} the union of all $\mathcal{V}_N^{\mathsf{r}}$ and use similar notation for the sets of all dependency relations and relational tables. For the members of these sets we sometimes omit the subscript indicating the scope when it is irrelevant or clear from the context. We then define a partial order \sqsubseteq^{r} on \mathcal{V}^{r} , which is analogous to the partial order \sqsubseteq on concrete values, except that equality on constants is replaced by subset inclusion on dependency relations:

$$\begin{aligned} r_1 \sqsubseteq^{\mathsf{r}} r_2 & \iff r_1 = \perp^{\mathsf{r}} \lor r_2 = \top^{\mathsf{r}} \lor (r_1, r_2 \in \mathcal{R}_N^{\mathsf{r}} \land r_1 \subseteq r_2) \lor \\ (r_1, r_2 \in \mathcal{T}_N^{\mathsf{r}} \land \forall n_{\mathsf{cs}}. \ r_1(n_{\mathsf{cs}}) \stackrel{.}{\sqsubseteq}^{\mathsf{r}} r_2(n_{\mathsf{cs}})) \end{aligned}$$

We note that aside from \perp^{r} and \top^{r} , only the relational values of identical scope N are comparable. This ordering induces for every N, a complete lattice $(\mathcal{V}_N^{\mathsf{r}}, \sqsubseteq^{\mathsf{r}}, \perp^{\mathsf{r}}, \perp^{\mathsf{r}}, \square^{\mathsf{r}}, \square^{\mathsf{r}})$. The definition of join over relational values is given below.

The meet \sqcap^r is defined similarly.

Lemma 4. Let $V \in \wp(\mathcal{V}_N^r)$. Then, $\sqcup^r V = lub(V)$ and $\sqcap^r V = glb(V)$ subject to \sqsubseteq^r .

The lattices on relational values can be lifted pointwise to a complete lattice on relational execution maps $(\mathcal{M}^{\mathsf{r}}, \sqsubseteq^{\mathsf{r}}, \mathcal{M}^{\mathsf{r}}_{\perp^{\mathsf{r}}}, \mathcal{M}^{\mathsf{r}}_{\top^{\mathsf{r}}}, \dot{\sqcup}^{\mathsf{r}}, \dot{\sqcap}^{\mathsf{r}}).$

Galois connections

We define the meaning of relational execution maps in terms of a Galois connection between \mathcal{M}^{r} and the complete lattice of sets of concrete execution maps $\wp(\mathcal{M})$, which we can then lift to properties $\mathcal{P} = \mathbb{N} \to \wp(\mathcal{M})$. To this end, we first formalize the intuitive meaning of relational values that we have provided above. We do this using a family of Galois connections between the complete lattices \mathcal{V}_N^r and the complete lattice of sets of pairs of execution maps and values, $\wp(\mathcal{M} \times \mathcal{V})$, which is ordered by subset inclusion. These Galois connections share the same concretization function $\gamma^r : \mathcal{V}^r \to \wp(\mathcal{M} \times \mathcal{V})$, which is defined as:

$$\gamma^{\mathsf{r}}(\bot^{\mathsf{r}}) \stackrel{\text{def}}{=} \mathcal{M} \times \{\bot\} \qquad \gamma^{\mathsf{r}}(\top^{\mathsf{r}}) \stackrel{\text{def}}{=} \mathcal{M} \times \mathcal{V}$$
$$\gamma^{\mathsf{r}}(R_{N}^{\mathsf{r}}) \stackrel{\text{def}}{=} \{\langle M, c \rangle \mid D^{\mathsf{r}} \in R_{N}^{\mathsf{r}} \wedge D^{\mathsf{r}}(\nu) = c \land$$
$$\forall n \in N. \ M(n) \in \gamma^{\mathsf{d}}(D^{\mathsf{r}}(n))\} \cup \gamma^{\mathsf{r}}(\bot^{\mathsf{r}})$$
$$\gamma^{\mathsf{r}}(T_{N}^{\mathsf{r}}) \stackrel{\text{def}}{=} \{\langle M, T \rangle \mid \forall n_{\mathsf{cs}}. \ T(n_{\mathsf{cs}}) = \langle v_{i}, v_{o} \rangle \land T^{\mathsf{r}}(n_{\mathsf{cs}}) = \langle r_{i}, r_{o} \rangle$$
$$\land \langle M_{i}, v_{i} \rangle \in \gamma^{\mathsf{r}}(r_{i}) \land \langle M_{o}, v_{o} \rangle \in \gamma^{\mathsf{r}}(r_{o})$$
$$\land M_{i} = \{n_{\mathsf{cs}}\} \ M \land M_{o} = M_{i}[n_{\mathsf{cs}} \mapsto v_{i}]\} \cup \gamma^{\mathsf{r}}(\bot^{\mathsf{r}})$$

Here, the function γ^{d} , which is used to give meaning to dependency relations, is defined by $\gamma^{\mathsf{d}}(c) = \{c\}$ and $\gamma^{\mathsf{d}}(\mathsf{F}) = \mathcal{T}$.

Lemma 5. For all scopes N, γ^{r} is a complete meet-morphism between $\mathcal{V}_N^{\mathsf{r}}$ and $\wp(\mathcal{M} \times \mathcal{V})$.

It follows that there exists a Galois connection between $\wp(\mathcal{M} \times \mathcal{V})$ and \mathcal{V}_N^r for every scope N which has γ^r as its right adjoint [17]. Let $\alpha_N^r : \wp(\mathcal{M} \times \mathcal{V}) \to \mathcal{V}_N^r$ be the corresponding left adjoint, which is uniquely determined by γ^r according to Proposition 1. The meaning of relational execution maps is then given in terms of γ_N^r by the function $\dot{\gamma}^r : \mathcal{M}^r \to \wp(\mathcal{M})$ as follows:

$$\dot{\gamma}^{\mathsf{r}}(M^{\mathsf{r}}) \stackrel{\text{\tiny def}}{=} \{ M \mid \forall n. \langle M, M(n) \rangle \in \gamma^{\mathsf{r}}(M^{\mathsf{r}}(n)) \}$$

From Lemma 5 it easily follows that $\dot{\gamma}^{r}$ is also a complete meet-morphism. We

denote by $\dot{\alpha}^r : \wp(\mathcal{M}) \to \mathcal{M}^r$ the left adjoint of the induced Galois connection.

Abstract domain operations

Before we can define the abstract transformer, we need to introduce a few primitive operations for constructing and manipulating relational values and execution maps (Fig. 3.5). The operation $r_N[n_1=n_2]$ strengthens r_N by enforcing equality

$$r_{N}[n_{1}=n_{2}] \stackrel{\text{def}}{=} \alpha_{N}^{\mathsf{r}}(\gamma^{\mathsf{r}}(r_{N}) \cap \{\langle M, v \rangle \mid M(n_{1})=M(n_{2})\})$$

$$r_{N}[\nu=n] \stackrel{\text{def}}{=} \alpha_{N}^{\mathsf{r}}(\gamma^{\mathsf{r}}(r_{N}) \cap \{\langle M, v \rangle \mid v=M(n)\})$$

$$r_{N}[n \leftarrow r'] \stackrel{\text{def}}{=} \alpha_{N}^{\mathsf{r}}(\gamma^{\mathsf{r}}(r_{N}) \cap \{\langle M'[n \mapsto v'], v \rangle \mid \langle M', v' \rangle \in \gamma^{\mathsf{r}}(r')\})$$

$$M^{\mathsf{r}}[n \leftarrow r_{N}] \stackrel{\text{def}}{=} \prod^{\mathsf{r}} \{M^{\mathsf{r}}[n' \mapsto M^{\mathsf{r}}(n')[n \leftarrow r]] \mid n' \in N\}$$

$$r_{N}[M^{\mathsf{r}}] \stackrel{\text{def}}{=} \prod^{\mathsf{r}} \{r[n \leftarrow M^{\mathsf{r}}(n)] \mid n \in N\}$$

$$r_{N}[N'] \stackrel{\text{def}}{=} \alpha_{N'}^{\mathsf{r}}(\gamma^{\mathsf{r}}(r_{N}) \cap \{\langle M, v \rangle \mid \forall n \in N' \backslash N. M(n) \notin \{\bot, \omega\}\})$$

Figure 3.5: Operations on relational abstract domains

between n_1 and n_2 in the dependency vectors. Similarly, $r_N[\nu=n]$ strengthens r_N by requiring that the described value is equal to the value at node n. We use the operation $r_N[n \leftarrow r']$ to strengthen r_N with r' for n. Logically, this can be viewed as computing $r_N \wedge r'[n/\nu]$. The operation $M^r[n \leftarrow r_N]$ lifts $r_N[n \leftarrow r']$ pointwise to relational execution maps. Note that here for every n' we implicitly rescope M(n')so that its scope includes n. The operation $r_N[M^r]$, conversely, strengthens r_N with all the relations describing the values at nodes N in M^r . Intuitively, this operation is used to push scope/environment assumptions to a relational value. Finally, the *rescoping* operation $r_N[N']$ changes the scope of r_N from N to N' while strengthening it with the information that all nodes in $N' \setminus N$ are neither \perp nor ω . We use the operation $\exists \hat{n}. r_N \stackrel{\text{def}}{=} r[N \setminus \{\hat{n}\}]$ to simply project out, i.e., remove a node \hat{n} from the scope of a given relational value r_N . We (implicitly) use the rescoping operations to enable precise propagation between relational values that have incompatible scopes.

The relational abstraction of a constant c in scope N is defined as $c_N^{\mathsf{r}} \stackrel{\text{def}}{=} \{ D^{\mathsf{r}} \in \mathcal{D}_N^{\mathsf{r}} \mid D^{\mathsf{r}}(\nu) = c \}$. We define $Bool_N^{\mathsf{r}} \stackrel{\text{def}}{=} true_N^{\mathsf{r}} \sqcup^{\mathsf{r}} false_N^{\mathsf{r}}$.

3.3.3 Abstract Transformer

The abstract transformer, $\operatorname{step}^{\mathsf{r}}$, of the relational semantics is shown in Fig. 3.7. It closely resembles the concrete transformer, step , where relational execution maps M^{r} take the place of concrete execution maps M and relational values r take the place of concrete values v. More precisely, the actual transformer $\operatorname{Step}_{k+1}^{\mathsf{r}} \llbracket e \rrbracket (E)(M^{\mathsf{r}}) \stackrel{\text{def}}{=} M^{\mathsf{r}} \stackrel{i}{\sqcup}^{\mathsf{r}} \operatorname{step}_{k+1}^{\mathsf{r}} \llbracket e \rrbracket (E)(M^{\mathsf{r}})$ simply runs $\operatorname{step}^{\mathsf{r}}$ on a given M^{r} and joins the result with M^{r} . Defining the actual transformer this way is non-essential, but simplifies our proofs.

The case for constant values c^{ℓ} effectively replaces the join on values in the concrete transformer using a join on relational values. The constant c is abstracted by the relational value $c^{r}[M^{r}]$, which establishes the relation between c and the other values stored at the nodes that are bound in the environment E. In general, the transformer *pushes* environment assumptions into relational values rather than leaving the assumptions implicit. We decided to treat environment assumptions this way so we can interpret relational values independently of the environment. This idea is implemented in the rules for variables and constants, as explained above, while it is done inductively for other constructs. In the case for variables x^{ℓ} , the abstract transformer replaces **prop** in the concrete transformer by the abstract propagation operator **prop**^r defined in Fig. 3.6. Note that the definition of

$$prop^{r}(T^{r}, \perp^{r}) \stackrel{\text{def}}{=} \langle T^{r}, T_{\perp^{r}}^{r} \rangle \quad prop^{r}(T^{r}, \top^{r}) \stackrel{\text{def}}{=} \langle \top^{r}, \top^{r} \rangle$$

$$prop^{r}(T_{1}^{r}, T_{2}^{r}) \stackrel{\text{def}}{=}$$

$$let T^{r} = \Lambda n_{cs}.$$

$$let \langle r_{1i}, r_{1o} \rangle = T_{1}^{r}(n_{cs}); \langle r_{2i}, r_{2o} \rangle = T_{2}^{r}(n_{cs})$$

$$\langle r_{2i}^{\prime}, r_{1i}^{\prime} \rangle = prop^{r}(r_{2i}, r_{1i})$$

$$\langle r_{1o}^{\prime}, r_{2o}^{\prime} \rangle = prop^{r}(r_{1o}[n_{cs} \leftarrow r_{2i}], r_{2o}[n_{cs} \leftarrow r_{2i}])$$

$$in (\langle r_{1i}^{\prime}, r_{1o} \sqcup^{r} r_{1o}^{\prime} \rangle, \langle r_{2i}^{\prime}, r_{2o} \sqcup^{r} r_{2o}^{\prime} \rangle)$$

$$in \langle \Lambda n_{cs}. \pi_{1}(T^{r}(n_{cs})), \Lambda n_{cs}. \pi_{2}(T^{r}(n_{cs})) \rangle$$

$$prop^{r}(r_{1}, r_{2}) \stackrel{\text{def}}{=} \langle r_{1}, r_{1} \sqcup^{r} r_{2} \rangle \quad (otherwise)$$

Figure 3.6: Value propagation in the relational data flow semantics

prop' assumes that its arguments range over the same scope. Therefore, whenever we use prop' in step', we implicitly use the rescoping operator to make the arguments compatible. For instance, in the call to prop' for variables x^{ℓ} , the scope of $M^{r}(E(x))$ is $N_{E(x)}$ which does not include the variable node E(x) itself. We therefore implicitly rescope $M^{r}(E(x))$ to $M^{r}(E(x))[N_{E \circ \ell}]$ before strengthening it with the equality $\nu = E(x)$. The other cases are fairly straightforward. Observe how in the case for conditionals $(e_{0}^{\ell_{0}} ? e_{1}^{\ell_{1}} : e_{2}^{\ell_{2}})^{\ell}$ we analyze both branches and combine the resulting relational execution maps with a join. In each branch, the map M_{0}^{r} is strengthened with the information about the truth value stored at $E \diamond \ell_{0}$, reflecting the path-sensitive nature of Liquid types.

Lemma 6. The function prop^r is monotone and increasing.

Lemma 7. For every $k \in \mathbb{N}$, $e \in \lambda^d$ and $E \in \mathcal{E}$ such that $\langle e, E \rangle$ is well-formed, then $\operatorname{Step}_k^r[\![e]\!](E)$ is monotone and increasing. $\operatorname{step}_{0}^{\mathsf{r}} \llbracket e^{\ell} \rrbracket (E) (M^{\mathsf{r}}) \stackrel{\text{def}}{=} M^{\mathsf{r}}$ $\operatorname{step}_{k+1}^{\mathsf{r}} \llbracket c^{\ell} \rrbracket (E)(M^{\mathsf{r}}) \stackrel{\text{def}}{=} M^{\mathsf{r}} [E \diamond \ell \mapsto M^{\mathsf{r}}(E \diamond \ell) \sqcup^{\mathsf{r}} c^{\mathsf{r}}[M^{\mathsf{r}}]]$ $\mathsf{step}_{k+1}^{\mathsf{r}}\llbracket x^{\ell} \rrbracket(E)(M^{\mathsf{r}}) \stackrel{\text{\tiny def}}{=}$ let $\langle r'_{r}, r' \rangle = \operatorname{prop}^{\mathsf{r}}(M^{\mathsf{r}}(E(x))[\nu = E(x)], M^{\mathsf{r}}(E \diamond \ell))$ in $M^{\mathsf{r}}[E(x) \mapsto r'_{x}, E \diamond \ell \mapsto r'[M^{\mathsf{r}}]]$ $\operatorname{step}_{h+1}^{\mathsf{r}} \llbracket (e_1^{\ell_1} e_2^{\ell_2})^{\ell} \rrbracket (E) (M^{\mathsf{r}}) \stackrel{\text{def}}{=}$ let $M_1^r = \text{Step}_k^r [\![e_1]\!](E)(M^r); r_1 =_{M_1^r} M_1^r(E \diamond \ell_1)$ in if $r_1 \notin \mathcal{T}^r$ then M_{Tr}^r else let $M_2^{\mathsf{r}} = \mathsf{Step}_k^{\mathsf{r}} \llbracket e_2 \rrbracket (E)(M_1^{\mathsf{r}}); r_2 =_{M_2^{\mathsf{r}}} M_2^{\mathsf{r}}(E \diamond \ell_2)$ in $\langle r'_1, T^{\mathsf{r}}_0 \rangle = \mathsf{prop}^{\mathsf{r}}(r_1, [E \diamond \ell_1 \mapsto \langle r_2, M^{\mathsf{r}}_2(E \diamond \ell) \rangle])$ $\langle r_2', r' \rangle = T_0^{\mathsf{r}}(E \diamond \ell_1)$ in $M_2^r[E \diamond \ell_1 \mapsto r'_1, E \diamond \ell_2 \mapsto r'_2, E \diamond \ell \mapsto r']$ $\operatorname{step}_{k+1}^{\mathsf{r}} \llbracket (\mu f. \lambda x. e_1^{\ell_1})^{\ell} \rrbracket (E) (M^{\mathsf{r}}) \stackrel{\text{def}}{=}$ let $T^{\mathsf{r}} =_{M^{\mathsf{r}}} M^{\mathsf{r}}(E \diamond \ell) \sqcup^{\mathsf{r}} T^{\mathsf{r}}_{\perp}$ in let $\overline{M^{\mathsf{r}}} = \Lambda n_{\mathsf{cs}} \in T^{\mathsf{r}}$. if $\pi_1(T^{\mathsf{r}}(n_{\mathsf{cs}})) = \top^{\mathsf{r}}$ then $M^{\mathsf{r}}_{\mathsf{T}}$ else let $n_x = E \diamond n_{cs} \diamond x$; $n_f = E \diamond n_{cs} \diamond f$; $E_1 = E[x \mapsto n_x, f \mapsto n_f]$ $\langle T_r^{\mathsf{r}}, T_1^{\mathsf{r}} \rangle = \mathsf{prop}^{\mathsf{r}}([n_{\mathsf{cs}} \mapsto \langle M^{\mathsf{r}}(n_r), M^{\mathsf{r}}(E_1 \diamond \ell_1)[n_{\mathsf{cs}} = n_r] \rangle], T^{\mathsf{r}})$ $\langle T_2^{\mathsf{r}}, T_f^{\mathsf{r}} \rangle = \mathsf{prop}^{\mathsf{r}}(T^{\mathsf{r}}, M^{\mathsf{r}}(n_f)); \langle r'_r, r'_1 \rangle = T_r^{\mathsf{r}}(n_{\mathsf{cs}})$ $M_1^{\mathsf{r}} = M^{\mathsf{r}}[E \diamond \ell \mapsto T_1^{\mathsf{r}} \sqcup^{\mathsf{r}} T_2^{\mathsf{r}}, n_f \mapsto T_f^{\mathsf{r}}, n_x \mapsto r_x', E_1 \diamond \ell_1 \mapsto r_1' [n_x = n_{\mathsf{cs}}]]$ in Step^r_k $[e_1](E_1)(M_1^r)$ in $M^{\mathsf{r}}[E \diamond \ell \mapsto T^{\mathsf{r}}] \dot{\sqcup}^{\mathsf{r}} | \overset{\cdot}{n}_{n \leftarrow T^{\mathsf{r}}} \overline{M^{\mathsf{r}}}(n_{\mathsf{cs}})$ $\operatorname{step}_{k+1}^{\mathsf{r}} \llbracket (e_0^{\ell_0} ? e_1^{\ell_1} : e_2^{\ell_2})^{\ell} \rrbracket (E) (M^{\mathsf{r}}) \stackrel{\text{def}}{=}$ let $M_0^{\mathsf{r}} = \mathsf{Step}_k^{\mathsf{r}} \llbracket e_0 \rrbracket (E)(M^{\mathsf{r}}); r_0 =_{M_0^{\mathsf{r}}} M_0^{\mathsf{r}}(E \diamond \ell_0)$ in if $r_0 \not \sqsubset^r Bool^r$ then $M^{\mathsf{r}}_{\top \mathsf{r}}$ else let $M_1^r = \operatorname{Step}_k^r \llbracket e_1 \rrbracket (E) (M_0^r \llbracket E \diamond \ell_0 \leftarrow true^r \sqcap^r r_0])$ $M_2^{\mathsf{r}} = \mathsf{Step}_k^{\mathsf{r}} \llbracket e_2 \rrbracket (E) (M_0^{\mathsf{r}} \llbracket E \diamond \ell_0 \leftarrow false^{\mathsf{r}} \sqcap^{\mathsf{r}} r_0 \rrbracket)$ $r_1 = M_1^{\mathsf{r}}(E \diamond \ell_1); r_2 = M_2^{\mathsf{r}}(E \diamond \ell_2)$ $\langle r'_1, r' \rangle = \operatorname{prop}^{\mathsf{r}}(r_1, M_1^{\mathsf{r}}(E \diamond \ell)); \langle r'_2, r'' \rangle = \operatorname{prop}^{\mathsf{r}}(r_2, M_2^{\mathsf{r}}(E \diamond \ell))$ $\text{ in } M_0^{\mathsf{r}} \stackrel{.}{\sqcup}^{\mathsf{r}} M_1^{\mathsf{r}}[E \diamond \ell_1 \mapsto r_1', E \diamond \ell \mapsto r'] \stackrel{.}{\sqcup}^{\mathsf{r}} M_2^{\mathsf{r}}[E \diamond \ell_2 \mapsto r_2', E \diamond \ell \mapsto r'']$

Figure 3.7: Abstract transformer for relational data flow semantics

3.3.4 Abstract Semantics

The abstract domain of the relational abstraction is given by relational properties $\mathcal{P}^{\mathsf{r}} \stackrel{\text{def}}{=} \mathbb{N} \to \mathcal{M}^{\mathsf{r}}$, which are ordered by $\stackrel{:}{=}^{\mathsf{r}}$, the pointwise lifting of $\stackrel{:}{=}^{\mathsf{r}}$ on \mathcal{M}^{r} to \mathcal{P}^{r} . We likewise lift the Galois connection $\langle \dot{\alpha}^{\mathsf{r}}, \dot{\gamma}^{\mathsf{r}} \rangle$ to a Galois connection $\langle \ddot{\alpha}^{\mathsf{r}}, \ddot{\gamma}^{\mathsf{r}} \rangle$ between the concrete lattice \mathcal{P} and the abstract lattice \mathcal{P}^{r} . The relational abstract semantics $\mathbf{C}^{\mathsf{r}} : \lambda^d \to \mathcal{P}^{\mathsf{r}}$ is then defined as the least fixpoint of Step^r as follows:

$$\mathbf{C}^{\mathsf{r}}\llbracket e \rrbracket \stackrel{\text{\tiny def}}{=} \mathbf{lfp}_{\Lambda k.\ M^{\mathsf{r}}_{\perp}\mathsf{r}}^{\overset{\text{\tiny lef}}{=}\mathsf{r}} \Lambda P^{\mathsf{r}}.\Lambda k.\ \mathsf{Step}^{\mathsf{r}}_{k}\llbracket e \rrbracket(\epsilon)(P^{\mathsf{r}}(k))$$

Theorem 1. The relational abstract semantics is sound, i.e., for all programs e, $\mathbf{C}[\![e]\!] \subseteq \ddot{\gamma}^{r}(\mathbf{C}^{r}[\![e]\!]).$

In the proof of Theorem 1 we show that **prop**^r and **step**^r are sound abstractions of their concrete counterparts **prop** and **step**, from which the claim then follows easily. For the most part, these proofs are straightforward inductions because the definitions of the abstract operators are so closely aligned with the concrete ones. Most intermediate steps just follow directly from properties of the involved Galois connections. The soundness proof of **step**^r is slightly more involved because it is not true that for all well-formed $(e, E), k \in \mathbb{N}, M \in \mathcal{M}, \text{ and } M^r \in \mathcal{M}^r$, if $M \in \dot{\gamma}^r(M^r)$ then $\operatorname{step}_k[\![e]\!](E)(M) \in \dot{\gamma}^r(\operatorname{step}_k^r[\![e]\!](E)(M^r))$. This is because step^r relies on certain invariants satisfied by the (E, M) that the concrete step maintains for the iterates of $\mathbf{S}[\![e]\!]$. For instance, to prove that the implicit uses of the rescoping operator r[N]in step^r are sound, we need the invariant that in the concrete, all variable nodes that are bound in E are not mapped to \perp or ω by M. This can be shown with a simple induction proof. Similar invariants are needed to prove the last step in the case for recursive functions, where we need to show that the join over all call sites $\dot{\sqcup}_{n_{es}\in T} \overline{M}(n_{cs})$ in the concrete transformer is abstracted by the join over all call sites $\bigsqcup_{n_{cs}\in T}^{\mathsf{r}} \overline{M^{\mathsf{r}}}(n_{cs})$ in step^r.

3.4 Collapsed Relational Data Flow Semantics

We now describe the key abstraction step in the construction of Liquid types, which we formalize in terms of a *collapsed (relational data flow) semantics*. This semantics abstracts the relational semantics by (1) collapsing the input/output relations for all call sites in each table into a single one, and (2) collapsing variable nodes by removing their call site components. That is, tables no longer keep track of where they are called and variable bindings no longer track the call stack.

3.4.1 Example

We first provide an intuition behind the collapsed data flow semantics through an example. For that purpose, use the program from Example 3 exhibiting higherorder features of λ^d .

The map computed by the collapsed data flow semantics is presented below. We only show map entries for a portion of reachable nodes.

$$\begin{split} g1, c \to \langle z, \{(\nu:1), \{(z:1,\nu:0)\}\rangle & g2, e \to \langle z, \{(\nu:1), \{(z:1,\nu:1)\}\rangle \\ b \to \langle g_1, \langle z, \{(\nu:1)\}, \{(z:1,\nu:0)\}\rangle, \{(g_1:\mathsf{F},\nu:1), (g_1:\mathsf{F},\nu:0)\}\rangle \\ d \to \langle g_2, \langle z, \{(\nu:1)\}, \{(z:1,\nu:1)\}\rangle, \{(g_2:\mathsf{F},\nu:1), (g_2:\mathsf{F},\nu:0)\}\rangle \\ h \to \langle g, \langle z, \{(\nu:1)\}, \{(z:1,\nu:1), (z:1,\nu:0)\}\rangle, \{(g:\mathsf{F},\nu:1), (g:\mathsf{F},\nu:0)\}\rangle \end{split}$$

As opposed to the concrete and relational semantics that keep track of execution points where the functions are being called in the program, the collapsed semantics completely disregards such information. Consider first the entry for g1 in the map that assigns g1 the *collapsed* table value $\langle z, \{(\nu:1), \{(z:1,\nu:0)\}\rangle$, stating that g1 is a function called at some concrete call site node, symbolically represented as z, with integer 1 as the input and 0 as the output. Note that the information about the actual call site node at which g1 is called is not maintained. Also, observe how z is used to explicitly remember the input value in the output relation. The same explanation can be given for the map entry for g2. The ramifications of not maintaining the precise call site information in tables are best understood by looking at the map entry for h. Since both g1 and g2 are inputs of h in the program and h cannot distinguish between these two inputs using their call site information that is absent, the analysis of the body of h must simultaneously consider both g1 and g_2 as the input argument g of h. That is, the semantics infers that the input g of h is the table value $\langle z, \{(\nu:1)\}, \{(z:1,\nu:1), (z:1,\nu:0)\} \rangle$ that can intuitively be seen as the join of values of g1 and g2. The semantics then concludes that the output relation for h is the relation $\{(g: \mathsf{F}, \nu: 1), (g: \mathsf{F}, \nu: 0)\}$, containing both 0 and 1 as the possible result of calling h. In effect, the result for both calls to h at nodes b and d can be either 0 or 1. Therefore, the collapsed semantics flags the program as potentially unsafe as the value of the divisor expression can possibly be 0. Note that the concrete and relational semantics render this program safe. Overall, by not maintaining precise call site information in tables, the collapsed semantics can introduce imprecision as the input relations of functions coming from different call sites must be merged together when analyzing the function body.

We next introduce the collapsed relational data flow semantics formally. We start by explaining the abstract domains used in defining the semantics.

3.4.2 Abstract Domains

The semantic domains of the collapsed semantics are defined as follows:

$$\begin{split} \hat{n} &\in \hat{\mathcal{N}} \stackrel{\text{def}}{=} \hat{\mathcal{E}} \times Loc & \text{abstract nodes} \\ \hat{E} &\in \hat{\mathcal{E}} \stackrel{\text{def}}{=} Vars \rightharpoonup_{fin} \hat{\mathcal{N}} & \text{abstract environments} \\ u &\in \mathcal{V}_{\hat{N}}^{\text{cr}} ::= \bot^{\text{cr}} \mid \top^{\text{cr}} \mid R_{\hat{N}}^{\text{cr}} \mid T_{\hat{N}}^{\text{cr}} & \text{collapsed values} \\ T_{\hat{N}}^{\text{cr}} &\in \mathcal{T}_{\hat{N}}^{\text{cr}} \stackrel{\text{def}}{=} \Sigma z \in DVar. \ \mathcal{V}_{\hat{N} \setminus \{z\}}^{\text{cr}} \times \mathcal{V}_{\hat{N} \cup \{z\}}^{\text{cr}} & \text{collapsed tables} \\ R_{\hat{N}}^{\text{cr}} &\in \mathcal{R}_{\hat{N}}^{\text{cr}} \stackrel{\text{def}}{=} \wp(\mathcal{D}_{\hat{N}}^{\text{cr}}) & \text{collapsed relations} \\ D_{\hat{N}}^{\text{cr}} &\in \mathcal{D}_{\hat{N}}^{\text{cr}} \stackrel{\text{def}}{=} (\hat{N} \cup \{\nu\}) \rightarrow (Cons \cup \{\mathsf{F}\}) & \text{collapsed dep. vectors} \\ M^{\text{cr}} &\in \mathcal{M}^{\text{cr}} \stackrel{\text{def}}{=} \Pi \hat{n} \in \hat{\mathcal{N}}. \ \mathcal{V}_{\hat{N}_{\hat{n}}}^{\text{cr}} & \text{collapsed maps} \end{split}$$

Since we abstract from call site information in variable nodes, the new semantics no longer distinguishes between expression and variable nodes, leading to the definitions of *abstract nodes* $\hat{n} \in \hat{\mathcal{N}}$ and *abstract environments* $\hat{E} \in \hat{\mathcal{E}}$. The operations *env* and *loc* on abstract nodes are defined as for concrete nodes.

In order to be able to express input-output dependencies in tables without keeping track of call sites explicitly, the collapsed semantics introduces *dependency variables* $z \in DVar$, which symbolically stand for the concrete input nodes of call sites. An *abstract scope* \hat{N} is then defined as a finite set of abstract nodes and dependency variables $\hat{N} \subseteq_{fin} (\hat{N} \cup DVar)$. Analogous to concrete scopes, we define $\hat{N}_{\hat{n}} = \operatorname{rng}(env(\hat{n}))$ as the abstract scope of an abstract node \hat{n} .

Collapsed values $u_{\hat{N}} \in \mathcal{V}_{\hat{N}}^{cr}$ are implicitly indexed by an abstract scope \hat{N} and closely resemble relational values. The only major difference is in the definition of tables, which now only track a single symbolic call site. A *collapsed table* $T_{N}^{cr} \in \mathcal{T}_{\hat{N}}^{cr}$ is a tuple $\langle z, u_1, u_2 \rangle$ where z cannot appear in the scope of the input relational value u_1 but is included in the scope of output relational value u_2 . We write $dx(T_N^{cr})$ to denote the dependency variable z and $io(T_N^{cr})$ to denote the pair $\langle u_1, u_2 \rangle$. The remaining definitions are as for relational values except that abstract nodes take the place of concrete nodes.

The ordering $\sqsubseteq_{\hat{N}}^{cr}$ on collapsed values $\mathcal{V}_{\hat{N}}^{cr}$ resembles the ordering on relational values:

$$u_1 \sqsubseteq^{\mathsf{cr}} u_2 \iff u_1 = \bot^{\mathsf{cr}} \lor u_2 = \top^{\mathsf{cr}} \lor (u_1, u_2 \in \mathcal{R}_{\hat{N}}^{\mathsf{cr}} \land u_1 \subseteq u_2)$$
$$\lor (u_1, u_2 \in \mathcal{T}_{\hat{N}}^{\mathsf{cr}} \land io(u_1) \stackrel{:}{\sqsubseteq}^{\mathsf{cr}} io(u_2))$$

We assume that when we compare two collapsed tables, we implicitly apply α renaming so that they range over the same dependency variable. Again, this
ordering induces, for every \hat{N} , a complete lattice $(\mathcal{V}_{\hat{N}}^{\mathsf{cr}}, \sqsubseteq^{\mathsf{cr}}, \bot^{\mathsf{cr}}, \square^{\mathsf{cr}}, \square^{\mathsf{cr}})$. The
definition of the join \sqcup^{cr} operator is as follows.

The meet \sqcap^{cr} is defined similarly.

Lemma 8. Let $V \in \wp(\mathcal{V}_{\hat{N}}^{cr})$. Then, $\sqcup^{cr}V = lub(V)$ and $\sqcap^{cr}V = glb(V)$ subject to \sqsubseteq^{cr} .

Technically, \sqsubseteq^{cr} is only a pre-order due to α -renaming of tables. We implicitly identify $\mathcal{V}_{\hat{N}}^{cr}$ with its quotient subject to the equivalence relation induced by \sqsubseteq^{cr} . These lattices are lifted pointwise to a complete lattice on collapsed execution maps $(\mathcal{M}^{cr}, \dot{\sqsubseteq}^{cr}, M_{\perp^{cr}}^{cr}, \dot{\Pi}_{\perp^{cr}}^{cr}, \dot{\Box}^{cr}, \dot{\Box}^{cr}).$

Galois connections

In order to provide the meaning of collapsed execution maps and values, we first need to formally relate concrete and abstract nodes, environments, and scopes. To this end, we define the node abstraction function $\rho : \mathcal{N} \to \hat{\mathcal{N}}$ as $\rho(n) \stackrel{\text{def}}{=} (\rho \circ env(n)) \diamond loc(n)$. In essence, ρ recursively removes call site information from concrete variable nodes. For instance, ρ maps two concrete variable nodes $E \diamond n_{cs} \diamond x$ and $E \diamond n'_{cs} \diamond x$ for distinct call sites n_{cs} and n'_{cs} to the same abstract node $(\rho \circ E) \diamond x$. Next, we want to define a Galois connection between relational values $\mathcal{V}_N^{\mathsf{rr}}$ and collapsed values $\mathcal{V}_N^{\mathsf{cr}}$ for given concrete and abstract scopes N and \hat{N} . More precisely, we want to abstract relations over concrete nodes in N by relations over abstract nodes in \hat{N} which may contain dependency variables standing for arbitrary expression nodes at call sites. The corresponding concretization function therefore needs to be defined with respect to a specific matching $\delta : \hat{N} \to N$ of abstract to concrete nodes. These matching functions need to be consistent with node abstraction. We formalize this consistency property by defining the set of all *scope consistent matchings* $\Delta(\hat{N}, N)$ as

$$\Delta(\hat{N}, N) \stackrel{\text{def}}{=} \{ \delta : \hat{N} \to N \mid \forall \hat{n} \in \hat{N} . (\hat{n} \in \hat{\mathcal{N}} \Rightarrow \rho(\delta(\hat{n})) = \hat{n}) \\ \wedge (\hat{n} \in DVar \Rightarrow \delta(\hat{n}) \in \mathcal{N}_e) \}$$

That is, we require that the matching of abstract nodes with concrete nodes is consistent with ρ . For dependency variables, we require that they are matched back to expression nodes. This second requirement formalizes the intuition that dependency variables symbolically stand for arbitrary concrete call site nodes. The meaning of collapsed values $\mathcal{V}_{\hat{N}}^{cr}$ with respect to relational values \mathcal{V}_{N}^{r} is then given subject to a scope consistent matching $\delta \in \Delta(\hat{N}, N)$ by the concretization function
$\gamma^{\mathsf{cr}}_{\delta}: \mathcal{V}^{\mathsf{cr}}_{\hat{N}} \to \mathcal{V}^{\mathsf{r}}_{N}$, which is defined as follows:

$$\begin{split} \gamma_{\delta}^{\mathsf{cr}}(\bot^{\mathsf{cr}}) &\stackrel{\text{def}}{=} \bot^{\mathsf{r}} \quad \gamma_{\delta}^{\mathsf{cr}}(\top^{\mathsf{cr}}) \stackrel{\text{def}}{=} \top^{\mathsf{r}} \quad \gamma_{\delta}^{\mathsf{cr}}(\langle z, u_{i}, u_{o} \rangle) \stackrel{\text{def}}{=} \Lambda n_{\mathsf{cs}}. \ \langle \gamma_{\delta}^{\mathsf{cr}}(u_{i}), \gamma_{\delta[z \mapsto n_{\mathsf{cs}}]}^{\mathsf{cr}}(u_{o}) \rangle \\ \gamma_{\delta}^{\mathsf{cr}}(R_{\hat{N}}^{\mathsf{cr}}) \stackrel{\text{def}}{=} \{ D^{\mathsf{r}} \mid D^{\mathsf{cr}} \in R^{\mathsf{cr}} \land D^{\mathsf{cr}}(\nu) = D^{\mathsf{r}}(\nu) \ \land \forall \hat{n} \in \hat{N}. \ D^{\mathsf{r}}(\delta(\hat{n})) = D^{\mathsf{cr}}(\hat{n}) \} \end{split}$$

The concretization function simply renames the abstract nodes and dependency variables to concrete nodes, as given by δ . For tables, it assigns the same pair of relational values to every possible call site (modulo renaming of n_{cs} in the output relation). Note that the extended matching function $\delta[z \mapsto n_{cs}]$ is again scope consistent for $\hat{N} \cup \{z\}$ and $N \cup \{n_{cs}\}$.

Lemma 9. For every $\delta \in \Delta(\hat{N}, N)$, γ_{δ}^{cr} is a complete meet-morphism between $\mathcal{V}_{\hat{N}}^{cr}$ and \mathcal{V}_{N}^{r} .

It follows then that there exists a unique Galois connection $\langle \alpha_{\delta}^{cr}, \gamma_{\delta}^{cr} \rangle$ between \mathcal{V}_{N}^{r} and $\mathcal{V}_{\hat{N}}^{cr}$, for every $\delta \in \Delta(\hat{N}, N)$ [17]. We then lift the γ_{δ}^{cr} to a concretization function $\dot{\gamma}^{cr}: M^{cr} \to M^{r}$ for collapsed execution maps:

$$\dot{\gamma}^{\mathsf{cr}}(M^{\mathsf{cr}}) \stackrel{\text{def}}{=} \Lambda n. \prod^{\mathsf{r}} \{\gamma_{\delta}^{\mathsf{cr}}(M^{\mathsf{cr}}(\rho(n))) \mid \delta \in \Delta(\hat{N}_{\rho(n)}, N_n)\}$$

That is, if $M^r \in \dot{\gamma}^{cr}(M^{cr})$, then for every node n, the relational value $M^r(n)$ must be subsumed by $\gamma_{\delta}^{cr}(M^{cr}(\rho(n)))$ for all possible consistent matchings δ between N_n and $\hat{N}_{\rho(n)}$. We note that for all concrete nodes n reachable in our transformer step^r, there is a unique $\delta \in \Delta(\hat{N}_{\rho(n)}, N_n)$. This is because $\hat{N}_{\rho(n)}$ contains no dependency variables. Moreover, the locations of the concrete variable nodes bound in reachable environments are the locations of the variables to which they are bound. Hence, they are pairwise distinct. Since ρ preserves the node location information, it is injective on N_n . Hence, δ is uniquely determined by ρ as per the definition of Δ . Since this uniqueness property does not hold in general for unreachable nodes, the meet over all δ is needed so that $\dot{\gamma}^{cr}$ is a complete meet-morphism, which follows directly from its definition and Lemma 9. As usual, the left adjoint of the induced Galois connection is denoted by $\dot{\alpha}^{cr} : \mathcal{M}^r \to \mathcal{M}^{cr}$.

Abstract domain operations

In the new abstract transformer we replace the strengthening and rescoping operations on relational values such as $r_N[n_1 = n_2]$ by their most precise abstractions on collapsed values $u_{\hat{N}}$. For technical reasons in the soundness proof of the abstract transformer, these abstractions are defined with respect to a Galois connection between $\mathcal{V}_{\hat{N}}^{cr}$ and the complete lattice $\Pi N. \Delta(\hat{N}, N) \to \mathcal{V}_N^r$ where the ordering on \mathcal{V}_N^r and the concretizations γ_{δ}^{cr} are simply lifted pointwise. For instance, strengthening with equality between abstract nodes, $u_{\hat{N}}[\hat{n}_1=\hat{n}_2]$, is defined as

$$u_{\hat{N}}[\hat{n}_1 = \hat{n}_2] \stackrel{\text{def}}{=} \bigsqcup^{\mathsf{cr}} \{ \alpha_{\delta}^{\mathsf{cr}}(\gamma_{\delta}^{\mathsf{cr}}(u)[\delta(\hat{n}_1) = \delta(\hat{n}_2)]) \mid N \subseteq \mathcal{N} \land \delta \in \Delta(\hat{N}, N) \}$$

This definition in terms of most precise abstraction should not obscure the simple nature of this operation. E.g., for collapsed dependency relations R^{cr} we simply have:

$$R_{\hat{N}}^{\mathsf{cr}}[\hat{n}_{1} = \hat{n}_{2}] = \{ D^{\mathsf{cr}} \in R_{\hat{N}}^{\mathsf{cr}} \mid \hat{n}_{1}, \hat{n}_{2} \in \hat{N} \Rightarrow D^{\mathsf{cr}}(\hat{n}_{1}) = D^{\mathsf{cr}}(\hat{n}_{2}) \}$$

as one might expect. The other strengthening operations on relational values are abstracted in a similar manner.

3.4.3 Abstract Transformer

The abstract transformer for the collapsed semantics, $\operatorname{step}^{\operatorname{cr}} : \lambda^d \to \hat{\mathcal{E}} \to \mathcal{M}^{\operatorname{cr}} \to \mathcal{M}^{\operatorname{cr}}$, closely resembles the relational abstract transformer $\operatorname{step}^{\operatorname{r}}$ where the operations on relational values are simply replaced by their abstractions on collapsed

values. The only significant difference is in the case for recursive function definitions, shown in Figure 3.8. As collapsed tables do not maintain per-call-site information, the function body e_1 is now evaluated only once for the abstract node \hat{n}_x that abstracts all variable nodes bound to x in the concrete. The propagation between collapsed values, shown in Figure 3.9, is similarly simplified. Note that, technically, propagation works on the representatives of equivalence classes of collapsed values modulo α -renaming of dependency variables in tables. We assume that whenever we propagate between two collapsed tables, the dependency variables in these tables are consistently renamed so that they match on both sides.

$$\begin{aligned} \operatorname{step}_{k+1}^{\operatorname{cr}} & \left[\left(\mu f.\lambda x.e_{1}^{\ell_{1}} \right)^{\ell} \right] (\hat{E}) (M^{\operatorname{cr}}) \stackrel{\operatorname{def}}{=} \\ & \operatorname{let} T^{\operatorname{cr}} =_{M^{\operatorname{cr}}} M^{\operatorname{cr}} (\hat{E} \diamond \ell) \sqcup^{\operatorname{cr}} T_{\perp}^{\operatorname{cr}} \operatorname{in} \\ & \operatorname{if} \pi_{1} (io(T^{\operatorname{cr}})) = \bot^{\operatorname{cr}} \operatorname{then} M^{\operatorname{cr}} [\hat{E} \diamond \ell \mapsto T^{\operatorname{cr}}] \operatorname{else} \\ & \operatorname{if} \pi_{1} (io(T^{\operatorname{cr}})) = \top^{\operatorname{cr}} \operatorname{then} M^{\operatorname{cr}}_{\top^{\operatorname{cr}}} \operatorname{else} \\ & \operatorname{let} \hat{n}_{x} = \hat{E} \diamond x; \, \hat{n}_{f} = \hat{E} \diamond f; \, \hat{E}_{1} = \hat{E} [x \mapsto \hat{n}_{x}, f \mapsto \hat{n}_{f}] \\ & \langle T_{x}^{\operatorname{cr}}, T_{1}^{\operatorname{cr}} \rangle = \operatorname{prop}^{\operatorname{cr}} (\langle dx(T^{\operatorname{cr}}), M^{\operatorname{r}}(n_{x}), M^{\operatorname{r}} (\hat{E}_{1} \diamond \ell_{1}) [dx(T^{\operatorname{cr}}) = \hat{n}_{x}] \rangle, T^{\operatorname{cr}}) \\ & \langle T_{2}^{\operatorname{cr}}, T_{f}^{\operatorname{cr}} \rangle = \operatorname{prop}^{\operatorname{cr}} (T_{1}^{\operatorname{cr}}, M^{\operatorname{cr}} (\hat{n}_{f})); \, \langle u'_{x}, u'_{1} \rangle = io(T_{x}^{\operatorname{cr}}) \\ & M_{1}^{\operatorname{cr}} = M^{\operatorname{cr}} [\hat{E} \diamond \ell \mapsto T_{1}^{\operatorname{cr}} \sqcup^{\operatorname{cr}} T_{2}^{\operatorname{cr}}, \hat{n}_{f} \mapsto T_{f}^{\operatorname{cr}}, \hat{n}_{x} \mapsto u'_{x}, \hat{E}_{1} \diamond \ell_{1} \mapsto u'_{1} [\hat{n}_{x} = dx(T^{\operatorname{cr}})]] \\ & \operatorname{in} \operatorname{Step}_{k}^{\operatorname{cr}} \llbracket e_{1} \rrbracket (\hat{E}_{1}) (M_{1}^{\operatorname{cr}}) \end{aligned}$$

Figure 3.8: Abstract transformer for collapsed semantics (recursive function definition rule)

Lemma 10. The function prop^{cr} is monotone and increasing.

Lemma 11. For every $k \in \mathbb{N}$, $e \in \lambda^d$ and $\hat{E} \in \hat{\mathcal{E}}$ such that $\langle e, \hat{E} \rangle$ is well-formed, then $\operatorname{Step}_k^{\operatorname{cr}}[\![e]\!](\hat{E})$ is monotone and increasing.

In the above, the well-formedness of a pair (e, \hat{E}) of an expression e and abstract

$$\begin{aligned} \mathsf{prop}^{\mathsf{cr}}(T^{\mathsf{cr}}, \bot^{\mathsf{cr}}) &\stackrel{\text{def}}{=} \langle T^{\mathsf{cr}}, \langle dx(T^{\mathsf{cr}}), \bot, \bot \rangle \rangle \\ \mathsf{prop}^{\mathsf{cr}}(T^{\mathsf{cr}}, \top^{\mathsf{cr}}) &\stackrel{\text{def}}{=} \langle \top^{\mathsf{cr}}, \top^{\mathsf{cr}} \rangle \\ \mathsf{prop}^{\mathsf{cr}}(T_{1}^{\mathsf{cr}}, T_{2}^{\mathsf{cr}}) &\stackrel{\text{def}}{=} \\ \mathbf{let} \langle z, u_{1i}, u_{1o} \rangle &= T_{1}^{\mathsf{cr}}; \langle z, u_{2i}, u_{2o} \rangle = T_{2}^{\mathsf{cr}} \\ \langle u'_{2i}, u'_{1i} \rangle &= \mathsf{prop}^{\mathsf{cr}}(u_{2i}, u_{1i}) \\ \langle u'_{1o}, u'_{2o} \rangle &= \mathsf{prop}^{\mathsf{cr}}(u_{1o}[z \leftarrow u_{2i}], u_{2o}[z \leftarrow u_{2i}]) \\ \mathbf{in} \langle \langle z, u'_{1i}, u_{1o} \sqcup^{\mathsf{cr}} u'_{1o} \rangle, \langle z, u'_{2i}, u_{2o} \sqcup^{\mathsf{cr}} u'_{2o} \rangle \rangle \\ \\ \mathsf{prop}^{\mathsf{cr}}(u_{1}, u_{2}) &\stackrel{\text{def}}{=} \langle u_{1}, u_{1} \sqcup^{\mathsf{cr}} u_{2} \rangle \quad (\mathsf{otherwise}) \end{aligned}$$

Figure 3.9: Propagation between values in the collapsed semantics

environment \hat{E} is defined in the similar way as for the concrete environments.

3.4.4 Abstract Semantics

The abstract domain of the collapsed abstract semantics is given by *collapsed* properties $\mathcal{P}^{cr} \stackrel{\text{def}}{=} \mathbb{N} \to \mathcal{M}^{cr}$, ordered by $\overset{\sim}{=}^{cr}$. The Galois connection $\langle \dot{\alpha}^{cr}, \dot{\gamma}^{cr} \rangle$ is lifted to the Galois connection $\langle \ddot{\alpha}^{cr}, \ddot{\gamma}^{cr} \rangle$ between the lattices \mathcal{P}^{r} and \mathcal{P}^{cr} . The collapsed semantics $\mathbf{C}^{cr}[\![\cdot]\!]: \lambda^{d} \to \mathcal{P}^{cr}$ is then defined as

$$\mathbf{C}^{\mathsf{cr}}\llbracket e \rrbracket \stackrel{\stackrel{\simeq}{=}}{=} \mathbf{lfp}_{\Lambda k.\ M_{\perp}^{\mathsf{cr}}}^{\stackrel{\simeq}{=}} \Lambda P^{\mathsf{cr}}.\Lambda k.\ \mathsf{Step}_{k}^{\mathsf{cr}}\llbracket e \rrbracket(\epsilon)(P^{\mathsf{cr}}(k))$$

where $\operatorname{Step}_{k}^{\operatorname{cr}} \llbracket e \rrbracket(\hat{E})(M^{\operatorname{cr}}) \stackrel{\text{\tiny def}}{=} M^{\operatorname{cr}} \stackrel{\operatorname{i}}{\sqcup}^{\operatorname{cr}} \operatorname{step}_{k}^{\operatorname{cr}} \llbracket e \rrbracket(\hat{E})(M^{\operatorname{cr}}).$

Theorem 2. The collapsed abstract semantics is sound, i.e., for all programs e, $\mathbf{C}^{\mathsf{r}}\llbracket e \rrbracket \overset{\sim}{\sqsubseteq} \overset{\mathsf{r}}{\gamma} \overset{\mathsf{cr}}{\mathsf{c}} (\mathbf{C}^{\mathsf{cr}}\llbracket e \rrbracket).$

The proof outline is as follows. First, the collapsed domain operations are by definition sound abstractions of the corresponding relational operations. It follows that prop^{cr} abstracts prop^r. These facts are then used to prove that step^{cr} overapproximates step^r. The only nontrivial case is the rule for recursive function defini-

tions. The proof first shows that after evaluating the body e_1 for every call site n_{cs} in the relational semantics, the resulting map $\overline{M^r}(n_{cs}) = \operatorname{Step}_{k+1}^r \llbracket e_1 \rrbracket (E_1)(M_1^r)$ is in $\dot{\gamma}^{cr}(M_2^{cr})$ where $M_2^{cr} = \operatorname{Step}_{k+1}^{cr} \llbracket e_1 \rrbracket (\hat{E}_1)(M_1^{cr})$. It then follows that $\dot{\bigsqcup}_{n_{cs} \in T^{cr}}^r \overline{M^r}(n_{cs}) \in \dot{\gamma}^{cr}(M_2^{cr})$.

One important implication of node abstraction is that the collapsed semantics only reaches finitely many abstract nodes.

Lemma 12. For all programs e and $k \in \mathbb{N}$, either $\mathbf{C}^{\mathsf{cr}}[\![e]\!](k)(\hat{n}) = M_{\top^{\mathsf{cr}}}^{\mathsf{cr}}$ or the set $\{\hat{n} \in \hat{\mathcal{N}} \mid \mathbf{C}^{\mathsf{cr}}[\![e]\!](k)(\hat{n}) \neq \bot^{\mathsf{cr}}\}$ is finite.

In fact, it can be easily observed from the definition of the abstract transformer that, given some initial abstract environment, every subexpression will be analyzed with at most one abstract environment. This is a consequence of the fact that the transformer in the case of recursive function definitions (where variables are bound) always extends the given environment with the same unique nodes. This implies that there is a one-to-one correspondence between expression/variable nodes and expression/variable locations for the nodes that are reachable.

Example. Consider again our running program e from Example 1. The stabilized collapsed map M_e^{cr} in the fixpoint of $\mathbf{C}^{cr}[\![e]\!]$ is as follows:

$$\begin{split} & \mathrm{id} \to \langle z, \{(\nu \colon 1), (\nu \colon 2)\}, \{(z \colon 1, \nu \colon 1), (z \colon 2, \nu \colon 2)\} \rangle \\ & \mathrm{u}, b, c \to \{(\mathrm{id} \colon \mathsf{F}, \nu \colon 1)\} \qquad e, f \to \{(\mathrm{id} \colon \mathsf{F}, \mathrm{u} \colon 1, \nu \colon 2)\} \\ & a \to \langle z, \{(\mathrm{id} \colon \mathsf{F}, \nu \colon 1)\}, \{(\mathrm{id} \colon \mathsf{F}, z \colon 1, \nu \colon 1)\} \rangle \\ & d \to \langle z, \{(\mathrm{id} \colon \mathsf{F}, \mathrm{u} \colon 1, \nu \colon 2)\}, \{(\mathrm{id} \colon \mathsf{F}, \mathrm{u} \colon 1, z \colon 2, \nu \colon 2)\} \rangle \\ & g^b, g^e \to \{(x \colon 1, \nu \colon 1), (x \colon 2, \nu \colon 2)\} \end{split}$$

Note that in this specific example there is no precision loss in terms of the inputoutput relations. In particular, for the nodes at locations c and f we have a collapsed relation that contains a single vector, which means that we still have precise information about the value obtained as the result of function calls of id. This happens since the input values at the two calls to id are different and function values keep track of precise dependencies on constant inputs. The output relation of id at each call site is strengthened inside prop^{cr} with the actual value of the corresponding input. For instance, at the call site identified with e, the output of id $\{(z: 1, \nu: 1), (z: 2, \nu: 2)\}$ is effectively strengthened with the input $\{(id: F, u: 1, z: 2)\}$ emitted at e, resulting in the value $\{(id: F, u: 1, \nu: 2)\}$.

3.5 Data Flow Refinement Type Semantics

Finally, we obtain a parametric generalization of Liquid types, which we refer to as the *data flow refinement type* semantics. This semantics is an abstraction of the collapsed relational semantics. The abstraction combines three ideas: (1) we abstract dependency relations over concrete constants by abstract relations drawn from some abstract domain of type refinements, (2) we fold the step-indexed properties in the collapsed semantics into single *type maps*, and (3) we introduce widening to enforce convergence of the fixpoint in a finite number of iterations. We will show that the Liquid type inference proposed in [78] is a specific instance of this parametric abstract semantics.

3.5.1 Abstract Domains

Our data flow refinement type analysis is parameterized by the choice of base types and the abstract domain representing type refinements. We assume we are given a set of base types $b \in B$ whose meaning is provided by the concretization function $\gamma_B : \mathsf{B} \to \wp(Cons)$. For simplicity, base types are assumed to be disjoint and every constant has a base type $\forall c, \exists \mathsf{b}. c \in \gamma_B(\mathsf{b})$.

Type refinements a are drawn from a family of relational abstract lattices $\langle \mathcal{A}_{\hat{N}}, \sqsubseteq^{a}, \perp^{a}, \top^{a}, \sqcup^{a}, \sqcap^{a} \rangle$ parameterized by abstract scopes \hat{N} . We assume that there exists a Galois connection $\langle \alpha_{\hat{N}}^{a}, \gamma_{\hat{N}}^{a} \rangle$ between $\mathcal{A}_{\hat{N}}$ and the complete lattice of collapsed dependency relations $\langle \mathcal{R}_{\hat{N}}^{cr}, \subseteq, \emptyset, \mathcal{D}_{\hat{N}}^{cr}, \cup, \cap \rangle$ for every \hat{N} . We further assume that the abstract domains $\mathcal{A}_{\hat{N}}$ come equipped with sound abstractions of the various strengthening operations on collapsed dependency relations. For instance, we assume that there exists an operation $a_{\hat{N}}[\hat{n}_{1} = \hat{n}_{2}]$ such that $\alpha_{\hat{N}}^{a}(\gamma_{\hat{N}}^{a}(a_{\hat{N}})[\hat{n}_{1} = \hat{n}_{2}]) \sqsubseteq^{a} a_{\hat{N}}[\hat{n}_{1} = \hat{n}_{2}]$. Finally, we assume that these abstract domains provide widening operators $\nabla_{\hat{N}}^{a} : \mathcal{A}_{\hat{N}} \times \mathcal{A}_{\hat{N}} \to \mathcal{A}_{\hat{N}}$.

The semantic domains of the data flow refinement type semantics are as follows:

$$t \in \mathcal{V}_{\hat{N}}^{\mathsf{t}} ::= \bot^{\mathsf{t}} \mid \top^{\mathsf{t}} \mid \mathcal{R}_{\hat{N}}^{\mathsf{t}} \mid T_{\hat{N}}^{\mathsf{t}} \qquad \text{data flow refinement types}$$
$$T_{\hat{N}}^{\mathsf{t}} \in \mathcal{T}_{\hat{N}}^{\mathsf{t}} \stackrel{\text{def}}{=} \Sigma z \in DVar. \ \mathcal{V}_{\hat{N} \setminus \{z\}}^{\mathsf{t}} \times \mathcal{V}_{\hat{N} \cup \{z\}}^{\mathsf{t}} \qquad \text{refinement function types}$$
$$R_{\hat{N}}^{\mathsf{t}} \in \mathcal{R}_{\hat{N}}^{\mathsf{t}} \stackrel{\text{def}}{=} \mathsf{B} \times \mathcal{A}_{\hat{N}} \qquad \text{refinement base types}$$
$$M^{\mathsf{t}} \in \mathcal{M}^{\mathsf{t}} \stackrel{\text{def}}{=} \Pi \hat{n} \in \hat{\mathcal{N}}. \ \mathcal{V}_{\hat{N}_{\hat{n}}}^{\mathsf{t}} \qquad \text{refinement type maps}$$

The abstract values of the new semantics, which we refer to as data flow refinement types, closely resembles collapsed abstract values. The major difference is in the treatment of dependency relations. The abstraction of a dependency relation is a refinement base type $R_{\hat{N}}^{t}$, which consists of a base type **b** and an abstract refinement relation *a*. We denote these pairs as $\{\nu : \mathbf{b} \mid a\}$ and write $\mathbf{b}(R_{\hat{N}}^{t})$ for **b** and $a(R_{\hat{N}}^{t})$ for *a*. Refinement function types abstract collapsed tables and consist of an input type t_i , which is bound to a dependency variable *z*, and an output type t_o . We use the familiar function type notation $z : t_i \to t_o$ for these values. As for collapsed tables, we use the functions dx and io for extracting the dependency variable, respectively, the pair of input/output types from function types. As expected, refinement type maps associate a type of appropriate scope with each abstract node. The refinement boolean true and false values are defined as $true^{t} \stackrel{\text{def}}{=} \{\nu :$ bool $| a_t \}$ and $false^{t} \stackrel{\text{def}}{=} \{\nu : \text{bool } | a_f \}$, respectively, where $Bool^{t} \stackrel{\text{def}}{=} \{\nu : \text{bool } | a_{\top} \}$.

The ordering on types is defined as:

$$t_1 \sqsubseteq^{\mathsf{t}} t_2 \iff t_1 = \bot^{\mathsf{t}} \lor t_2 = \top^{\mathsf{t}} \lor (t_1, t_2 \in \mathcal{R}^{\mathsf{t}}_{\hat{N}} \land \mathsf{b}(t_1) = \mathsf{b}(t_2) \land \ a(t_1) \sqsubseteq^{\mathsf{a}} a(t_2))$$
$$\lor (t_1, t_2 \in \mathcal{T}^{\mathsf{t}}_{\hat{N}} \land io(t_1) \stackrel{:}{\sqsubseteq}^{\mathsf{t}} io(t_2))$$

Again, we implicitly quotient types modulo renaming of dependency variables to obtain a partial order that induces a family of complete lattices $(\mathcal{V}_{\hat{N}}^{t}, \sqsubseteq^{t}, \perp^{t}, \top^{t}, \sqcup^{t}, \sqcap^{t})$. The definition of the join \sqcup^{t} operator is as follows.

$$\{\nu: \mathbf{b} \mid a_1\} \sqcup^{\mathbf{t}} \{\nu: \mathbf{b} \mid a_2\} \stackrel{\text{def}}{=} \{\nu: \mathbf{b} \mid a_1 \sqcup^{\mathbf{a}} a_2\} \qquad \bot^{\mathbf{t}} \sqcup^{\mathbf{t}} t \stackrel{\text{def}}{=} t \qquad t \sqcup^{\mathbf{t}} \bot^{\mathbf{t}} \stackrel{\text{def}}{=} t$$
$$z: t_1 \to t_2 \sqcup^{\mathbf{t}} z: t_3 \to t_4 \stackrel{\text{def}}{=} z: t_1 \sqcup^{\mathbf{t}} t_3 \to t_2 \sqcup^{\mathbf{t}} t_4 \qquad t_1 \sqcup^{\mathbf{t}} t_2 \stackrel{\text{def}}{=} \top^{\mathbf{t}} \quad \text{(otherwise)}$$

The meet \sqcap^t is defined similarly.

Lemma 13. Let $V \in \wp(\mathcal{V}_{\hat{N}}^{\mathsf{t}})$. Then, $\sqcup^{\mathsf{t}} V = lub(V)$ and $\sqcap^{\mathsf{t}} V = glb(V)$ subject to \sqsubseteq^{t} .

These are lifted pointwise to a complete lattice of refinement type maps $(\mathcal{M}^t, \stackrel{:}{\sqsubseteq}^t, \mathcal{M}^t_{\perp^t}, \mathcal{M}^t_{\top^t}, \stackrel{:}{\sqcup}^t, \stackrel{:}{\Pi}^t).$

Galois connections

The meaning of refinement types is given by the function $\gamma^{t} : \mathcal{V}_{\hat{N}}^{t} \to \mathcal{V}_{\hat{N}}^{cr}$ that uses the concretization functions for base types and type refinements to map a refinement type to a collapsed value:

$$\begin{split} \gamma^{\mathsf{t}}(\bot^{\mathsf{t}}) &\stackrel{\text{def}}{=} \bot^{\mathsf{cr}} & \gamma^{\mathsf{t}}(\top^{\mathsf{t}}) \stackrel{\text{def}}{=} \top^{\mathsf{cr}} \\ \gamma^{\mathsf{t}}(z \colon t_i \to t_o) &\stackrel{\text{def}}{=} \langle z, \gamma^{\mathsf{t}}(t_i), \gamma^{\mathsf{t}}(t_o) \rangle \\ \gamma^{\mathsf{t}}(\{\nu \colon \mathsf{b} \mid a\}) \stackrel{\text{def}}{=} \{D^{\mathsf{cr}} \mid D^{\mathsf{cr}} \in \gamma^{\mathsf{a}}(a) \land D^{\mathsf{cr}}(\nu) \in \gamma_B(\mathsf{b})\} \end{split}$$

Lemma 14. For all \hat{N} , $\gamma_{\hat{N}}^{\mathsf{t}}$ is a complete meet-morphism between $\mathcal{V}_{\hat{N}}^{\mathsf{t}}$ and $\mathcal{V}_{\hat{N}}^{\mathsf{cr}}$.

We follow the usual notations for the Galois connections induced by $\gamma_{\hat{N}}^{t}$, which we lift pointwise to a Galois connection $\langle \dot{\alpha}^{t}, \dot{\gamma}^{t} \rangle$ between type maps and collapsed maps.

3.5.2 Abstract propagation and transformer

As usual, the domain operations on refinement types can be defined as the best abstractions, using the above introduced Galois connections, of domain operations of the collapsed semantics. Here, we provide a more direct structural definition of refinement type operations by relying on the corresponding operations of the given domain $\mathcal{A}_{\hat{N}}$ of type refinements. First, $t_{\hat{N}}[z \leftarrow t] \stackrel{\text{def}}{=} t$ whenever $z \notin \hat{N}$. Otherwise, the definition goes as shown in Figure 3.10. The strengthening of a bottom and

$$t[z \leftarrow \bot^{t}] \qquad \stackrel{\text{def}}{=} \bot^{t}$$

$$t[z \leftarrow t'] \qquad \stackrel{\text{def}}{=} t \qquad \text{if } t \in \{\bot^{t}, \top^{t}\} \text{ or } t' \in \{\top^{t}\} \cup \mathcal{T}^{t}$$

$$\{\nu : \mathsf{b}_{1} \mid a_{1}\}[z \leftarrow \{\nu : \mathsf{b}_{2} \mid a_{2}\}] \qquad \stackrel{\text{def}}{=} \{\nu : \mathsf{b}_{1} \mid a_{1}[z \leftarrow a_{2}]\}$$

$$(z : t_{1} \rightarrow t_{2})[z_{1} \leftarrow t] \qquad \stackrel{\text{def}}{=} z : (t_{1}[z_{1} \leftarrow t]) \rightarrow (t_{2}[z_{1} \leftarrow \exists z.t]) \qquad t \in \mathcal{R}^{t}$$

Figure 3.10: A structural definition of strengthening operations on refinement types

top type in fact does not perform any proper strengthening. The same applies

for the cases where we strengthen with \top^t or function types. The definition of the strengthening operations for the mentioned cases is guided by simplicity and it is in line with the strengthening operations performed by Liquid type inference that does not explicitly model \perp^t and \top^t types. The strengthening of a function type is defined structurally by performing the strengthening on the corresponding input and output types. Note that the definition for strengthening an output type ensures that the dependency of the output values on the input are not changed by projecting out the dependency variable z from the scope of the type t used to perform the strengthening. The strengthening of a basic refinement type reduces to strengthening of the corresponding refinement domain elements. The structural definitions of other operations on types can be defined similarly.

Lemma 15. Structural strengthening operations are sound.

The propagation operation $\operatorname{prop}^{\mathsf{t}}$ on refinement types is then obtained from $\operatorname{prop}^{\mathsf{cr}}$ by replacing all operations on collapsed values with their counterparts from the refinement type semantics. We omit the definition as it is straightforward. Also, the refinement type abstracting a constant c with abstract scope \hat{N} is defined as $c_{\hat{N}}^{\mathsf{t}} \stackrel{\text{def}}{=} \{\nu : ty(c) \mid \top^{\mathsf{a}}[\nu = c^{\mathsf{t}}]\}$ where $\top^{\mathsf{a}} \in \mathcal{A}_{\hat{N}}$. We define the Boolean refinement type as $Bool_{\hat{N}}^{\mathsf{t}} \triangleq false_{\hat{N}}^{\mathsf{t}} \sqcup^{\mathsf{t}} true_{\hat{N}}^{\mathsf{t}}$.

We obtain the new abstract transformer $\operatorname{Step}^{\mathsf{t}}$ from the collapsed transformer $\operatorname{step}^{\mathsf{cr}}$ in the same fashion. However, the signature of the new transformer is now $\operatorname{Step}^{\mathsf{t}} : \lambda^d \to \hat{\mathcal{E}} \to \mathcal{M}^{\mathsf{t}} \to \mathcal{M}^{\mathsf{t}}$. That is, we additionally eliminate the step-indexing parameter k. The transformer is given as $\operatorname{Step}^{\mathsf{t}} \llbracket e^{\ell} \rrbracket (\hat{E})(M^{\mathsf{t}}) \stackrel{\text{def}}{=} M^{\mathsf{t}} \stackrel{i}{\sqcup}^{\mathsf{t}} \operatorname{step}^{\mathsf{t}} \llbracket e^{\ell} \rrbracket (\hat{E})(M^{\mathsf{t}})$ where $\operatorname{step}^{\mathsf{t}}$ is shown in Figure 3.11. As the transformer definition is parametric in the operations on the abstract domain of type refinements, different choices of refinements do not require changes in the transformer. Moreover, the analysis ab $\mathsf{step}^{\mathsf{t}}\llbracket c^{\ell} \rrbracket(\hat{E})(M^{\mathsf{t}}) \stackrel{\text{\tiny def}}{=} M^{\mathsf{t}}[\hat{E} \diamond \ell \mapsto M^{\mathsf{t}}(\hat{E} \diamond \ell) \sqcup^{\mathsf{t}} c^{\mathsf{t}}[M^{\mathsf{t}}]]$ $\operatorname{step}^{\mathsf{t}} \llbracket x^{\ell} \rrbracket (\hat{E})(M^{\mathsf{t}}) \stackrel{\text{\tiny def}}{=}$ let $\langle t'_{x}, t' \rangle = \operatorname{prop}^{\mathsf{t}}(M^{\mathsf{t}}(\hat{E}(x))[\nu = \hat{E}(x)], M^{\mathsf{t}}(\hat{E} \diamond \ell))$ in $M^{t}[\hat{E}(x) \mapsto t'_{x}, \hat{E} \diamond \ell \mapsto t'[M^{t}]]$ $\mathsf{step}^{\mathsf{t}}\llbracket(e_1^{\ell_1}e_2^{\ell_2})^{\ell}\rrbracket(\hat{E})(M^{\mathsf{t}}) \stackrel{\text{\tiny def}}{=}$ let $M_1^t = \text{Step}^t [\![e_1]\!](\hat{E})(M^t); t_1 =_{M_1^t} M_1^t(\hat{E} \diamond \ell_1)$ in if $t_1 \notin \mathcal{T}^t$ then $M_{\top t}^t$ else let $M_2^t = \text{Step}^t [\![e_2]\!](\hat{E})(M_1^t); t_2 =_{M_2^t} M_2^t(\hat{E} \diamond \ell_2)$ in let $\langle t'_1, T_0^t \rangle = \operatorname{prop}^t(t_1, dx(t_1) : t_2 \to M_2^t(\hat{E} \diamond \ell))$ $\langle t'_2, t' \rangle = io(T_2^t)$ in $M_2^t[\hat{E} \diamond \ell_1 \mapsto t'_1, \hat{E} \diamond \ell_2 \mapsto t'_2, \hat{E} \diamond \ell \mapsto t']$ $\operatorname{step}^{\mathsf{t}} \llbracket (\mu f. \lambda x. e_1^{\ell_1})^{\ell} \rrbracket (\hat{E}) (M^{\mathsf{t}}) \stackrel{\text{def}}{=}$ let $T^{\mathsf{t}} =_{M^{\mathsf{t}}} M^{\mathsf{t}}(\hat{E} \diamond \ell) \sqcup^{\mathsf{t}} T^{\mathsf{t}}_{\perp}$ in if $\pi_1(io(T^t)) = \bot^t$ then $M^t[\hat{E} \diamond \ell \mapsto T^t]$ else if $\pi_1(io(T^t)) = \top^t$ then $M^t_{\top t}$ else let $\hat{n}_x = \hat{E} \diamond x$; $\hat{n}_f = \hat{E} \diamond f$; $\hat{E}_1 = \hat{E}[x \mapsto \hat{n}_x, f \mapsto \hat{n}_f]$ $\langle T_r^t, T_1^t \rangle = \operatorname{prop}^t(dx(T^t): M^t(n_r) \to M^t(\hat{E}_1 \diamond \ell_1)[dx(T^t) = \hat{n}_r], T^t)$ $\langle T_2^{\mathsf{t}}, T_f^{\mathsf{t}} \rangle = \mathsf{prop}^{\mathsf{t}}(T^{\mathsf{t}}, M^{\mathsf{t}}(\hat{n}_f)); \langle t'_x, t'_1 \rangle = io(T_x^{\mathsf{t}})$ $M_1^{\mathsf{t}} = M^{\mathsf{t}}[\hat{E} \diamond \ell \mapsto T_1^{\mathsf{t}} \sqcup^{\mathsf{t}} T_2^{\mathsf{t}}, \hat{n}_f \mapsto T_f^{\mathsf{t}}, \hat{n}_x \mapsto t'_x, \hat{E}_1 \diamond \ell_1 \mapsto t'_1[\hat{n}_x = dx(T^{\mathsf{t}})]]$ in Step^t $\llbracket e_1 \rrbracket (\hat{E}_1)(M_1^t)$ $step^{t} [\![(e_{0}^{\ell_{0}}?e_{1}^{\ell_{1}}:e_{2}^{\ell_{2}})^{\ell}]\!](\hat{E})(M^{t}) \stackrel{\text{def}}{=}$ let $M_0^t = \text{Step}^t [\![e_0]\!](\hat{E})(M^t); t_0 =_{M^t} M_0^t(\hat{E} \diamond \ell_0)$ in if $t_0 \not \sqsubset^t Bool^t$ then $M_{\top t}^t$ else let $M_1^t = \mathsf{Step}^t \llbracket e_1 \rrbracket (\hat{E}) (M_0^t [\hat{E} \diamond \ell_0 \leftarrow true^t \sqcap^t t_0])$ $M_2^{\mathsf{t}} = \mathsf{Step}^{\mathsf{t}} \llbracket e_2 \rrbracket (\hat{E}) (M_0^{\mathsf{t}} [\hat{E} \diamond \ell_0 \leftarrow false^{\mathsf{t}} \sqcap^{\mathsf{t}} t_0])$ $t_1 = M_1^{t}(\hat{E} \diamond \ell_1); t_2 = M_2^{t}(\hat{E} \diamond \ell_2)$ $\langle t'_1, t' \rangle = \operatorname{prop}^{\mathsf{t}}(t_1, M_1^{\mathsf{t}}(\hat{E} \diamond \ell)); \langle t'_2, t'' \rangle = \operatorname{prop}^{\mathsf{t}}(t_2, M_2^{\mathsf{t}}(\hat{E} \diamond \ell))$ in $M_0^t \dot{\sqcup}^t M_1^t [\hat{E} \diamond \ell_1 \mapsto t'_1, \hat{E} \diamond \ell \mapsto t'] \dot{\sqcup}^t M_0^t [\hat{E} \diamond \ell_2 \mapsto t'_2, \hat{E} \diamond \ell \mapsto t'']$

Figure 3.11: Abstract transformer for data flow refinement semantics

stracts from the specific representation of type refinements. If one uses, e.g., polyhedra for type refinements, the actual implementations can choose between constraint and generator representations of polyhedra elements [19], providing more flexibility for tuning the analysis' precision/efficiency trade-off.

Lemma 16. The function prop^t is increasing. Assuming that refinement domain operations are monotone, prop^t is monotone as well.

Lemma 17. For every $k \in \mathbb{N}$, $e \in \lambda^d$ and $\hat{E} \in \hat{\mathcal{E}}$ such that $\langle e, \hat{E} \rangle$ is well-formed, then $\mathsf{Step}^t[\![e]\!](\hat{E})$ is increasing. If the refinement domain operations are monotone, then $\mathsf{Step}^t[\![e]\!](\hat{E})$ is also monotone.

Widening

To ensure that the type analysis converges in finitely many steps, we need a widening operator for $\mathcal{V}_{\hat{N}}^{\mathsf{t}}$. We construct this widening operator from the widening operator $\nabla_{\hat{N}}^{\mathsf{a}}$ on the domains of refinement relations and a *shape widening* operator. In order to define the latter, first define the *shape* of a type using the function $\mathsf{sh}: \mathcal{V}_{\hat{N}}^{\mathsf{t}} \to \mathcal{V}_{\hat{N}}^{\mathsf{t}}$

$$\mathsf{sh}(\bot^{\mathsf{t}}) \stackrel{\text{def}}{=} \bot^{\mathsf{t}} \quad \mathsf{sh}(\top^{\mathsf{t}}) \stackrel{\text{def}}{=} \top^{\mathsf{t}} \quad \mathsf{sh}(R^{\mathsf{t}}) \stackrel{\text{def}}{=} \bot^{\mathsf{t}} \quad \mathsf{sh}(z : t_i \to t_o) \stackrel{\text{def}}{=} z : \mathsf{sh}(t_i) \to \mathsf{sh}(t_o)$$

A shape widening operator is a function $\nabla_{\hat{N}}^{\mathsf{sh}} : \mathcal{V}_{\hat{N}}^{\mathsf{t}} \times \mathcal{V}_{\hat{N}}^{\mathsf{t}} \to \mathcal{V}_{\hat{N}}^{\mathsf{t}}$ such that (1) $\nabla_{\hat{N}}^{\mathsf{sh}}$ is an upper bound operator and (2) for every infinite ascending chain $t_0 \sqsubseteq_{\hat{N}}^{\mathsf{t}} t_0 \sqsubseteq_{\hat{N}}^{\mathsf{t}} \dots$, the chain $\mathsf{sh}(t'_0) \sqsubseteq_{\hat{N}}^{\mathsf{t}} \mathsf{sh}(t'_1) \sqsubseteq_{\hat{N}}^{\mathsf{t}} \dots$ stabilizes, where $t'_0 \stackrel{\text{def}}{=} t_0$ and $t'_i \stackrel{\text{def}}{=} t'_{i-1} \nabla_{\hat{N}}^{\mathsf{sh}} t_i$ for i > 0. In what follows, let $\nabla_{\hat{N}}^{\mathsf{sh}}$ be a shape widening operator. First, we lift $\nabla_{\hat{N}}^{\mathsf{a}}$ to an upper bound operator $\nabla_{\hat{N}}^{\mathsf{ra}}$ on $\mathcal{V}_{\hat{N}}^{\mathsf{t}}$:

$$\begin{split} t \nabla_{\hat{N}}^{\mathsf{ra}} t' &\stackrel{\text{def}}{=} \\ \{\nu : \mathsf{b} \mid a \nabla_{\hat{N}}^{\mathsf{a}} a'\} & \qquad \mathbf{if} \ t = \{\nu : \mathsf{b} \mid a\} \land t' = \{\nu : \mathsf{b} \mid a'\} \\ z : (t_i \nabla_{\hat{N}}^{\mathsf{ra}} t'_i) \to (t_o \nabla_{\hat{N} \cup \{z\}}^{\mathsf{ra}} t'_o) & \qquad \mathbf{if} \ t = z : t_i \to t_o \land t' = z : t'_i \to t'_o \\ t \sqcup_{\hat{N}}^{\mathsf{t}} t' & \qquad \mathbf{otherwise} \end{split}$$

We then define the operator $\nabla_{\hat{N}}^{\mathsf{t}}: \mathcal{V}_{\hat{N}}^{\mathsf{t}} \times \mathcal{V}_{\hat{N}}^{\mathsf{t}} \to \mathcal{V}_{\hat{N}}^{\mathsf{t}}$ as the composition of $\nabla_{\hat{N}}^{\mathsf{sh}}$ and $\nabla_{\hat{N}}^{\mathsf{ra}}$, that is, $t \nabla_{\hat{N}}^{\mathsf{t}} t' \stackrel{\text{def}}{=} t \nabla_{\hat{N}}^{\mathsf{ra}} (t \nabla_{\hat{N}}^{\mathsf{sh}} t')$.

Lemma 18. $\nabla^{\mathsf{t}}_{\hat{N}}$ is a widening operator.

3.5.3 Abstract semantics

We identify the properties \mathcal{P}^{t} of the data flow refinement type semantics with type maps, $\mathcal{P}^{\mathsf{t}} \stackrel{\text{def}}{=} \mathcal{M}^{\mathsf{t}}$. Then we connect \mathcal{P}^{t} to collapsed properties $\mathcal{P}^{\mathsf{cr}}$ via the Galois connection induced by the concretization function $\ddot{\gamma}^{\mathsf{t}} \stackrel{\text{def}}{=} \Lambda M^{\mathsf{t}} . \Lambda k. \ \dot{\gamma}^{\mathsf{t}}(M^{\mathsf{t}})$. We lift the widening operators $\nabla^{\mathsf{t}}_{\hat{N}}$ pointwise to an upper bound operator $\dot{\nabla}^{\mathsf{t}} : \mathcal{P}^{\mathsf{t}} \times \mathcal{P}^{\mathsf{t}} \to \mathcal{P}^{\mathsf{t}}$. The data flow refinement semantics $\mathbf{C}^{\mathsf{t}}[\![\cdot]\!] : \lambda^d \to \mathcal{P}^{\mathsf{t}}$ is then defined as the least fixpoint of the widened iterates of $\mathsf{Step}^{\mathsf{t}}$:

$$\mathbf{C}^{\mathsf{t}}\llbracket e \rrbracket \stackrel{\mathrm{\tiny def}}{=} \mathbf{lfp}_{M_{\perp}^{\mathsf{t}}}^{\overset{\mathrm{\tiny L}}{=}^{\mathsf{t}}} \Lambda M^{\mathsf{t}}. \ (M^{\mathsf{t}} \ \dot{\bigtriangledown}^{\mathsf{t}} \ \mathsf{Step}^{\mathsf{t}}\llbracket e \rrbracket(\epsilon)(M^{\mathsf{t}}))$$

Theorem 3. The refinement type semantics is sound and terminating, i.e., for all programs e, $\mathbf{C}^{\mathsf{t}}[\![e]\!]$ converges in finitely many iterations. Moreover, $\mathbf{C}^{\mathsf{cr}}[\![e]\!] \sqsubseteq^{\mathsf{cr}}$ $\ddot{\gamma}^{\mathsf{t}}(\mathbf{C}^{\mathsf{t}}[\![e]\!])$.

The soundness proof and termination follow straightforwardly from the properties of the involved Galois connections, Lemma 18, Lemma 12, and Proposition 3. We say that a type map M^{t} is *safe* if for all abstract nodes \hat{n} , $M^{t}(\hat{n}) \neq \top^{t}$. Let further $\ddot{\gamma} \stackrel{\text{def}}{=} \ddot{\gamma}^{r} \circ \ddot{\gamma}^{cr} \circ \ddot{\gamma}^{t}$. The next lemma states that safe type maps yield safe properties. It follows immediately from the definitions of the concretizations.

Lemma 19. For all safe type maps M^t , $\ddot{\gamma}(M^t) \stackrel{.}{\subseteq} P_{\mathsf{safe}}$.

A direct corollary of this lemma and the soundness theorems for our abstract semantics is that we can use the data flow refinement type semantics to prove the safety of programs.

Corollary 1. For all programs e, if $C^t[\![e]\!]$ is safe then so is e.

3.5.4 Instantiating the framework

In order to obtain an instance of our data flow refinement type analysis, one has to choose basic types, the abstract domain of type refinements with the accompanying domain operations and widening operator, and the shape widening operator.

3.6 Liquid Type Semantics

We now recast Liquid types as an instantiation of the parametric data flow refinement type inference.

Abstract domains. The set of basic types consists of types such as int, bool, etc. The abstract domain of type refinements $\mathcal{A}_{\hat{N}}$ is induced by a given set of qualifiers Q by instantiating them with the given abstract scope \hat{N} . Hence, the abstract domain of type refinements is $\wp(Q_{\hat{N}})$ whose elements are interpreted conjunctively and quotiented by the pre-order induced by entailment modulo some appropriate decidable first-order theory (e.g., linear arithmetic with uninterpreted functions as in [78]). This quotient is a complete lattice.

Domain Operations. The strengthening operations on our model of Liquid types are defined structurally, as in Figure 3.10, by relying on the strengthening operations provided by the Liquid domain of type refinements. These operations correspond to the logical operations used in the original Liquid types work: a refinement domain element a can be represented as a logical predicate. Based on this representation, $a[\nu=\hat{n}]$ corresponds to $a \wedge \nu = \hat{n}$ and $a[\hat{n} \leftarrow a']$ is $a \wedge a'[\hat{n}/\nu]$ where the latter operation is a syntactic substitution of \hat{n} for ν in 'a. The operation $\exists z.t$ of projecting out a dependency variable z from a scope of a type t can be defined as removing the conjuncts of t that involve z. In practice, Liquid types avoid using the rescoping operation in the strengthening operations by renaming dependency variables, thus making sure that variables in types are always unique.

Widenings. We note that the Liquid Type inference first infer basic types and only then it infers type refinements. Our data flow refinement semantics performs the two inference phases simultaneously, thus not forcing the analysis to make an a priori decision about how the shapes of the types of recursive functions should be approximated. To obtain Liquid Types in our framework, we need a shape widening operator that mimics the approximation of recursive functions performed by the Hindley-Milner type system on which Liquid Types are based. The connection between this approximation and widening is already well understood [15]. Since the refinement domain $\mathcal{A}_{\hat{N}}$ has finite height, the type refinement widening operator $\nabla^{\mathbf{a}}_{\hat{N}}$ is simply the join operator on $\wp(Q_{\hat{N}})$. **Connection to original Liquid types.** We now connect our formalization of Liquid types with the one from the original work [78]. The subtyping relation on our model of Liquid types is defined as follows

$$\{\nu \colon \mathsf{b} \mid a_1\} <:^q \{\nu \colon \mathsf{b} \mid a_2\} \stackrel{\text{def}}{\iff} a_1 \sqsubseteq^{\mathsf{a}} a_2$$
$$z : t_1 \rightarrow t_2 <:^q z : t_3 \rightarrow t_4 \stackrel{\text{def}}{\iff} t_3 <:^q t_1 \land t_2[z \leftarrow t_3] <:^q t_4[z \leftarrow t_3]$$

Our subtyping relation coincides with the original one [78, p. 6]. The main difference is the original subtyping is defined subject to typing environment assumptions that we push directly into the type refinements. It can be shown that subtyping between types implies that propagation between these types has reached a fixpoint.

Lemma 20. If $t_1 <:^q t_2$, then $\langle t_1, t_2 \rangle = \text{prop}^{t}(t_1, t_2)$.

We use this fact to show that a valid derivation of the Liquid typing judgment relation can be obtained from the results of our analysis. To this end, we first define a typing relation \vdash_q adapted from [78, p. 6]. As with Liquid types, we define the typing rules with a more standard notion of typing environments $\Gamma \in \Xi \stackrel{\text{def}}{=} Vars \rightarrow \mathcal{V}^{\mathsf{t}}$ that map program variables to types. In order to connect the typing rules with the results produced by our abstract interpretation, we define the rules subject to an injective mapping $\kappa : Vars \rightarrow \hat{\mathcal{N}}$ from variables to abstract nodes. Let Ξ^{κ} be the set of liquid typing environments whose domain is $\mathsf{dom}(\kappa)$. The abstract scope induced by κ is then simply $\hat{N}_{\kappa} \stackrel{\text{def}}{=} \mathsf{rng}(\kappa)$. Also, the induced abstract environment is $\kappa(\Gamma) \stackrel{\text{def}}{=} \Lambda x \in \mathsf{dom}(\Gamma).\kappa(x)$. Given an encoding of abstract environment nodes κ , our Liquid typing relation $\vdash_q^{\kappa}: \Sigma\Gamma \in \Xi^{\kappa}.\lambda^d \times \mathcal{V}_{\hat{N}_{\kappa}}^{\mathsf{t}}$ is given by the typing rules shown in Figure 3.12. We assume we are given an existing base typing function $ty : Cons \rightarrow \mathsf{B}$ (e.g., Hindley-Milner) for constants and an existing strengthening operation $a[\nu = c]$ on type refinements that sets ν to c in the concrete (implemented

$$\begin{array}{c|c} \displaystyle \frac{x \in \operatorname{dom}(\Gamma) \quad \Gamma(x) <:^{q} t \quad t = \Gamma(x)[\nu = \kappa(x)][\Gamma]^{\kappa}}{\Gamma \vdash_{q}^{\kappa} x : t} & \overline{\Gamma \vdash_{q}^{\kappa} c : c^{\mathsf{t}}[\Gamma]^{\kappa}} \\ \\ \displaystyle \frac{\Gamma \vdash_{q}^{\kappa} e_{1} : t_{1} \quad \Gamma \vdash_{q}^{\kappa} e_{2} : t_{2} \quad t_{1} <:^{q} dx(t_{1}) : t_{2} \to t}{\Gamma \vdash_{q}^{\kappa} e_{1} e_{2} : t} \\ \\ \displaystyle \frac{t <:^{q} t_{f} \qquad dx(t) : t_{x} \to t_{1} <:^{q} t}{\Gamma \vdash_{q}^{\kappa} uf_{2} : t} \\ \\ \displaystyle \frac{\kappa_{1} = \kappa.x : \kappa(\Gamma) \diamond x.f : \kappa(\Gamma) \diamond f \quad \Gamma.x : t_{x}.f : t_{f} \vdash_{q}^{\kappa} e : t_{1}}{\Gamma \vdash_{q}^{\kappa} \mu f \lambda x. e : t} \\ \\ \Gamma \vdash_{q}^{\kappa} x : t_{0} \quad t_{0} <:^{q} Bool^{\mathsf{t}} \qquad t_{1} <:^{q} t \quad t_{2} <:^{q} t \\ \Gamma_{1} = \Lambda x \in \operatorname{dom}(\Gamma). \Gamma(x)[\kappa(x) \leftarrow t_{0} \sqcap^{\mathsf{t}} true^{\mathsf{t}}] \qquad \Gamma_{1} \vdash_{q}^{\kappa} e_{1} : t_{1} \\ \\ \Gamma_{2} = \Lambda x \in \operatorname{dom}(\Gamma). \Gamma(x)[\kappa(x) \leftarrow t_{0} \sqcap^{\mathsf{t}} false^{\mathsf{t}}] \qquad \Gamma_{2} \vdash_{q}^{\kappa} e_{2} : t_{2} \\ \hline \Gamma \vdash_{q}^{\kappa} (x ? e_{1} : e_{2}) : t \end{array}$$

Figure 3.12: Adapted Liquid typing rules

as adding a logical conjunct $\nu = c$ to a). We use these operations and κ to properly strengthen constants with the environment assumptions. More precisely, $c^{\mathsf{t}}[\Gamma]^{\kappa}$ is a shorthand for $c^{\mathsf{t}}[M^{\mathsf{t}}]$ where $M^{\mathsf{t}} = \Lambda \hat{n}$. $\hat{n} \in \mathsf{rng}(\kappa) ? \Gamma(\kappa^{-1}(\hat{n})) : \bot^{\mathsf{t}}$.

The first obvious difference between our typing rules and the original Liquid typing rules is that we need to maintain the mapping κ between program variables and environment nodes. Again, we use this mapping to easier connect the typing rules to our semantics and properly push environment assumptions to types that happens on the leaves of typing derivations. Second, the rules in the original paper are nondeterministic: the subtyping rule can be applied at any point in the typing derivation. Our rules, on the other hand, are deterministic as we directly push subtyping constraints at any point where there is a potential flow of data. Lastly, we assume for simplicity that conditionals are A-normalized, as is the case with the original Liquid type inference algorithm implementation [78, sec. 4.4].

It can then be shown that the Liquid typing rules specify a fixpoint of our

Liquid type semantics. Let *valid* be typing environments that do not contain \perp^t and \top^t types. We can now formalize the result that a solution computed by type inference based on the Liquid typing rules can be obtained using our abstract interpretation.

Theorem 4. Let e^{ℓ} be an expression, κ an injective encoding of abstract variable nodes, $\Gamma \in \Xi^{\kappa}$ a valid typing environment, and $t \in \mathcal{V}_{\hat{N}_{\kappa}}^{\mathsf{t}}$. If $\Gamma \vdash_{q}^{\kappa} e^{\ell} : t$, then there exists a safe type map M^{t} such that $\mathsf{Step}^{\mathsf{t}}[\![e^{\ell}]\!](\kappa(\Gamma) \diamond \ell)(M^{\mathsf{t}}) = M^{\mathsf{t}}$ and $M^{\mathsf{t}}(\kappa(\Gamma) \diamond \ell) = t$.

The proof goes by showing that one can construct a safe type map from the typing derivation corresponding to $\Gamma \vdash_q^{\kappa} e^{\ell} : t$. This map is then shown to be a fixpoint of the abstract transformer.

Design space. We note that our Liquid type inference is sound by construction and the points where the analysis loses precision are made explicit in our choice of widening operators and Galois connections used in the relational, collapsed, and refinement type semantics. This is in contrast to the original Liquid types work where the soundness of the inference is argued as follows. Given an expression e, a Liquid type environment Γ , and a qualifier set Q, it is first shown that for a Liquid type τ inferred by the algorithm Solve (Γ, e, Q) , the original Liquid typing relation holds $\Gamma \vdash_Q e : \tau$ [78, Theorem 2]. Then, this relation is shown to be conservative subject to an "exact" typing relation $\Gamma \vdash e : \tau$ [78, Theorem 1], which is in turn proven to be sound subject to a standard operational semantics elsewhere [51]. This indirect soundness argument obscures where the analysis loses precision. In particular, it does not make formally explicit how \vdash_Q abstracts \vdash and how \vdash abstracts the concrete semantics.

3.7 Related Work

We now discuss work related to refinement type inference.

Refinement type systems. One of the earliest works on refinement types goes back to [28]. The authors present a type system for ML where a programmerannotated definition of an algebraic datatype induces a finite set of its subtypes, called refinements. These refinements constitute a finite-height (non-relational) abstract domain of type refinements. The presented type system uses intersection and union types in order to perform modular refinement type inference. Xi and Pfenning [96] introduce a modular and semi-automatic dependent type analysis for ML where refinement types (called index types) are defined over a parametric constraint domain. Techniques used in the original Liquid types [78] work were extended to support Haskell [92], abstractions over type refinements [91, 90], and inference of heap invariants for C programs [79]. Our results can serve as the basis for reformulating these works in terms of abstract interpretation. We also note that the work in [92] provides a denotational semantic model of Liquid types. However, this model is used for proving the soundness of Liquid typing rules and the inference algorithm. This is a different task than systematically and constructively designing Liquid type inference as an abstract semantics of programs, for which denotational semantics is not a good fit as it is inherently modular and it does not make explicit how data flows in programs. Zhu and Jagannathan [99] present an algorithm that infers data flow invariants of functional programs using a refinement type inference algorithm that solves subtyping constraints per function body and then lazily propagates the information from call sites to corresponding functions, akin to our data flow refinement type system. Moreover, their system

uses intersection types labeled by call site information to infer the type of a function based on different call sites. This suggests that their system can be formalized using our framework where the collapsed semantics is refined by a partition of call sites. Their inference algorithm, however, is not guaranteed to terminate. Viewed using our framework, this is a consequence of the fact that the abstract domain of type refinements is of infinite height and the fixpoint computation is not using a widening operator. Knowles and Flanagan [51] introduce a general refinement type system that infers data flow invariants for functional programs where refinements are conjunctions and disjunctions of unrestricted Boolean expressions drawn from the language of programs. Similar to Liquid types, their system first computes the basic type structure of the program and then sets up constraints encoding the program data flow. When viewed through the lens of our work, their inference algorithm works over an infinite-height abstract domain and uses certain widening operator to ensure finite convergence. When analyzing recursive functions, the resulting function type is itself defined as a least fixpoint, since the fixpoint operator is anyhow available in the language. The inference algorithm proposed in [43] can be seen as an implementation of our data flow refinement type system where (collapsed) data flow equations of a higher-order program are transformed into a first-order program that is then analyzed using various abstract domains. The results are then pushed back to the original higher-order program. Refinement type techniques have also been used in the context of dynamic languages [93, 48, 53]. The work in [53] uses Galois connections to formalize the intuition that gradual refinement types abstract static refinement types. However, the resulting gradual type system does not take full advantage of the framework of abstract interpretation, i.e., the type system is not calculationally constructed as an abstraction

of a concrete program semantics. We note that none of the above works formally describe their refinement type systems as abstract interpretations in this original sense.

Semantics of higher-order programs. Work in [41] and similarly [76] use flow graphs to assign concrete meaning to higher-order programs. Nodes in the graph correspond to nodes of our data flow semantics. Their semantics models functions as expression nodes storing the location of the function definition expression. Argument values then flow directly from call sites to the function. Hence, this semantics has to make non-local changes when analyzing function applications. In contrast, our data flow semantics treats functions as tables and propagates values backwards in lockstep with the program structure, which is more suitable for building type systems. The minimal function graph semantics [46, 45] models functions as tables whose domains only contain inputs observed during program execution, similar to our data flow semantics. However, minimal function graphs do not explicitly model flow of information in a program, which is not well suited for designing data flow refinement type systems. More precisely, similar to minimal function graphs, our new semantics computes for every function, the pairs of input and output values for the calls to that function that are actually observed during program execution. While minimal function graph semantics accumulates this information in a separate dedicated global structure, our semantics stores information on function calls more locally. That is, at each point in the flow graph where a function value is observed, we record those calls to the function that happen from that point onward as the function value continues to flow through the program. The semantics makes explicit how input values propagate backwards from the call sites of a function to the function definition and, conversely, how output values propagate from the definition back to the call sites. Our semantics precisely captures the properties abstracted by Liquid types and is defined structurally on the program syntax, allowing for an easier formal connection to the typing rules underlying Liquid type inference.

There is a large body of work on control and data flow analysis on higher-order programs and the concrete semantics they overapproximate [58, 69, 65, 44, 18]. However, these works either do not make program data flow explicit or they treat program functions as graph nodes, continuations, or closures, all of which are not well-suited for developing data flow refinement type systems as an abstract semantics since they do not make explicit how data flows in functional programs, i.e., how inputs flow from function call sites to the function definition and how the outputs flow back in the other direction. To the best of our knowledge, the data flow semantics introduced in this work is the first concrete semantics where the notion of data flow is made explicit and the concrete transformer only makes local changes to the global execution map, enabling abstractions that naturally lead to structural typing rules.

Types as abstract interpretations. The formal connection between types and abstract interpretation is studied by Cousot [15]. His paper shows how to construct standard polymorphic type systems via a sequence of abstractions starting from a concrete denotational call-by-value semantics. Monsuez [62] similarly shows how to use abstract interpretation to design polymorphic type systems for call-by-name semantics and model advanced type systems such as System F [60, 61, 63, 64]. Techniques presented in [32, 31] use the abstract interpretation view of types

to design new type inference algorithms for ML-like languages by incorporating widening operators that infer more precise types for recursive functions. Unlike these works that focus on modular type inference, this thesis is instead focused on refinement type systems that perform a whole program analysis, an example of which are Liquid types. The work in [42] uses abstract interpretation techniques to develop a type system for JavaScript that is able to guarantee absence of common errors such as invoking a non-function value or accessing an undefined property. Further, the work in [29] uses Galois connections to relate gradual types with abstract static types to build a gradual type system from a static one. Harper [37] introduces a framework for constructing dependent type systems from operational semantics based on the PER model of types. Although this work does not build on abstract interpretation, it relies on the idea that types overapproximate program behaviors and that they can be derived using a suitable notion of abstraction.

Chapter 4

Type Error Localization

The contents of this chapter are previously published in a slightly modified form in [72, 73]. Proofs for all the theorems, lemmas, and corollaries stated in this chapter can be found in Appendix B, unless already presented in the chapter.

4.1 Overview

In this section, we provide an overview of our approach to the problem of type error localization introduced in § 1 and explain it through several illustrative examples.

The high-level execution flow of our constraint-based type error localization is shown in Figure 4.1. The framework can be viewed as a compiler plug-in. When type inference takes place, our algorithm starts by generating typing assertions for the given input program. The constraint generation incorporates a compilerspecific criterion for ranking type error sources by appropriately assigning weights to the assertions. After the constraint generation finishes, the produced annotated assertions, constituting a typing constraint, are passed to a weighted MaxSMT solver. If the input program has a type error, the generated constraint is unsatis-



Figure 4.1: High-level overview of constraint-based type error localization. Thick arrows represent a looping interaction between a compiler and the SMT solver.

fiable. In this case, the MaxSMT solver finds all error sources that are minimum subject to the specified ranking criterion. The compiler can then iteratively interact with the solver to further rank and filter the error sources to generate an appropriate error message. This way, our framework can support interaction with the programmer. In particular, the compiler can take feedback from the programmer and pass it to the solver to guide the search for the error source the programmer needs. In the following, we describe the actual algorithm in more detail and highlight the main features of our approach.

4.1.1 Minimum Error Sources

First, let us make the notion of minimum error source more precise. An error source is a set of program locations that need to be fixed so that the erroneous program becomes well typed. Usually, not all error sources are equally likely to cause the error, and the compiler might prefer some error sources over others. In our framework, the compiler provides a criterion for ranking error sources by assigning weights to program locations. Our algorithm then finds the error sources with minimum cumulative weight. As an example, consider the following OCaml program:

1 let f x = print_int x in
2 let g x = x + 1 in
3 let x = "hi" in
4 f x;
5 g x

Note that the program cannot be typed since the functions f and g expect an argument of type int, but both are applied to x whose type is string. The program has several possible error sources. One error source is the declaration of x. Namely, changing the declaration of x to, say, an integer constant would make the program well typed. Another error source is given by the two function applications of f and g. For example, replacing both f and g by print_string would yield another well-typed program.

Now consider a ranking criterion that assigns each program location equal weight 1. Then the cumulative weight of the first error source is 1, while the weight of the second error source is 2. Our algorithm would therefore report the first error source and discard the second one. This simple ranking criterion thus minimizes the number of program locations that need to be modified so that the program becomes well typed. Later in this chapter, we discuss more elaborate ranking criteria that are useful for type error localization.

4.1.2 Reduction to MaxSMT

Next, we show how the reduction to weighted MaxSMT works through another example. Consider the following OCaml program:

let x = "hi" in not x

Clearly, the program is not well typed as the operation **not** on Booleans is applied to a variable **x** of type **string**.

Our constraint generation procedure takes the program and generates a set of typing assertions using the OCaml type inference rules. This set of assertions constitutes a typing constraint. A formal description of the constraint generation can be found in § 4.3.2. For our example program, the constraint generation produces the following set of assertions:

$$\alpha_{not} = \mathsf{fun}(\mathsf{bool}, \mathsf{bool}) \qquad \qquad [\text{Def. of not}] \qquad (4.1)$$

$$\alpha_{app} = \mathsf{fun}(\alpha_i, \alpha_o) \qquad \qquad \mathsf{not} \ \mathsf{x} \tag{4.2}$$

$$\alpha_{app} = \alpha_{not} \qquad \text{not} \qquad (4.3)$$

$$\alpha_i = \alpha_x \qquad \qquad \mathbf{x} \qquad (4.4)$$

$$\alpha_x = \text{string} \qquad \qquad \mathbf{x} = \text{"hi"} \qquad (4.5)$$

The constraint encoding preserves the mapping between assertions and associated program locations. That is, each assertion comes from a particular program expression shown to the right of the assertion. The assertion (4.1) is generated from the definition of the function **not** in OCaml's standard library [55]. It specifies the type α_{not} of **not** as a function type from **bool** to **bool**. The assertions (4.2) to (4.4) specify the types of all subexpressions of the expression **not x** in the program. Finally, the assertion (4.5) specifies the type of **x**, according to its declaration.

The generated typing constraint is interpreted in the theory of inductive data types, which is supported by many SMT solvers [5, 21, 23]. In this theory, the terms α_{app} , α_i , α_o , α_{not} , and α_x stand for type variables, while the terms **boo**l

and string are type constants. The function fun is a type constructor that maps an argument type α_i and a result type α_o to a function type from α_i to α_o . The theory interprets fun as an injective constructor. Hence, the assertions (4.1) to (4.3) together imply $\alpha_i = \text{bool}$. Type constants such as bool and string are interpreted as pairwise distinct values. The assertions (4.4) and (4.5) imply $\alpha_i = \text{string}$, so we conclude bool = string together with the assertions (4.1) to (4.3) – a contradiction. Consequently, the generated typing constraint is unsatisfiable, confirming there is a type error.

As in the previous example, we first assume a ranking criterion that assigns each program location equal weight 1. The problem of finding the minimum error sources in the program is then encoded into a weighted MaxSMT problem by assigning weight 1 to each assertion in the generated constraint. The weighted MaxSMT solver computes all subsets of assertions in the typing constraint that are satisfiable and have maximum cumulative weight. Thus, the complements of those sets encode the minimum error sources.

For the running example, there are altogether five complement sets of maximum satisfiable subsets, each consisting of a single assertion (4.1)-(4.5). Removing any single assertion leaves the remaining assertions satisfiable. In other words, by fixing the program expression associated with the removed assertion, the input program becomes well typed. For example, consider the assertion (4.3). Removal of that assertion can be interpreted as using a different function instead of not in the application not x. Removal of the assertion (4.1) can be seen as changing the internals of not so that it applies to string values. Hence, the meaning associated with removing a certain assertion can be utilized for suggesting possible error fixes.

4.1.3 Ranking Criteria

Not all error sources are equally useful for fixing the type error. In the previous example, it is unlikely that the true error source is located in the implementation of the function not, since it is part of the standard library. Compilers might want to eliminate such error sources from consideration. Similarly, the error source corresponding to the removal of assertion (4.2) implies that the programmer should use some other expression than not x in the program. This error source seems less specific than the remaining error sources, so compilers might want to rank it as less likely. In general, compilers might want to rank error sources by some definition of usefulness. In our framework, such rankings are encoded by assigning appropriate weights to program locations, and hence assertions in the generated typing constraint.

Hard Assertions. One way of incorporating ranking criteria is to specify that certain assertions must hold, no matter what. Such assertions are commonly referred to as *hard assertions*. Assertions that the MaxSMT solver is allowed to remove from the assertion set are called *soft*. The compiler may, for instance, annotate all assertions as hard that come from the definitions of functions provided in the standard library. This would encode that all type error sources must be located in user-defined functions. If we apply this criterion to our previous example, we annotate the assertion (4.1) as hard. This eliminates the error source implying that the implementation of not must be modified. Other assertions that might be considered hard are assertions that come from user-provided type annotations in the program.

Weighted Clauses. Our approach supports more sophisticated ranking criteria than distinguishing only between hard and soft assertions. Such criteria are encoded by assigning weights to assertions. Going back to our running example, compilers might favor error sources consisting of smaller program expressions. For instance, each assertion can be assigned a weight equal to the size of the corresponding expression in the abstract syntax tree. This way, **not** \mathbf{x} in our previous example is not reported, as desired.

To demonstrate the power of weighted assertions, we consider another example from the student benchmarks in [54], which served as the motivating example for the type error localization technique presented in [97]:

```
let f(lst:move list):
1
2
       (float*float) list = ...
     let rec loop lst x y dir acc =
3
       if lst = [] then
4
         acc
5
       else
6
         print_string "foo"
7
     in
8
     List.rev
9
       (loop lst 0.0 0.0 0.0 [(0.0,0.0)])
10
```

The standard OCaml compiler reports a type error in the shaded expression on line 10. However, the actual error cause lies in the misuse of the print_string function on line 7, as reported by the authors of [97]. The technique of Zhang and Myers [97] correctly reports the expression on line 7 as the most likely error cause. With the ranking criteria that we introduced so far, our algorithm generates

two error sources that both have minimum weight. These error sources can be interpreted as follows:

- 1. Replace the function print_string on line 7.
- 2. Replace the function loop on line 10.

However, the second error source is not likely the actual cause of the error: using some other function than loop on line 10 would mean that the loop function is not used at all in the program. In fact, the OCaml compiler would generate a warning for the fixed program.

The compiler can thus analyze the produced error sources and find those sources that correspond to the removal of entire functions from the program. Assertions associated with such error sources can then be set as hard. In the running example, the solver will then produce just a single error source, indicating the actual error cause.

This additional ranking criterion can also be encoded directly in the typing constraint without a second round of interaction with the solver. Suppose the program contains a user-defined function \mathbf{f} whose type is described by a type variable α_f . Suppose further that \mathbf{f} is used n times in the program and for i, $1 \leq i \leq n$, let the type variable α_i indicate the type imposed on \mathbf{f} by the context of the *i*-th usage of \mathbf{f} . Then we can add the following additional hard assertion to the typing constraint: $\alpha_1 = \alpha_f \vee \ldots \vee \alpha_n = \alpha_f$. This assertion expresses that \mathbf{f} needs to be used at least once, effectively encoding the additional criterion.

4.2 Problem

We now formally define the problem of computing minimum error sources for a given ranking criterion.

Language and Semantics. For the purposes of type error localization, we assume in this chapter that constants of our language besides Booleans also include integers $n \in Int$ and a constant hole \perp . Please note that the constant \perp here plays a different role than the same symbol in Chapter 3. We explain the role of holes in more detail below. We are also only interested in non-recursive lambda abstractions $\lambda x.e$. This simplification is introduced for presentation purposes without loss of generality. We name the resulting language λ^{\perp} . We also assume standard operational semantics for λ^{\perp} .

As for other Hindley-Milner type systems, type inference is decidable for λ^{\perp} . Expressions therefore do not require explicit type annotations. Typing judgments take the form $\Gamma \vdash e : \tau$. We say that a program $p \in \lambda^{\perp}$ is *well typed* iff there exists a type σ such that $\emptyset \vdash p : \sigma$.

The actual typing rules are shown in Figure 4.2. The rules are standard, with the exception of the rule [HOLE], which we describe in more detail. The value \perp has the polytype $\forall \alpha.\alpha$. Therefore, the typing rule [HOLE] assigns a fresh unconstrained type variable to each usage of \perp . Intuitively, \perp is a placeholder for another program expression. It can be safely used in any context without causing a type error. Changes to hole expressions in an ill-typed program cannot make the program well typed¹.

 $^{^{1}}$ A similar concept was previously used in [54] where a hole in an OCaml program is represented by an expression that raises an exception.

$$\frac{x:\forall \vec{\alpha}. \tau \in \Gamma \quad \vec{\beta} \text{ new}}{\Gamma \vdash x: \tau[\vec{\beta}/\vec{\alpha}]} \text{ [VAR]} \qquad \frac{\Gamma \vdash e_1:\tau_1 \to \tau_2 \quad \Gamma \vdash e_2:\tau_1}{\Gamma \vdash e_1 e_2:\tau_2} \text{ [APP]}$$

$$\frac{b \in Bool}{\Gamma \vdash b: \text{ bool}} \text{ [BOOL]} \qquad \frac{\Gamma.x:\tau_1 \vdash e:\tau_2}{\Gamma \vdash \lambda x.e:\tau_1 \to \tau_2} \text{ [ABS]}$$

$$\frac{\Gamma \vdash e_1: \text{bool}}{\Gamma \vdash e_1:e_2:e_3:\tau} \text{ [COND]} \qquad \frac{n \in Int}{\Gamma \vdash n: \text{ int}} \text{ [INT]}$$

$$\frac{\Gamma \vdash e_1:\tau_1 \quad \Gamma.x:\forall \vec{\alpha}.\tau_1 \vdash e_2:\tau_2 \quad \vec{\alpha} = fv(\tau_1) \setminus fv(\Gamma)}{\Gamma \vdash \text{ let } x = e_1 \text{ in } e_2:\tau_2} \text{ [LET]}$$

$$\frac{\alpha \text{ new}}{\Gamma \vdash \bot : \alpha} \text{ [HOLE]}$$

Figure 4.2: Typing rules for λ^{\perp}

Minimum Error Sources. Without loss of generality, we assume that locations of expressions are implemented as sequences of natural numbers. That is, we call a path $\ell \in \mathbb{N}^*$ in the abstract syntax tree representation of e a *location* of e. Now, let *mask* be the function that maps an expression e and a location $\ell \in Loc(e)$ to the expression obtained from e by replacing $e(\ell)$ with \perp . We extend *mask* to sets of locations in the expected way.

Definition 1 (Error source). Let p be a program. A set of locations $L \subseteq Loc(p)$ is an error source of p if

- 1. mask(p, L) is well typed
- 2. for all strict subsets L' of L, mask(p, L') is not well typed.

A ranking criterion allows the compiler to favor error sources of particular interest by assigning appropriate weights to locations. Formally, a *ranking criterion* is a function R that maps a program p to a partial function $R(p) : Loc(p) \to \mathbb{N}_+$. The locations in $Loc(p) \setminus \text{dom}(R(p))$ are considered hard locations, i.e., R disregards these locations as causes of type errors. We extend R(p) to a set of locations $L \subseteq Loc(p)$ by defining

$$R(p)(L) = \sum_{\ell \in \operatorname{dom}(R(p)) \cap L} R(p)(\ell) \ .$$

Minimum error sources are error sources that minimize the given ranking criterion.

Definition 2 (Minimum error source). Let R be a ranking criterion and p a program. An error source $L \subseteq Loc(p)$ is called minimum error source of p subject to R if for all other error sources L' of p, $R(p)(L) \leq R(p)(L')$.

We are interested in the problem of finding a minimum error source for a given program p subject to a given ranking criterion R, respectively, finding all such minimum error sources.

4.3 Algorithm

In this section, we present our algorithm for computing the minimum error sources. More precisely, we show how the problem of finding minimum type error sources can be reduced to the weighted MaxSMT problem that can in turn be solved offthe-shelf using automated theorem provers. We first formally define the weighted MaxSMT problem and then explain the actual reduction.

Weighted MaxSMT. The MaxSAT problem takes as input a finite set of propositional clauses C and finds an assignment that maximizes the number of clauses K that are simultaneously satisfied [56]. MaxSAT can alternatively be viewed as finding the largest subset \mathcal{C}' of clauses \mathcal{C} such that \mathcal{C}' is satisfiable and \mathcal{C}' is a maximum satisfiable subset, $|\mathcal{C}'| = K$. Partial MaxSAT partitions \mathcal{C} into hard and soft clauses. The hard clauses \mathcal{C}_H are assumed to hold and the goal is to find a maximizing subset \mathcal{C}' of the soft clauses such that $\mathcal{C}' \cup \mathcal{C}_H$ is satisfiable. Weighted Partial MaxSAT, for simplicity referred to as only weighted MaxSAT, adds an integer weight $w_i = w(C_i)$ to each soft clause C_i and asks for a satisfiable subset \mathcal{C}' that maximizes the weighted score:

$$\sum_{C_i \in \mathcal{C}'} w_i \text{ subject to } \mathcal{C}_H \cup \mathcal{C}' \text{ is satisfiable.}$$
(4.6)

The MaxSMT problem generalizes MaxSAT from working over propositional clauses to a set of assertion formulas \mathcal{A} where each assertion belongs to a fixed first-order theory \mathcal{T} . Most concepts directly generalize from MaxSAT to MaxSMT: satisfiability is replaced by Satisfiability Modulo Theories [6], partial MaxSMT has hard and soft assertions (\mathcal{A}_H and \mathcal{A}_S), weighted partial MaxSMT assigns an integer weight to each soft assertion. We represent weighted MaxSMT instances as tuples ($w, \mathcal{A}_H, \mathcal{A}_S$) where w is the weight function assigning the weights to the soft assertions.

Theory of Inductive Data Types. Our reduction to the weighted MaxSMT problem generates a typing constraint from the given input program. This constraint is satisfiable iff the input program is well typed. The constraint is then passed to the MaxSMT solver, which computes the minimum error sources. The generated typing constraint hence needs to be expressed in terms of assertions that are interpreted in an appropriate first order theory. We use the theory of inductive data types [7] for this purpose. In this theory, we can define our own inductive

data types and then state equality constraints over the terms of that data type. For our encoding, we define an inductive data type **Types** that represents the set of all ground monotypes of λ^{\perp} :

$$t \in \mathsf{Types} ::= \mathsf{int} \mid \mathsf{bool} \mid \mathsf{fun}(t,t)$$

Here, the term constructor fun is used to encode the ground function types.

The terms in **Types** are interpreted in the models of the theory of inductive data types. A model of this theory is a first-order structure that interprets the type constructors such that the following axioms are satisfied:

$$\begin{aligned} & \mathsf{int} \neq \mathsf{bool} \\ & \forall \alpha, \beta \in \mathsf{Types.} \ \mathsf{fun}(\alpha, \beta) \neq \mathsf{int} \land \mathsf{fun}(\alpha, \beta) \neq \mathsf{bool} \\ & \forall \alpha, \beta, \gamma, \delta \in \mathsf{Types.} \ \mathsf{fun}(\alpha, \beta) \!=\! \mathsf{fun}(\gamma, \delta) \Rightarrow \alpha \!=\! \gamma \land \beta \!=\! \delta \end{aligned}$$

That is, the term constructor fun must be interpreted by an injective function, and the interpretation of the terms int and bool is distinct from the interpretation of all other terms. Hence, the axioms exactly encode the equality relation on ground monotypes of λ^{\perp} . For the type systems of actual languages such as OCaml, we extend Types with additional type constructors, e.g., to encode user-defined algebraic data types in OCaml programs.

Overview of the Reduction. Next, we describe the actual reduction to weighted MaxSMT. In the following, let p be the program under consideration and let R be the ranking criterion of interest.

At the heart of our reduction is a constraint generation procedure that traverses p and generates a set of assertions in our theory of Types. These assertions encode
the typing rules shown in Figure 4.2. The constraint generation procedure can produce multiple assertions associated with a single location ℓ of p. Suppose R considers a location ℓ to be soft. That is, while searching for minimum error sources the solver must consider two possibilities. If ℓ is in the minimum error source, then none of the assertions generated from the expression $p(\ell)$ need to be satisfied. If on the other hand ℓ is not in the minimum error source, then the type of the expression $p(\ell)$ must be consistent with its context. That is, all the assertions directly associated with ℓ must be satisfied simultaneously. However, some of the assertions generated from subexpressions of $p(\ell)$ may be dropped because the corresponding locations may still be in the minimum error source. Thus, to properly encode the problem of finding minimum error sources, we need to link the assertions of location ℓ together so that the solver considers them as a single logical unit. Also, we need to capture the relationship between the locations according to the structure of p. For this purpose, we associate a propositional variable T_{ℓ} with every location ℓ . By setting T_{ℓ} to true, the solver decides that T_{ℓ} is not in the minimum error source. That is, each assertion associated with a location ℓ_o takes the following form:

$$T_{\ell_n} \Rightarrow \dots \Rightarrow T_{\ell_1} \Rightarrow T_{\ell_0} \Rightarrow t_1 = t_2$$
 (4.7)

Here, ℓ_1, \ldots, ℓ_n are the locations that are reached along the path to ℓ_0 in p (i.e., all proper prefixes of ℓ_0). The terms t_1 and t_2 are terms in our theory of Types. These terms may include logical variables that need to be unified by the MaxSMT solver. Note that the assertion (4.7) ensures that if any of the T_{ℓ_i} is set to false (i.e., ℓ_i is in the minimum error source), then all the assertions that depend on ℓ_0 are immediately satisfied. In order to simplify the presentation in § 4.1, we have omitted the additional propositional variables in our discussion of the examples.

The assertions that encode the typing rules are considered hard. To reduce the problem of finding minimum error sources to weighted MaxSMT, we add to these hard assertions an additional set of clauses, each of which consists of one of the propositional variables T_{ℓ} . For each ℓ , if $R(p)(\ell)$ is defined, the clause T_{ℓ} is soft with associated weight $R(p)(\ell)$. Otherwise, T_{ℓ} is hard. Each solution to this weighted MaxSMT instance then yields a minimum error source by taking all locations ℓ whose clause T_{ℓ} is not in the solution set.

We now explain these steps in more detail, starting with the generation of the assertions that encode the typing rules.

4.3.1 Assertion Generation

We build on existing work for constrained-based type checking [87, 1, 71] and adapt it for our purposes.

We formalize the assertion generation in terms of a constraint typing relation $\mathcal{A}, \Gamma \vdash p : \alpha$. Here, \mathcal{A} is a set of typing assertions of the form (4.7). The rules that define the constraint typing relation are given in Figure 4.3. They are defined recursively on the structure of expressions. To relate the generated typing assertions to p's locations, we assume that every expression is annotated with its location ℓ in p. Note that for a set of assertions \mathcal{A} , we write $T_{\ell} \Rightarrow \mathcal{A}$ to mean $\{T_{\ell} \Rightarrow \mathcal{A} \mid A \in \mathcal{A}\}.$

We focus on the rules [A-LET] and [A-VAR] as they are the most complex ones. The [A-LET] rule handles let bindings let $x = e_1$ in e_2 by first computing a set of typing assertions \mathcal{A}_1 for e_1 . The assertions in \mathcal{A}_1 encode all typing assertions imposed by e_1 under the current environment Γ . In particular, they capture the

$$\frac{\mathcal{A}, \Gamma.x: \alpha \vdash e: \beta \quad \gamma \text{ new}}{T_{\ell} \Rightarrow (\{\gamma = \mathsf{fun}(\alpha, \beta)\} \cup \mathcal{A}), \Gamma \vdash (\lambda x.e)^{\ell}: \gamma} \text{ [A-ABS]}$$
$$\frac{\mathcal{A}_1, \Gamma \vdash e_1: \alpha \quad \mathcal{A}_2, \Gamma \vdash e_2: \beta \quad \gamma \text{ new}}{T_{\ell} \Rightarrow (\{\alpha = \mathsf{fun}(\beta, \gamma)\} \cup \mathcal{A}_1 \cup \mathcal{A}_2), \Gamma \vdash (e_1 \ e_2)^{\ell}: \gamma} \text{ [A-APP]}$$

$$\frac{\mathcal{A}_1, \Gamma \vdash e_1 : \alpha \quad \mathcal{A}_2, \Gamma \vdash e_2 : \beta \quad \mathcal{A}_3, \Gamma \vdash e_3 : \gamma \quad \delta \text{ new}}{\mathcal{A} = \{ (T_{\ell_1} \Rightarrow \alpha = \mathsf{bool}), (T_{\ell_2} \Rightarrow \beta = \delta), (T_{\ell_3} \Rightarrow \gamma = \delta) \} \cup \mathcal{A}_1 \cup \mathcal{A}_2 \cup \mathcal{A}_3} \quad [A-\text{COND}]}{T_\ell \Rightarrow \mathcal{A}, \Gamma \vdash (e_1^{\ell_1} ? e_2^{\ell_2} : e_3^{\ell_3})^\ell : \delta} \quad [A-\text{COND}]$$

$$\frac{\alpha \text{ new}}{\emptyset, \Gamma \vdash \bot : \alpha} \text{ [A-HOLE]} \qquad \qquad \frac{x : \forall \vec{\alpha}. (\mathcal{A} \Rightarrow \alpha) \in \Gamma \quad \vec{\beta}, \gamma \text{ new}}{T_{\ell} \Rightarrow (\{\gamma = \alpha[\vec{\beta}/\vec{\alpha}]\} \cup \mathcal{A}[\vec{\beta}/\vec{\alpha}]), \Gamma \vdash x^{\ell} : \gamma} \text{ [A-VAR]}$$

$$\frac{b \in Bool \quad \alpha \text{ new}}{\{T_{\ell} \Rightarrow \alpha = \mathsf{bool}\}, \Gamma \vdash b^{\ell} : \alpha} \text{ [A-BOOL]} \qquad \frac{n \in Int \quad \alpha \text{ new}}{\{T_{\ell} \Rightarrow \alpha = \mathsf{int}\}, \Gamma \vdash n^{\ell} : \alpha} \text{ [A-INT]}$$

$$\frac{\mathcal{A}_1, \Gamma \vdash e_1 : \alpha_1 \quad \mathcal{A}_2, \Gamma.x : \forall \vec{\alpha}.(\mathcal{A}_1 \Rightarrow \alpha_1) \vdash e_2 : \alpha_2 \quad \vec{\alpha} = fv(\alpha) \setminus fv(\Gamma) \quad \vec{\beta}, \gamma \text{ new}}{T_\ell \Rightarrow (\{\gamma = \alpha_2\} \cup \mathcal{A}_1[\vec{\beta}/\vec{\alpha}] \cup \mathcal{A}_2), \Gamma \vdash (\texttt{let } x = e_1 \texttt{ in } e_2)^\ell : \gamma}$$
[A-LET]

Figure 4.3: Rules defining the constraint typing relation for λ^{\perp}

assertions on the type of e_1 itself, which is described by a fresh type variable α_1 . To compute the typing assertions \mathcal{A}_2 for e_2 , the variable x is added to the typing environment Γ . The type of x is described by a typing schema of the form:

$$\forall \vec{\alpha}. (\mathcal{A}_1 \Rightarrow \alpha_1)$$

The typing schema is needed to properly handle polymorphism. It remembers the set of assertions \mathcal{A}_1 along with the type variable α_1 that represents the type of e_1 and x. Note that the schema quantifies over all type variables $\vec{\alpha}$ that have been freshly introduced when analyzing e_1 (including α_1). Whenever x is used inside the body e_2 of the let binding, the [A-VAR] rule instantiates \mathcal{A}_1 by replacing the type variables $\vec{\alpha}$ with fresh type variables $\vec{\beta}$. The instantiated copy is then added to the other generated assertions. The fresh instances of \mathcal{A}_1 ensure that each usage of x in e_2 is consistent with the typing assertions imposed by e_1 .

The following lemma states the correctness of the constraint typing relation.

Lemma 21. Let p be a program, $L \subseteq Loc(p)$, \mathcal{A} a set of typing assertions, and α a type variable such that $\mathcal{A}, \emptyset \vdash p : \alpha$. Then mask(L, p) is well typed iff $\mathcal{A} \cup \{T_{\ell} \mid \ell \notin L\}$ is satisfiable in the theory of Types.

The proof of Lemma 21 closely follows that of [87, Lemma 2], modulo reasoning about the auxiliary propositional variables.

4.3.2 Reduction to Weighted Partial MaxSMT

To compute minimum error sources, we generate a weighted MaxSMT instance $I(p, R) = (w, \mathcal{A}_H, \mathcal{A}_S)$ as follows. Let \mathcal{A} be a set of assertions such that $\mathcal{A}, \emptyset \vdash p : \alpha$ for some type variable α . Then define:

$$\mathcal{A}_{H} = \mathcal{A} \cup \{ T_{\ell} \mid \ell \notin \operatorname{dom}(R(p)) \}$$
$$\mathcal{A}_{S} = \{ T_{\ell} \mid \ell \in \operatorname{dom}(R(p)) \}$$
$$w(T_{\ell}) = R(p)(\ell), \text{ for all } T_{\ell} \in \mathcal{A}_{S}$$

Let \mathcal{S} be a solution of I(p, R). Then define $E(\mathcal{S}) = \{ \ell \in Loc(p) \mid T_{\ell} \notin \mathcal{S} \}$. The following theorem states the correctness of our reduction.

Theorem 5. Let p be a program and R a ranking criterion. Then $L \subseteq Loc(p)$ is a minimum error source of p subject to R iff there exists a solution S of I(p, R)such that L = E(S).



Figure 4.4: Labeled abstract syntax tree for the program p

We conclude the presentation of our algorithm with an example that demonstrates how our approach deals with polymorphic functions.

Example 4. Let p be the following ill-typed program:

let
$$f = \lambda x y. y x$$
 in $f = 10$

The most general type that can be inferred for f from its defining expression is

$$f: \forall \alpha \beta. \alpha \to (\alpha \to \beta) \to \beta$$
.

However, the type of 0 is int and not a function type. Hence, applying 0 as second argument to f in the body of the let violates the typing rules.

Figure 4.4 shows the abstract syntax tree of the program p with each node labeled by its location. Applying the constraint generation rules from Figure 4.3 then yields the following set of assertions \mathcal{A} :

$$\begin{split} T_{\ell_0} \Rightarrow & \{\alpha_0 = \alpha_6, \ C_f(\alpha_1'', \alpha_2'', \alpha_3'', \alpha_4'', \alpha_5'', \beta_1'', \beta_2''), \\ & T_{\ell_6} \Rightarrow \{\alpha_7 = \mathsf{fun}(\alpha_8, \alpha_6), \\ & T_{\ell_7} \Rightarrow \{\alpha_9 = \mathsf{fun}(\alpha_{10}, \alpha_7), \\ & T_{\ell_8} \Rightarrow \alpha_8 = \mathsf{int}, \\ & T_{\ell_9} \Rightarrow \{\alpha_9 = \alpha_1'', \\ & C_f(\alpha_1', \alpha_2', \alpha_3', \alpha_4', \alpha_5', \beta_1', \beta_2')\} \\ & \}, \\ & T_{\ell_{10}} \Rightarrow \alpha_{10} = \mathsf{int}\} \} \end{split}$$

where

$$C_{f}(\alpha_{1}, \alpha_{2}, \alpha_{3}, \alpha_{4}, \alpha_{5}, \beta_{1}, \beta_{2}) \equiv$$

$$T_{\ell_{1}} \Rightarrow \{ \alpha_{1} = \operatorname{fun}(\beta_{1}, \alpha_{2}),$$

$$T_{\ell_{2}} \Rightarrow \{ \alpha_{2} = \operatorname{fun}(\beta_{2}, \alpha_{3}),$$

$$T_{\ell_{3}} \Rightarrow \{ \alpha_{4} = \operatorname{fun}(\alpha_{5}, \alpha_{3}),$$

$$T_{\ell_{4}} \Rightarrow \alpha_{4} = \beta_{2},$$

$$T_{\ell_{5}} \Rightarrow \alpha_{5} = \beta_{1} \} \}$$

Note that we introduced a predicate C_f to serve as a shorthand for the instantiated assertions that are associated with the bound variable f.

The assertions in \mathcal{A} are satisfiable in Types if one of the following location variables is assigned false: T_{ℓ_4} , T_{ℓ_8} , T_{ℓ_9} , or any other T_{ℓ_i} where ℓ_i is on the path to one of the locations ℓ_4 , ℓ_8 , or ℓ_9 .

Now consider the ranking criterion R that defines ℓ_9 hard and all other locations ℓ_i soft with rank $R(p)(\ell_i) = n_i$, where n_i is the number of nodes in the subtree of the AST whose root ℓ_i identifies. Solving the weighted MaxSMT instance I(p, R) yields two minimum error sources, each of which contains one of the locations ℓ_4 and ℓ_8 .

4.4 Implementation and Evaluation

We have implemented our weighted MaxSMT based algorithm for computing minimum type error sources for the Caml subset of the OCaml language and evaluated it on the OCaml benchmark suite from [54]. The implementation of our algorithm was able to find a minimum error source for 98% of the benchmarks in a reasonable time using a typical ranking criterion.

4.4.1 Implementation

Our implementation bundles together the EasyOCaml [24] tool and the MaxSMT solver νZ [11, 10]. The νZ solver is available as a branch of the SMT solver Z3 [21]. We use EasyOCaml for generating typing constraints for OCaml programs. Once we convert the constraints to the weighted MaxSMT instances, we use Z3's weighted MaxRes [66] algorithm to compute a minimum error source.

Assertion Generation. EasyOCaml implements the type error localization approach described in [33] for a subset of the OCaml language. Similar to our approach, EasyOCaml uses constraint solving for error localization. We tapped into the corresponding code of EasyOCaml and modified it to produce the weighted MaxSMT instances from the input program. The implementation begins by running EasyOCaml on the input program. EasyOCaml produces typing assertions, which we annotate with the locations of the corresponding program expressions in the input source code. These assertions only encode the typing relation. In addition, we output the structure of the input program in terms of its abstract syntax tree edges where the nodes are the locations associated with the typing assertions. This information is sufficient to generate the final typing constraint as described

in § 4.3.2. In our evaluation, we used a fixed ranking criterion that is defined as follows: (1) all assertions that come from expressions in external libraries are set as hard; (2) all assertions that come from user-provided type annotations are set as hard; and (3) all remaining assertions have a weight equal to the size of the corresponding expression in the abstract syntax tree. We encode the generated assertions in an extension of the SMT-LIB 2 language [8] to handle the theory of inductive data types.

Solving the Weighted MaxSMT Instances. We compute the weighted partial MaxSMT solution for the encoded typing constraints by using Z3's weighted partial MaxSMT facilities. In particular, we configure the solver to use the MaxRes [66] algorithm for solving the weighed partial MaxSMT problem.

4.4.2 Evaluation

We now discuss in more details the effectiveness of minimum error sources as well as the efficiency of computing them using our implementation.

Quality of Computed Minimum Error Sources. In our first experiment, we assessed how good the produced minimum error sources are at pinpointing the true error source. To this end, we chose 20 programs from the benchmark suite at random, identified the true error source in each program, and compared it against the first minimum error source that was computed by our implementation. As mentioned earlier, we used the ranking criterion that favors error sources of smaller code size and that excludes possible errors sources coming from external functions and type annotations. To identify the true error source we used additional information that was provided by the benchmark suite. Namely, the benchmark suite does not consist of individual programs. Instead, it consists of sequences of modified versions of programs. Each such program sequence was recorded in a single session of interactions between a student and the OCaml compiler. By comparing subsequent program versions within one sequence, we can identify the changes that the student made to fix the reported type errors and thereby infer the true error sources.

We classified the result of the comparison as either "hit", "close", or "miss". Here, "hit" means that the computed minimum error source exactly matches the true error source, and "miss" means that there is no relationship between the two. A "close" result means that the locations in the computed minimum error source were close enough to the true error locations so that a programmer could easily understand and fix the type error. For example, if the minimum error source reported the function in a function application instead of the argument of the application (or vice versa), then we considered the result to be close to the true error source.

We then repeated the same experiment but this time using OCaml's type checker instead of our implementation. For each program, we recorded the result of the two experiments. Figure 4.5 shows the number of obtained result pairs, aggregated over the 20 benchmark programs. As can be seen, on the randomly chosen programs, our approach identifies the true error source more often than the OCaml type checker, even though we were using a rather simplistic ranking criterion. Specifically, our approach missed the true error source in only three programs, whereas OCaml did so in six programs. Despite the subjective nature of true error sources, we believe that this experiment shows the potential of our

AST size criterion	OCaml	# of outcomes
hit	miss	3
hit	close	2
hit	hit	4
close	miss	1
close	close	6
close	hit	1
miss	miss	2
miss	close	1
miss	hit	0

Figure 4.5: Quality of AST ranking criterion compared to OCaml's type checker in pinpointing the true error source in 20 benchmark programs

approach to providing helpful error reports to the user.

Computing minimum error sources. The last experiment we performed had the goal of measuring the time spent computing the minimum error sources. We broke the student benchmark suite consisting of 356 programs into 8 groups according to their size in the number of lines of code. The first group includes programs consisting of 0 and 50 lines of code, the second group includes programs of size 50 to 100, and so on as shown in the first column of Figure 4.6. The numbers in parenthesis indicate the number of programs in each group. The figure shows statistics about the execution time for computing a minimum error source. We show the minimum, average, and maximum execution times for each program family. As it can be seen, the execution times are reasonably small. However, for some programs the framework spends around a dozen seconds computing the solution.

The slowdown in our implementation mostly occurs in the cases where a huge number of constraints has been generated. This is confirmed by statistics shown in Figure 4.7. As it can be seen, even for small to moderate programs the number of



Figure 4.6: Maximum, average, and minimum execution times for computing a minimum error source



Figure 4.7: Maximum, average, and minimum number of generated assertions (in thousands) for computing a minimum error source

generated assertions can reach more than ten thousand. In general, the number of assertions does not grow linearly with the size of the program. In fact, our tool can generate an exponential number of assertions for certain programs. As we discuss next, this phenomena is tied to an inherent property of polymorphic typing.

More experiments and accompanying statistics can be found in the original paper [72].

4.5 Taming The Exponential Explosion

As observed in the previous section, the number of typing constraints can grow exponentially in the size of the program. This is because the constraints associated with polymorphic functions are duplicated each time these functions are used, as illustrated in the discussion of Example 4.

Exponential Explosion Problem. Let us again briefly provide the intuition behind copying constraints for polymorphic functions. Consider the typing constraints for the OCaml function let f(a, b, c) = (a, c) that simply takes a first and last element of a triple and returns them as a pair.

$$\alpha_f = \mathsf{fun}(\alpha_i, \alpha_o)$$
$$\alpha_i = \mathsf{triple}(\alpha_a, \alpha_b, \alpha_c)$$
$$\alpha_o = \mathsf{pair}(\alpha_a, \alpha_c)$$

The above constraints state that the type of \mathbf{f} , represented by the type variable α_f , is a function type that accepts some triple and returns a pair whose constituent types are equal to the type of the first and last component of that triple. When a polymorphic function, such as \mathbf{f} , is called in the program, the associated set of typing constraints needs to be *instantiated* and the new copy has to be added to the whole set of typing constraints. Instantiation of typing constraints involves

copying the constraints and replacing free type variables in the copy with fresh type variables. If the constraints of polymorphic function were not freshly instantiated for each usage of the function, the same type variable would be constrained by the context of each usage, potentially resulting in a spurious type error. Introducing a fresh copy of typing constraints for functions is a standard approach for dealing with polymorphism in purely constraint based type inference and checking [74].

The exponential explosion in the constraint size does not seem to be avoidable. The type inference problem for polymorphic type systems is known to be EXPTIME-complete [57, 49]. However, in practice, compilers successfully avoid the explosion by computing the principal type [20] of each polymorphic function, such as the one shown in Example 4, and then instantiating a fresh copy of this type for each usage. The resulting constraints are much smaller in practice. Since these constraints are equisatisfiable with the original constraints, the resulting algorithm is a decision procedure for the type checking problem [20]. Unfortunately, this technique cannot be applied immediately to the optimization problem of type error localization. If the minimum cost error source is located inside of a polymorphic function, then abstracting the constraints of that function by its principle type will hide this error source. Thus, this approach can yield incorrect results.

Iterative Expansion. In this section, we propose an improved reduction to the MaxSMT problem that abstracts polymorphic functions by principal types and still guarantees the optimum solution. Our reduction is based on the optimistic assumption that type error sources only involve few polymorphic functions, even for large programs. Our new reduction to MaxSMT relies on principal type abstraction that is done in such a way that all potential error sources involving the definition

of an abstracted function are represented by a single error source whose cost is smaller or equal to the cost of all these potential error sources. The algorithm then iteratively computes minimum error sources for abstracted constraints. If an error source involves a usage of a polymorphic function, the corresponding instantiations of the principal type of that function are expanded to the actual typing constraints. Usages of polymorphic functions that are exposed by the new constraints are expanded if they are relevant for the minimum error source in the next iteration. The algorithm eventually terminates when the computed minimum error source no longer involves any usages of abstracted polymorphic functions. Such error sources are guaranteed to have minimum cost for the fully expanded constraints, even if the final constraint is not yet fully expanded.

Our core technical contribution is a refinement of the typing relation introduced in § 4.3.1. The novelty of this new typing relation is the ability to specify a set of variable usage locations whose typing constraints are abstracted as the principal type of the variable. We then describe an algorithm that iteratively uses this typing relation to find a minimum error source while expanding only those principal type usages that are relevant for the minimum source.

4.5.1 Assertion Generation

For the rest of this section, we assume a fixed program $p \in \lambda^{\perp}$. We first formally introduce the notion of principal types.

Principal Types. Standard type inference implementations handle expressions of the form let $x = e_1$ in e_2 by computing the *principal* type of e_1 , binding x to the principal type σ_p in the environment $\Gamma . x : \sigma_p$, and proceeding to perform type

$$\begin{split} \frac{\Pi.x:\alpha,\Gamma.x:\alpha\vdash_{\mathfrak{L}}e:\beta\mid\mathcal{A}\qquad\gamma \text{ new}}{\Pi,\Gamma\vdash_{\mathfrak{L}}(\lambda x.e)^{\ell}:\gamma\mid\{T_{\ell}\Rightarrow(\{\gamma=\text{fun}(\alpha,\beta)\}\cup\mathcal{A})\}} \text{ [A-ABS]}\\ \frac{\Pi,\Gamma\vdash_{\mathfrak{L}}e:\alpha\mid\alpha\mid\mathcal{A}_{1}\qquad\Pi,\Gamma\vdash_{\mathfrak{L}}e:q:\beta\mid\mathcal{A}_{2}\qquad\gamma \text{ new}}{\Pi,\Gamma\vdash_{\mathfrak{L}}e(\{e:1e:2)^{\ell}:\gamma\mid\{T_{\ell}\Rightarrow(\{\alpha=\text{fun}(\beta,\gamma)\}\cup\mathcal{A}_{1}\cup\mathcal{A}_{2})\}} \text{ [A-APP]}\\ \Pi,\Gamma\vdash_{\mathfrak{L}}e:q:a:\alpha\mid\mathcal{A}_{1}\qquad\Pi,\Gamma\vdash_{\mathfrak{L}}e:q:\alpha_{2}\mid\mathcal{A}_{2}\qquad\Pi,\Gamma\vdash_{\mathfrak{L}}e:a:\alpha_{3}\mid\mathcal{A}_{3}\qquad\gamma \text{ new}\\ \mathcal{A}_{4}=\{(T_{\ell_{1}}\Rightarrow\alpha_{1}=\text{bool}),(T_{\ell_{2}}\Rightarrow\alpha_{2}=\gamma),(T_{\ell_{3}}\Rightarrow\alpha_{3}=\gamma)\}\\ \Pi,\Gamma\vdash_{\mathfrak{L}}(e_{1}^{\ell_{1}}?e_{2}^{\ell_{2}}:e_{3}^{\ell_{3}})^{\ell}:\gamma\mid\{T_{\ell}\leftrightarrow(\mathcal{A}_{1}\cup\mathcal{A}_{2}\cup\mathcal{A}_{3}\cup\mathcal{A}_{4})\}\\ \hline\frac{\alpha \text{ new}}{\Pi,\Gamma\vdash_{\mathfrak{L}}e:a:\alpha\mid\emptyset} \text{ [A-HOLE]}\\ \frac{b\in Bool}{\Pi,\Gamma\vdash_{\mathfrak{L}}e!a:\alpha\mid\{\mathcal{A}_{\ell}\Rightarrow\alpha=\text{bool}\}} \text{ [A-BOOL]}\\ \hline\frac{n\in Int}{\Pi,\Gamma\vdash_{\mathfrak{L}}n^{\ell}:\alpha\mid\{T_{\ell}\Rightarrow\alpha=\text{int}\}} \text{ [A-INT]}\\ \\\frac{\ell\in\mathfrak{L}}{\Pi,\Gamma\vdash_{\mathfrak{L}}n^{\ell}:\gamma\mid\{T_{\ell}\Rightarrow(\{\gamma=\alpha[\vec{\beta}/\vec{\alpha}]\}\cup\mathcal{A}[\vec{\beta}/\vec{\alpha}])\}} \text{ [A-VAR-EXP]}\\ \\\frac{\ell\notin\mathfrak{L}}{\Pi,\Gamma\vdash_{\mathfrak{L}}x^{\ell}:\gamma\mid\{T_{\ell}\Rightarrow(\{\gamma=\alpha[\vec{\beta}/\vec{\alpha}]\}\cup\mathcal{A}[\vec{\beta}/\vec{\alpha}])\}} \text{ [A-VAR-PRIN]}\\ \ell_{1}\in\mathfrak{L} \end{split}$$

$$\begin{split} \Pi, \Gamma \vdash_{\mathfrak{L}} e_{1} &: \alpha_{1} \mid \mathcal{A}_{1} \qquad \vec{\alpha} = fv(\mathcal{A}_{1}) \setminus fv(\Gamma) \qquad \tau_{exp} = \forall \vec{\alpha}.(\mathcal{A}_{1} \Rrightarrow \alpha_{1}) \\ \Pi, \Gamma.x &: \tau_{exp} \vdash_{\mathfrak{L}} e_{2} &: \alpha_{2} \mid \mathcal{A}_{2} \qquad \vec{\beta}, \gamma \text{ new} \\ \hline \Pi, \Gamma \vdash_{\mathfrak{L}} (\texttt{let } x = e_{1}^{\ell_{1}} \texttt{ in } e_{2})^{\ell} &: \gamma \mid \{T_{\ell} \Rightarrow (\{\gamma = \alpha_{2}\} \cup \mathcal{A}_{1}[\vec{\beta}/\vec{\alpha}] \cup \mathcal{A}_{2})\} \end{split}$$
[A-LET-EXP]

$$\begin{split} \ell_{1} \notin \mathfrak{L} \\ \rho(\Pi, e_{1}) &= \forall \vec{\delta}.\tau_{p} \qquad \alpha \text{ new} \qquad \tau_{prin} = \forall \alpha, \vec{\delta}.(\{P_{\ell_{1}} \Rightarrow \alpha = \tau_{p}\} \Rightarrow \alpha) \\ \Pi, \Gamma \vdash_{\mathfrak{L}} e_{1} : \alpha_{1} \mid \mathcal{A}_{1} \qquad \vec{\alpha} = fv(\mathcal{A}_{1}) \setminus fv(\Gamma) \qquad \tau_{exp} = \forall \vec{\alpha}.(\mathcal{A}_{1} \Rightarrow \alpha_{1}) \\ \frac{\Pi.x : \tau_{prin}, \ \Gamma.x : \tau_{exp} \vdash_{\mathfrak{L}} e_{2} : \alpha_{2} \mid \mathcal{A}_{2} \qquad \vec{\beta}, \gamma \text{ new} \\ \overline{\Pi, \Gamma \vdash_{\mathfrak{L}} (\operatorname{let} x = e_{1}^{\ell_{1}} \text{ in } e_{2})^{\ell} : \gamma \mid \{T_{\ell} \Rightarrow (\{\gamma = \alpha_{2}\} \cup \mathcal{A}_{1}[\vec{\beta}/\vec{\alpha}] \cup \mathcal{A}_{2})\}} \quad [A-\operatorname{Let-Prin}] \end{split}$$

Figure 4.8: Rules defining the refined constraint typing relation for λ^\perp

inference on e_1 [20]. Given an environment Γ , the type σ_p is the principal type for e if $\Gamma \vdash e : \sigma_p$ and for any other σ such that $\Gamma \vdash e : \sigma$ then σ is a generic instance of σ_p . Note that a principal type is unique, subject to e and Γ , up to the renaming of bound type variables in σ_p . We define ρ to be a partial function accepting an expression e and a typing environment Γ where $\rho(\Gamma, e)$ returns a principal type of e subject to Γ . If e is not typeable in Γ , then $(\Gamma, e) \notin \operatorname{dom}(\rho)$. We assume that the principal type function ρ is precomputed for p. We do not provide a detailed algorithms for computing ρ since it is well-known from the literature. For instance, the ρ function can be implemented using the classical W algorithm [20].

The main idea behind our improved algorithm is to iteratively discover which principal type usages must be expanded to compute a minimum error source. The technical core of the algorithm is a new typing relation that produces typing constraints subject to a set of locations where principal type usages must be expanded.

As before, we use \mathcal{A} to denote a set of logical assertions in the signature of Types that represent typing constraints. Henceforth, when we refer to types we mean terms over Types. Expanded locations are a set of locations \mathfrak{L} such that $\mathfrak{L} \subseteq Loc(p)$. Intuitively, this is a set of locations corresponding to usages of let variables xwhere the typing of x in the current iteration of the algorithm is expanded into the corresponding typing constraints. Those locations of usages of x that are not expanded will treat x using its principal type. We also introduce a set of locations whose usages must be expanded \mathfrak{L}_0 . We will always assume $\mathfrak{L}_0 \subseteq \mathfrak{L}$. Formally, \mathfrak{L}_0 is the set of all program locations in p except the locations of well-typed let variables and their usages. This definition enforces that usages of variables that have no principal type are always expanded. In summary, $\mathfrak{L}_0 \subseteq \mathfrak{L} \subseteq Loc(p)$.

We define a typing relation $\vdash_{\mathfrak{L}}$ over $(\Pi, \Gamma, e, \alpha, \mathcal{A})$ which is parameterized by

 \mathfrak{L} . The relation is given by judgments of the form:

$$\Pi, \Gamma \vdash_{\mathfrak{L}} e : \alpha \mid \mathcal{A}.$$

Intuitively, the relation holds iff expression e in p has type α under typing environment Γ if we solve the constraints \mathcal{A} for α . (We make this statement formally precise later.) The relation depends on \mathfrak{L} , which controls whether a usage of a let variable is typed by the principal type of the let definition or the expanded typing constraints of that definition.

For technical reasons, the principal types are computed in tandem with the expanded typing constraints. This is because both the expanded constraints and the principal types may refer to type variables that are bound in the environment, and we have to ensure that both agree on these variables. We therefore keep track of two separate typing environments:

- the environment Π binds let variables to the principal types of their defining expressions if the principal type exists with respect to Π, and
- the typing environment Γ binds let variables to their expanded typing constraints (modulo £).

The typing relation ensures that the two environments are kept synchronized. To properly handle polymorphism, the bindings in Γ are represented by typing schemas:

$$x: \forall \vec{\alpha}. (\mathcal{A} \Rrightarrow \alpha)$$

The schema states that x has type α if we solve the typing constraints \mathcal{A} for the variables $\vec{\alpha}$. To simplify the presentation, we also represent bindings in Π as type schemas. Note that we can represent an arbitrary type t by the schema $\forall \alpha.(\{\alpha = t\} \Rightarrow \alpha)$ where $\alpha \notin fv(t)$. The symbol \Rightarrow is used here to suggest, but keep syntactically separate, the notion of logical implication \Rightarrow that is implicitly present in the schema.

The typing relation $\Pi, \Gamma \vdash_{\mathfrak{L}} e : \alpha \mid \mathcal{A}$ is defined in Figure 4.8. It can be seen as a constraint generation procedure that goes over an expression e at location ℓ and generates a set of typing constraints \mathcal{A} . Again, we associate with each location ℓ a propositional variable T_{ℓ} for the same purpose of identifying minimum error sources as defined in § 4.2.

The rules A-LET-PRIN and A-LET-EXP govern the computation and binding of typing constraints and principal types for let definitions $(\text{let } x = e_1^{\ell_1} \text{ in } e_2)^{\ell}$. If e_1 has no principal type under the current environment Π , then $\ell_1 \in \mathfrak{L}$ by the assumption that $\mathfrak{L}_0 \subseteq \mathfrak{L}$. Thus, when rule A-LET-PRIN applies, $\rho(\Pi, e_1)$ is defined. The rule then binds x in Π to the principal type and binds x in Γ to the expanded typing constraints obtained from e_1 .

The [A-LET-PRIN] rule binds x in both Π and Γ as it is possible that in the current iteration some usages of x need to be typed with principal types and some with expanded constraints. For instance, our algorithm can expand usages of a function, say f, in the first iteration, and then expand all usages of, say g, in the next iteration. If g's defining expression in turn contains calls to f, those calls will be typed with principal types. This is done because there may exist a minimum error source that does not require that the calls to f in g are expanded.

After extending the typing environments, the rule recurses to compute the typing constraints for the body e_2 with the extended environments. Note that the rule introduces an auxiliary propositional variable P_{ℓ_1} that guards all the typing constraints of the principal type before x is bound in Π . This step is crucial for the

correctness of the algorithm. We refer to the variables as *principal type correctness* variables. That is, if P_{ℓ_1} is **true** then this means that the definition of the variable bound at ℓ_1 is not involved in the minimum error source and the principal type safely abstracts the associated unexpanded typing constraints.

The rule A-LET-EXP applies whenever $\ell_1 \in \mathfrak{L}$. The rule is almost identical to the A-LET-PRIN rule, except that it does not bind x in Π to τ_{prin} (the principal type). This will have the effect that for all usages of x in e_2 , the typing constraints for e_1 to which x is bound in Γ will always be instantiated. By the way the algorithm extends the set \mathfrak{L} , $\ell_1 \in \mathfrak{L}$ implies that $\ell_1 \in \mathfrak{L}_0$, i.e., the defining expression of x is ill-typed and does not have a principal type.

The A-VAR-PRIN rule instantiates the typing constraints of the principal type of a let variable x if x is bound in Π and the location of x is not marked to be expanded. Instantiation is done by substituting the type variables $\vec{\alpha}$ that are bound in the schema of the principle type with fresh type variables $\vec{\beta}$. The A-VAR-EXP rule is again similar, except that it handles all usages of let variables that are marked for expansion, as well as all usages of variables that are bound in lambda abstractions.

The remaining rules are relatively straightforward. The rule A-ABS is noteworthy as it simultaneously binds the abstracted variable x to the same type variable α in both typing environments. This ensures that the two environments consistently refer to the same bound type variables when they are used in the subsequent constraint generation and principal type computation within e.

4.5.2 Reduction to Weighted Partial MaxSMT

The reduction to MaxSMT follows the reduction introduced in § 4.3.2. Additionally, the new reduction also must take care of assertions generated for usages of principal types.

Given a cost function R for program p and a set of locations \mathfrak{L} where $\mathfrak{L}_0 \subseteq \mathfrak{L}$, we generate a WPMaxSMT instance $I(p, R, \mathfrak{L}) = (\mathcal{A}_H, \mathcal{A}_S, w)$ as follows. Let $\mathcal{A}_{p,\mathfrak{L}}$ be a set of constraints such that $\emptyset, \emptyset \vdash_{\mathfrak{L}} p : \alpha \mid \mathcal{A}_{p,\mathfrak{L}}$ for some type variable α . Then define

$$\mathcal{A}_{H} = \mathcal{A}_{p,\mathfrak{L}} \cup \{ T_{\ell} \mid \ell \notin \operatorname{dom}(R(p)) \} \cup PDefs(p)$$
$$\mathcal{A}_{S} = \{ T_{\ell} \mid \ell \in \operatorname{dom}(R(p)) \}$$
$$w(T_{\ell}) = R(p)(\ell), \text{ for all } T_{\ell} \in \mathcal{A}_{S}$$

The set of assertions PDefs(p) contains the definitions for the principal type correctness variables P_{ℓ} . For a let variable x that is defined at some location ℓ , the variable P_{ℓ} is defined to be true iff

- each location variable $T_{\ell'}$ for a location ℓ' in the defining expression of x is true, and
- each principal type correctness variable $P_{\ell'}$ for a let variable that is defined at ℓ' and used in the defining expression of x is true.

Formally, PDefs(p) defines the set of formulas

$$PDefs(p) = \{ PDef_{\ell} \mid \ell \in \mathsf{dom}(Uloc_p) \}$$
$$PDef_{\ell} = \left(P_{\ell} \Leftrightarrow \bigwedge_{\ell' \in Loc(\ell)} T_{\ell'} \land \bigwedge_{\ell' \in Vloc(\ell)} P_{dloc(\ell')} \right)$$

Setting the P_{ℓ} to **false** thus captures all possible error sources that involve some of the locations in the defining expression of x, respectively, the defining expressions of other variables that x depends on. Recall that the propositional variable P_{ℓ} is used to guard all the instances of the principal types of x in $\mathcal{A}_{p,\mathfrak{L}}$. Thus, setting P_{ℓ} to **false** will make all usage locations of x well-typed that have not yet been expanded and are thus constrained by the principal type. By the way P_{ℓ} is defined, the cost of setting P_{ℓ} to **false** will be the minimum weight of all the location variables for the locations of x's definition and its dependencies. Thus, P_{ℓ} conservatively approximates all the potential minimum error sources that involve these locations.

We denote by SOLVE the procedure that given p, R, and \mathfrak{L} returns some model M that is a solution of $I(p, R, \mathfrak{L})$. This procedure can be, for instance, implemented using the previously introduced algorithm for computing minimum error sources in § 4.3.

Lemma 22. SOLVE is total.

Lemma 22 follows from our assumption that R is defined for the root location ℓ_p of the program p. That is, $I(p, R, \mathfrak{L})$ always has some solution since \mathcal{A}_H holds in any model M where $M \not\models T_{\ell_p}$.

Given a model $M = \text{SOLVE}(p, R, \mathfrak{L})$, we define L_M to be the set of locations excluded in M:

$$L_M = \{ \ell \in Loc(p) \mid M \models \neg T_\ell \}.$$

4.5.3 Iterative Algorithm

Next, we present our iterative algorithm for computing minimum type error sources.

In order to formalize the termination condition of the algorithm, we first need

to define the set of usage locations of let variables in program p that are in the scope of the current expansion \mathfrak{L} . We denote this set by $Scope(p, \mathfrak{L})$. Intuitively, $Scope(p, \mathfrak{L})$ consists of all those usage locations of let variables that either occur in the body of a top-level let declaration or in the defining expression of some other let variable which has at least one expanded usage location in \mathfrak{L} . Formally, $Scope(p, \mathfrak{L})$ is the largest set of usage locations in p that satisfies the following condition: for all $\ell \in \operatorname{dom}(Uloc_p)$, if $Uloc_p(\ell) \cap \mathfrak{L} = \emptyset \wedge Uloc_p(\ell) \neq \emptyset$, then $Loc(\ell) \cap Scope(p, \mathfrak{L}) = \emptyset$.

For $M = \text{SOLVE}(p, R, \mathfrak{L})$, we then define $Usages(p, \mathfrak{L}, M)$ to be the set of all usage locations of the let variables in p that are in scope of the current expansions and that are marked for expansion. That is, $\ell \in Usages(p, \mathfrak{L}, M)$ iff

- 1. $\ell \in Scope(p, \mathfrak{L})$, and
- 2. $M \not\models P_{dloc(\ell)}$

Note that if the second condition holds, then a potentially cheaper error source exists that involves locations in the definition of the variable x used at ℓ . Hence, that usage of x should not be typed by x's principal type but by the expanded typing constraints generated from x's defining expression.

We say that a solution L_M , corresponding to the result of $SOLVE(p, R, \mathfrak{L})$, is proper if $Usages(p, \mathfrak{L}, M) \subseteq \mathfrak{L}$, i.e., L_M does not contain any usage locations of let variables that are in scope and still typed by unexpanded instances of principal types.

Algorithm 1 shows the iterative algorithm. It takes an ill-typed program p and a cost function R as input and returns a minimum error source. The set \mathfrak{L} of locations to be expanded is initialized to \mathfrak{L}_0 . In each iteration, the algorithm first

Algorithm 1 Iterative algorithm for computing a minimum error source

1: procedure IterMinError (p, R)
2: $\mathfrak{L} \leftarrow \mathfrak{L}_0$
3: loop
4: $M \leftarrow \text{SOLVE}(p, R, \mathfrak{L})$
5: $L_u \leftarrow Usages(p, \mathfrak{L}, M)$
6: if $L_u \subseteq \mathfrak{L}$ then
7: return L_M
8: end if
9: $\mathfrak{L} \leftarrow \mathfrak{L} \cup L_u$
10: end loop
11: end procedure

computes a minimum error source for the current expansion using the procedure SOLVE from the previous section. If the computed error source is proper, the algorithm terminates and returns the current solution L_M . Otherwise, all usage locations of let variables involved in the current minimum solution are marked for expansion and the algorithm continues.

Correctness. We devote the remainder of this section to arguing the correctness of our iterative algorithm. In a nutshell, we show by induction that the solutions computed by our algorithm are also solutions of the naive algorithm that expands all usages of let variables immediately as in [72].

We start with the base case of the induction where we fully expand all constraints, i.e., $\mathfrak{L} = Loc(p)$.

Lemma 23. Let p be a program and R a cost function and let M = SOLVE(p, R, Loc(p)). Then $L_M \subseteq Loc(p)$ is a minimum error source of p subject to R.

Lemma 23 follows from [72, Theorem 1] because if $\mathfrak{L} = Loc(p)$, then we obtain exactly the same reduction to WPMaxSMT as in our previous work. More precisely, in this case the A-VAR-PRIN rule is never used. Hence, all usages of

let variables are typed by the expanded typing constraints according to rule A-VAR-EXP. The actual proof requires a simple induction over the derivations of the constraint typing relation defined in Figure 4.8, respectively, the constraint typing relation defined in 4.3.

We next prove that in order to achieve full expansion it is not necessary that $\mathfrak{L} = Loc(p)$. To this end, define the set \mathfrak{L}_p , which consists of \mathfrak{L}_0 and all usage locations of let variables in p:

$$\mathfrak{L}_p = \mathfrak{L}_0 \ \cup \ \bigcup_{l \in \mathsf{dom}(Uloc_p)} Uloc_p(l).$$

Then $\vdash_{\mathfrak{L}}$ generates the same constraints as $\vdash_{Loc(p)}$ as stated by the following lemma.

Lemma 24. For any p, Π , Γ , α , and \mathcal{A} , we have Π , $\Gamma \vdash_{\mathfrak{L}_p} p : \alpha \mid \mathcal{A}$ iff Π , $\Gamma \vdash_{Loc(p)} p : \alpha \mid \mathcal{A}$.

Lemma 24 can be proved using a simple induction on the derivations of $\vdash_{\mathfrak{L}_p}$, respectively, $\vdash_{Loc(p)}$. First, note that $Loc(p) \setminus \mathfrak{L}_p$ is the set of locations of welltyped let variable definitions in p. Hence, the derivations using $\vdash_{\mathfrak{L}_p}$ will never use the A-LET-EXP rule, only A-LET-PRIN. However, the A-LET-PRIN rule updates both Π and Γ , so applications of A-VAR-EXP (A-VAR-PRIN is never used in either case) will be the same as if $\vdash_{Loc(p)}$ is used.

The following lemma states that if the iterative algorithm terminates, then it computes a correct result.

Lemma 25. Let p be a program, R a cost function, and \mathfrak{L} such that $\mathfrak{L}_0 \subseteq \mathfrak{L} \subseteq \mathfrak{L}_p$. Further, let $M = \text{SOLVE}(p, R, \mathfrak{L})$ such that L_M is proper. Then, L_M is a minimum error source of p subject to R.

The proof of Lemma 25 can be found in Appendix A.7.3. For brevity, we

provide here only the high-level argument. The basic idea is to show that adding each of the remaining usage locations to \mathfrak{L} results in typing constraints for which L_M is again a proper minimum error source. More precisely, we show that for each set \mathfrak{D} such that $\mathfrak{L}_0 \subseteq \mathfrak{L} \subseteq \mathfrak{L} \cup \mathfrak{D} \subseteq \mathfrak{L}_p$, if M is the maximum model of $I(p, R, \mathfrak{L})$ from which L_M was computed, then M can be extended to a maximum model M'of $I(p, R, \mathfrak{L} \cup \mathfrak{D})$ such that $L_{M'} = L_M$. That is, L_M is again a proper minimum error source for $I(p, R, \mathfrak{L} \cup \mathfrak{D})$. The proof goes by induction on the cardinality of the set \mathfrak{D} . Therefore, by the case $\mathfrak{L} \cup \mathfrak{D} = \mathfrak{L}_p$, Lemma 23, and Lemma 24 we have that L_M is a true minimum error source for p subject to R.

Finally, note that the iterative algorithm always terminates since \mathfrak{L} is bounded from above by the finite set \mathfrak{L}_p and \mathfrak{L} grows in each iteration. Together with Lemma 25, this proves the total correctness of the algorithm.

Theorem 6. Let p be a program and R a cost function. Then, ITERMINERROR(p, R) terminates and computes a minimum error source for p subject to R.

4.5.4 Implementation and Evaluation

We now show the results of evaluating our improved algorithm for finding minimum type error sources on the OCaml student benchmark suite from [54] introduced in § 4.4. Additionally, we also evaluate our implementation on larger programs extracted from the GRASShopper [75] program verification tool written in OCaml.

The prototype implementation of the new algorithm was uniformly faster than the naive approach in our experiments. Most importantly, the number of generated typing constraints produced by the algorithm is almost an order of magnitude smaller than when using the naive approach. Consequently, the new algorithm also ran faster in the experiments. We note that the new algorithm and the algorithm from § 4.3 provide the same formal guarantees. Hence, we do not repeat here the experiment on quality of type error sources.

Implementation. Our implementation of the new algorithm directly follows the definition in 4.5.3. That is, our implementation is a simple Python wrapper around the implementation of the naive algorithm.

The generation of typing constraints for each iteration in our algorithm directly follows the typing rules in Figure 4.3. In addition, we perform a simple optimization that reduces the total number of typing constraints. When typing an expression let $x = e_1$ in e_2 , the A-Let-Prin and A-Let-Exp rules always add a fresh instance of the constraint \mathcal{A}_1 for e_1 to the whole set of constraints. This is to ensure that type errors in e_1 are not missed if x is never used in e_2 . We can avoid this duplication of \mathcal{A}_1 in certain cases. If a principal type was successfully computed for the let variable beforehand, the constraints \mathcal{A}_1 must be consistent. If the expression e_1 refers to variables in the environment that have been bound by lambda abstraction, then not instantiating \mathcal{A}_1 at all could make the types of these variables under-constrained. However, if \mathcal{A}_1 is consistent and e_1 does not contain variables that come from lambda abstractions, then we do not need to include a fresh instance of \mathcal{A}_1 in A-Let-Prin. Similarly, if e_1 has no principal type because of a type error and the variable x is used somewhere in e_2 , then the algorithm ensures that all such usages are expanded and included in the whole set of typing constraints. Therefore, we can safely omit the extra instance of \mathcal{A}_1 as well.

Student benchmarks. In our first experiment, we collected statistics for finding a single minimum error source in the student benchmarks with our iterative algorithm and the naive algorithm. We measured the number of typing constraints generated (Fig. 4.9), the execution times (Fig. 4.10), and the number of expansions and iterations taken by our algorithm (Table 4.1).

Figure 4.9 shows the statistics for the total number of generated typing assertions. By typing assertions we mean logical assertions, encoding the typing constraints, that we pass to the weighted MaxSMT solver. The number of typing assertions roughly corresponds to the sum of the total number of locations, constraints attached to each location due to copying, and the number of well typed let definitions. Again, all 8 groups of programs are shown on the x axis in Figure 4.9. The numbers in parenthesis indicate the number of programs in each group. For each group and each approach (naive and iterative), we plot the maximum, minimum and average number of typing assertions. To show the general trend for how both approaches are scaling, lines have been drawn between the averages for each group. (All of the figures in this section follow this pattern.) As can be seen, our algorithm reduces the total number of generated typing assertions. This number grows exponentially with the size of the program for the naive approach. With our approach, this number seems to grow at a much slower rate since it does not expand every usage of a let variable unless necessary. These results make us cautiously optimistic that the number of assertions the iterative approach expands will be polynomial in practice. Note that the total number of typing assertions produced by our algorithm is the one that is generated in the last iteration of the algorithm.

The statistics for execution times are shown in Figure 4.10. The iterative algorithm is consistently faster than the naive solution. We believe this to be a direct consequence of the fact that our algorithm generates a substantially smaller num-



Figure 4.9: Maximum, average, and minimum number of typing assertions for computing a minimum error source by naive and iterative approach

ber of typing constraints. The difference in execution times between our algorithm and the naive approach increases with the size of the input program. Note that the total times shown are collected across all iterations.

We also measured the statistics on the number of iterations and expansions taken by our algorithm. The number of expansions corresponds to the total number of usage locations of **let** variables that have been expanded in the last iteration of our algorithm. The results, shown in Table 4.1, indicate that the total number of iterations required does not substantially change with the input size. We hypothesize that this is due to the fact that type errors are usually tied only to a small portion of the input program, whereas the rest of the program is not relevant to the error.



Figure 4.10: Maximum, average, and minimum execution times for computing a minimum error source by naive and iterative approach

It is worth noting that both the naive and iterative algorithm compute single error sources. The algorithms may compute different solutions for the same input since the fixed cost function does not enforce unique solutions. Both approaches are complete and would compute identical solutions for the all error sources problem [72]. The iterative algorithm does not attempt to find a minimum error source in the least number of iterations possible, but rather it expands let definitions on-demand as they occur in the computed error sources. This means that the algorithm sometimes continues expanding let definitions even though there exists a proper minimum error source for the current expansion. In our future work, we plan to consider how to enforce the search algorithm so that it first finds those minimum error sources that require less iterations and expansions.

	iterations			expansions		
	min	avg	max	min	avg	max
0-50	0	0.49	2	0	1.7	11
50-100	0	0.29	3	0	0.88	13
100-150	0	0.49	4	0	1.37	32
150-200	0	0.44	3	0	1.82	19
200-250	0	0.49	2	0	3.11	30
250-300	0	0.36	2	0	6.04	45
300-350	0	0.67	2	0	3.33	10
350-400	0	0	0	0	0	0

Table 4.1: Statistics for the number of expansions and iterations when computing a single minimum error source

GRASShopper benchmarks. Since we lacked an appropriate corpus of larger ill-typed user written programs, we generated ill-typed programs from the source code of the GRASShopper tool [75]. We chose GRASShopper because it contains non-trivial code that mostly falls into the OCaml fragment supported by Easy-OCaml. For our experiments, we took several modules from the GRASShopper source code and put them together into four programs of 1000, 1500, 2000, and 2500 lines of code, respectively. These modules include the core data structures for representing the abstract syntax trees of programs and specification logics, as well as complex utility functions that operate on these data structures. We included comments when counting the number of program lines. However, comments were generally scars. The largest program with 2500 lines comprised 282 top-level let definitions and 567 let definitions in total. We then introduced separately five distinct type errors to each program, obtaining a new benchmarks suite of 20 programs in total. We introduced common type mismatch errors such as calling a function or passing an argument with an incompatible type.

We repeated the previous experiments on the generated GRASShopper bench-

marks. The benchmarks are grouped by code size. There are four groups of five programs corresponding to programs with 1000, 1500, 2000, and 2500 lines.

Figure 4.11 shows the total number of generated typing assertions subject to the code size. This figure follows the conventions of Fig. 4.9 except that the number of constraints is given on a logarithmic scale. Note that the minimum, maximum, and average points are plotted in Figures 4.11 and 4.12 for each group and algorithm, but these are relatively close to each other and hence visually mostly indistinguishable. The total number of assertions generated by our algorithm is consistently an order of magnitude smaller than when using the naive approach. The naive approach expands all let defined variables where the iterative approach expands only those let definitions that are needed to find the minimum error source. Consequently, the times taken by our algorithm to compute a minimum error source are smaller than when using the naive one, as shown in Figure 4.12. Beside solving a larger weighed MaxSMT instance, the naive approach also has to spend more time generating typing assertions than our iterative algorithm. Finally, Table 4.2 shows the statistics on the number of iterations and expansion our algorithm made while computing the minimum error source. Again, the total number of iterations appears to be independent of the size of the input program.

	iterations			expansions			
	min	avg	max	min	avg	max	
1000	0	0.2	1	0	0.2	1	
1500	0	0.4	2	0	2.8	14	
2000	0	0.6	2	0	53.8	210	
2500	0	0.2	1	0	3	15	

Table 4.2: Statistics for the number of expansions and iterations when computing a single minimum error source for larger programs



Figure 4.11: Maximum, average, and minimum number of typing assertions for computing a minimum error source by naive and iterative approach for larger programs

Comparison to other tools. Our algorithm also outperforms the approach by Myers and Zhang [97] in terms of speed on the same student benchmarks. While our algorithm ran always under 5 seconds, their algorithm took over 80 seconds for some programs. We also ran their tool SHErrLoc [82] on one of our GRASSHopper benchmark programs of 2000 lines of code. After approximately 3 minutes, their tool ran out of memory. We believe this is due to the exponential explosion in the number of typing constraints due to polymorphism. For that particular program, the total number of typing constraints their tool generated was roughly 200,000. On the other hand, their tool shows high precision in correctly pinpointing the actual source of type errors. These results nicely exemplify the nature of type error localization. In order to solve the problem of producing high quality type



Figure 4.12: Maximum, average, and minimum execution times for computing a minimum error source by naive and iterative approach for larger programs

error reports, one needs to consider the whole typing data. However, the size of that data can be impractically large, making the generation of type error reports slow to the point of being not usable. One benefit of our approach is that these two problems can be studied independently. We mainly focused on the second problem, i.e., how to make the search for high-quality type error sources practically fast.

4.6 Related Work

Previous work on localization of type errors has mainly focused on designing concrete systems for generating quality type error messages. The approach taken in [94, 22] computes a trace of type inference, i.e., a set of decisions made by type inference algorithm that led to a type error. Similarly, type error localization algorithms presented in [89, 30, 77] show a program slice involved in the error. The obvious difference between these approaches and the one taken in this thesis concerns the used notion of a type error source. In our work, type error sources are minimal sets of expressions that, once corrected, yield a well-typed program and are guaranteed to satisfy the compiler or user provided optimality criterion. In contrast, modeling error sources as traces and slices often results in redundant information being reported to the programmer, such as statements or type inference decisions that are related to the error but do not cause it. The works presented in [14, 13, 67] design completely new type systems that, while performing the type inference, also collect additional information that can be reported to the programmer in the case of type errors. Our work, on the other hand, does not require changes to the existing type systems yet it simply piggybacks on top of them. More closely related to our approach is the **Seminal** [54] tool, which computes several possible error sources by repeated calls to the type checker rather than an SMT solver. However, the search for error causes is based on fixed heuristics and provides no formal guarantees that all error sources are found, respectively, that they are ranked according to some criterion. Sulzmann [86, 84] presents a type error debugger for Haskell that presents to the programmer a minimal set of locations whose corresponding expressions are responsible for the type error. Our work also considers a minimal set of expressions inducing a type error as a potential error source, but we also allow incorporating heuristics for choosing which minimal set will be reported to the programmer as there can be many such sets. Zhang and Myers [97] encode typing information for Hindley-Milner type systems in terms of constraint graphs. The generated graphs are then analyzed to find most likely error

sources by using Bayesian inference. It is unclear how this approach would support more expressive type systems. Also, the constraint graph generation suffers from the exponential blow-up due to polymorphism, which our work addresses by the lazy expansion of constraints induced by typing polymorphic functions (§ 4.5). On the other hand, the proposed technique is designed for general diagnosis of static errors, whereas we focus specifically on diagnosis of type errors. In a follow up work, the authors also apply their techniques to Haskell [98].

Previous approaches based on constraint solving [33, 85] produce minimal but not minimum error sources and consider specific ranking criteria for specific type systems. Hage and Heeren [35] build a type error debugging system for Haskell by incorporating several fixed and expert designed heuristics for explaining type errors. Our work, in contrast, provides a general framework for type error debugging where the heuristic for identifying the root cause of a type error is merely an input to the system. The work by the same authors that is more related to our work can be found in [36]. Here, the authors present a framework for constraint-based type inference that allows compilers to plug in various different strategies for solving the constraints, consequently allowing more flexibility for producing type error reports. These strategies allow different constraint solving decision based on the context and can be used, for instance, to manipulate the order in which the constraints are solved. In contrast to our work, it is unclear whether their system produces the optimal solution given a constraint-solving input strategy. Seidel et al. devise a general machine learning algorithm that blames program expressions for causing type errors by offline learning from the corpus of ill-typed programs [81]. While their approach can support various type systems and ranking criteria, it does not provide a formal guarantee that the blamed program expressions indeed constitute an actual error source. The problem of type error localization has also been studied in the context of Java generics [12], domain-specific languages [34], and dependently typed languages [25].

Following our work initially published in [72, 73], the MaxSMT approach to type inference has also been applied to Python [39]. Here, a MaxSMT solver is used to infer types for as large as possible portion of the given program, which is particularly useful since Python is a dynamically typed language and inferring static types for the whole program is consequently often impossible. We note that our approach is in part inspired by the Bug-Assist tool [47], which uses a MaxSAT procedure for fault localization in imperative programs. However, the problem we are solving in this thesis is quite different.
Chapter 5

Conclusions

The goal of this thesis is to provide a mathematically rigorous framework for the systematic development of type inference algorithms that are convenient to use by the programmers. Towards achieving that goal, we focused on two important problems surrounding type inference: (1) how to constructively design type inference algorithms that improve over the state-of-the-art and (2) how to automatically debug type errors that arise during inference. This thesis approaches these two specific problems based on the close connection between type inference and program analysis techniques. We addressed the first problem by using abstract interpretation to constructively design Liquid type inference, an advanced family of algorithms that combine classical typing disciplines and known static analyses to prove various safety properties of functional programs. By rigorously modeling Liquid type inference using the primitives of abstract interpretation, we unveiled the design space of such algorithms and generalized it in a way that allows easy construction of novel type inference algorithms that are sound by construction. Regarding the second problem, we have shown how the problem of type

error localization for type inference algorithms in the style of Hindley-Milner can be cast as an optimization problem expressed in a formal logic. We then showed how this problem can be solved using automated theorem provers. Finally, we experimentally illustrated how our approach yields type error debugging systems that outperform the state-of-the-art tools in correctly identifying the root causes of type errors while still being efficient, even for large programs.

The results presented in this thesis naturally open several interesting research directions. An exciting future work project is to use our abstract interpretation model of Liquid types to obtain and implement type inference algorithms that are more expressive than the original Liquid type inference. This could be achieved by instantiating our generalization of Liquid types with, say, the polyhedra abstract domain or automata-based domains. Another interesting research problem is to extend our framework for type error debugging to support type inference algorithms that are more expressive than the Hindley-Milner algorithm, such as Liquid types for instance. Due to the close connection between type inference and static analysis, the problem of building a general framework for debugging static analyses naturally arises. In general, we believe this thesis provides important results and generally paves the way for the systematic study and development of enhanced type analyses that are usable by the programmers.

Appendix A

Data Flow Refinement Type Inference

A.1 Example Concrete Map

Let us reconsider the example from § 3.1. The code is again given below, but this time we also annotate expressions of interest with location labels.

1 let dec y = y - 1 in 2 let f x g = if x >= 0 then $g^o x^p$ else x^q 3 in $f^a (f^b 1^c dec^d)^j dec^m$

The concrete execution map M_e for this program, restricted to the execution nodes

of interest is as follows:

$$\begin{array}{ll} q^{a}\mapsto \bot & q^{b}\mapsto \bot & c\mapsto 1 & p^{a}\mapsto 0 & p^{b}\mapsto 1 \\ j\mapsto 0 & d\mapsto [p^{b}\mapsto (1,\ 0)] & m\mapsto [p^{a}\mapsto (0,\ -1)] \\ dec\mapsto [p^{b}\mapsto (1,\ 0),\ p^{a}\mapsto (0,\ -1)] \\ a\mapsto [j\mapsto (0,\ [m\mapsto ([p^{a}\mapsto (0,\ -1)],\ -1)])] \\ b\mapsto [c\mapsto (1,\ [d\mapsto ([p^{b}\mapsto (1,\ 0)],\ 0)])] \\ f\mapsto [j\mapsto (0,\ [m\mapsto ([p^{a}\mapsto (0,\ -1)],\ -1)]), \\ c\mapsto (1,\ [d\mapsto ([p^{b}\mapsto (1,\ 0)],\ 0)])] \end{array}$$

For this specific program, the execution will always reach the program expressions with locations a, b, c, d, j, and m with the same environment. For this reason, we identify these expression nodes with their locations. Expressions with labels o, p, and q will be reachable with two environments, each of which corresponds to executing the body of f with the value passed to it at call sites with locations a and b. We use this call site information to distinguish the two evaluations of the body of f. For instance, the expression node associated with the x^q expression when performing the analysis for the call to f made at location a is denoted by q^a in the above table.

A.2 Concrete Semantics Proofs

Lemma 26. prop is increasing.

Proof. Let $\langle v'_1, v'_2 \rangle = \mathsf{prop}(v_1, v_2)$. The proof goes by mutual structural induction on v_1 and v'_2 .

Base cases.

• $v_1 = \bot$. $v_1' = \bot, v_2' = v_2$ def. of prop $v_1 \sqsubseteq v'_1, v_2 \sqsubseteq v'_2$ def. of \sqsubseteq • $v_1 = \omega$. $v_1' = v_2' = \omega$ def. of prop $v_1 \sqsubseteq v'_1, v_2 \sqsubseteq v'_2$ def. of \sqsubseteq • $v_1 = c, v_2 = \bot$. $v_1' = v_2' = c$ def. of prop $v_1 \sqsubseteq v'_1, v_2 \sqsubseteq v'_2$ def. of \sqsubseteq • $v_1 = c, v_2 = c.$ $v_1' = v_2' = c$ def. of prop $v_1 \sqsubseteq v'_1, v_2 \sqsubseteq v'_2$ def. of \sqsubseteq • $v_1 = c_1, v_2 = c_2, c_1 \neq c_2.$ $v_1' = v_1, v_2' = \omega$ def. of prop $v_1 \sqsubseteq v'_1, v_2 \sqsubseteq v'_2$ def. of \sqsubseteq • $v_1 = c_1, v_2 = \omega$.

$v_1' = v_1, v_2' = \omega$	def. of prop
$v_1 \sqsubseteq v'_1, v_2 \sqsubseteq v'_2$	def. of \sqsubseteq

• $v_1 = c, v_2 \in \mathcal{T}$.

 $v'_1 = v_1, v'_2 = \omega$ def. of prop $v_1 \sqsubseteq v'_1, v_2 \sqsubseteq v'_2$ def. of \sqsubseteq

• $v_1 \in \mathcal{T}, v_2 = \bot$.

$$v'_1 = v_1, v'_2 = T_\perp$$
 def. of prop
 $v_1 \sqsubseteq v'_1, v_2 \sqsubseteq v'_2$ def. of \sqsubseteq

• $v_1 \in \mathcal{T}, v_2 = c$. $v'_1 = v_1, v'_2 = \omega$ def. of prop $v_1 \sqsubseteq v'_1, v_2 \sqsubseteq v'_2$ def. of \sqsubseteq

• $v_1 \in \mathcal{T}, v_2 = \omega.$

 $v_1' = v_2' = \omega$ def. of prop $v_1 \sqsubseteq v_1', v_2 \sqsubseteq v_2'$ def. of \sqsubseteq

Induction case: $v_1 = T_1, v_2 = T_2$. By definition of prop, we need to consider two cases of call-sites subject to the table T_2 .

• $n_{cs} \notin T_2$.

$$v_{1i} = v'_{1i}, v_{2i} = v'_{2i}, v_{1o} = v'_{1o}, v_{20} = v'_{2o}$$
 def. of prop
$$v_{1i} \sqsubseteq v'_{1i}, v_{2i} \sqsubseteq v'_{2i}, v_{1o} \sqsubseteq v'_{1o}, v_{20} \sqsubseteq v'_{2o}$$
 def. of \sqsubseteq

• $n_{cs} \in T_2$.

$$v_{1i} \sqsubseteq v'_{1i}, v_{2i} \sqsubseteq v'_{2i}, v_{1o} \sqsubseteq v'_{1o}, v_{20} \sqsubseteq v'_{2o} \qquad \text{def. of prop and i.h.}$$
(A.1)

Lemma 27. step is increasing.

Proof. By structural induction on program expressions. The result follows from the increasing property of the join \sqcup and Lemma 26.

Lemma 28. prop is monotone.

Proof. Let $\langle v'_1, v'_2 \rangle = \operatorname{prop}(v_1, v_2)$, $\langle v'_3, v'_4 \rangle = \operatorname{prop}(v_3, v_4)$, $v_1 \sqsubseteq v_3$, and $v_2 \sqsubseteq v_4$. We show that $v'_1 \sqsubseteq v'_3$ and $v'_2 \sqsubseteq v'_4$. The proof goes by mutual structural induction on v_3 and v_4 .

Base cases.

• $v_3 = \omega$.

 $\begin{array}{ll} v_3' = \omega, v_4' = \omega & \quad \mbox{def. of prop} \\ v_1' \sqsubseteq v_3', v_2' \sqsubseteq v_4' & \quad \mbox{def. of } \sqsubseteq \end{array}$

• $v_4 = \omega$.

$v_4' = \omega$	def. of prop
$v_2' \sqsubseteq v_4'$	by ⊑

 $-v_3 = c$

by $v_1 \sqsubseteq v_3$	$v_1 = \bot \lor v_1 = c$
def. of prop	$v_3' = c, v_1' = \bot \lor v_1' = c$
by 	$v_1' \sqsubseteq v_3'$

 $-v_3 \in \mathcal{T}$

$v_3' = \omega$	def. of prop
$v_1' \sqsubseteq v_3'$	by

• $v_3 = \bot$.

by $v_1 \sqsubseteq v_3$	$v_1 = \bot$
def. of prop	$v_4' = v_4, v_2' = v_2, v_3' = v_3, v_1' = v_1$
def. of \sqsubseteq	$v_1' \sqsubseteq v_3', v_2' \sqsubseteq v_4'$

• $v_4 = \bot$.

$$v_2 = \bot$$
 by $v_2 \sqsubseteq v_4$
 $v'_1 = v_1, v'_3 = v_3$ def. of prop $v'_1 \sqsubseteq v'_3$ def. of \sqsubseteq

 $-v_3 = c$

$$v_1 = \bot \lor v_1 = c \qquad \qquad \text{by } v_1 \sqsubseteq v_3$$
$$v'_4 = c, v'_2 = \bot \lor v'_2 = c \qquad \qquad \text{def. of prop}$$
$$v'_2 \sqsubseteq v'_4 \qquad \qquad \text{by } \sqsubseteq$$

 $-v_3 \in \mathcal{T}$

$$v_{1} = \bot \lor v_{1} \in \mathcal{T} \qquad \qquad \text{by } v_{1} \sqsubseteq v_{3}$$
$$v'_{4} = T_{\bot}, v'_{2} = \bot \lor v'_{2} = T_{\bot} \qquad \qquad \text{def. of prop}$$
$$v'_{2} \sqsubseteq v'_{4} \qquad \qquad \text{by } \sqsubseteq$$

•
$$v_3 = c, v_4 = c.$$

 $v'_4 = v_4 = v_3 = v'_3 = c$ def. of prop
 $v_1 = \bot \lor v_1 = c, v_2 = \bot \lor v_2 = c$ by $v_1 \sqsubseteq v_3, v_2 \sqsubseteq v_4$
 $v'_1 = \bot \lor v'_1 = c, v'_2 = \bot \lor v'_2 = c$ def. of prop
 $v'_1 \sqsubseteq v'_3, v'_2 \sqsubseteq v'_4$ def. of \sqsubseteq

• $v_3 = c_1, v_4 = c_2, c_1 \neq c_2.$

$$\begin{aligned} v_4' &= \omega, v_3 = v_3' = c_1 & \text{def. of prop} \\ v_1 &= \bot \lor v_1 = c_1 & \text{by } v_1 \sqsubseteq v_3, v_2 \sqsubseteq v_4 \\ v_1' &= \bot \lor v_1' = c_1 & \text{def. of prop} \\ v_1' &\sqsubseteq v_3', v_2' \sqsubseteq v_4' & \text{def. of } \sqsubseteq \end{aligned}$$

141

• $v_3 = c, v_4 \in \mathcal{T}$.

def. of prop	$v_4' = \omega, v_3 = v_3' = c$
by $v_1 \sqsubseteq v_3, v_2 \sqsubseteq v_4$	$v_1 = \bot \lor v_1 = c$
def. of prop	$v_1' = \bot \lor v_1' = c$
def. of \sqsubseteq	$v_1' \sqsubseteq v_3', v_2' \sqsubseteq v_4'$

• $v_3 \in \mathcal{T}, v_4 = c.$

def. of prop	$v_4' = \omega, v_3 = v_3'$
by $v_1 \sqsubseteq v_3, v_2 \sqsubseteq v_4$	$v_1 = \bot \lor v_1 \in \mathcal{T}, v_2 = \bot \lor v_2 = c$
def. of prop	$v_1' = v_1$
def. of \sqsubseteq	$v_1' \sqsubseteq v_3', v_2' \sqsubseteq v_4'$

Induction case: $v_3 = T_3$ and $v_4 = T_4$. If $v_1 = \bot$, then $v'_1 = \bot$ and $v'_1 \sqsubseteq v'_3$ trivially follows. When $v_2 = \bot$, we have $v'_2 = \bot \lor v'_2 = T_{\bot}$ and the argument follows from Lemma 26.

We are hence left with the case when $v_1 = T_1$ and $v_2 = T_2$. First, we have that $n_{cs} \in T_2 \Longrightarrow n_{cs} \in T_4$ by the $T_2 \sqsubseteq T_4$ assumption. We have three cases to consider.

- $n_{cs} \notin T_2$ and $n_{cs} \notin T_4$. Trivial, as there is no propagation.
- $n_{cs} \notin T_2$ and $n_{cs} \in T_4$. By prop, there is no propagation between the pairs $T_1(n_{cs})$ and $T_2(n_{cs})$. The result then follows from Lemma 26.
- $n_{cs} \in T_2$ and $n_{cs} \in T_4$. The result follows from the induction hypothesis.

Lemma 29. step is monotone.

Proof. Let $M' \sqsubseteq M''$ and $k \in \mathbb{N}$. For any E and e, we show that $M'_f \sqsubseteq M''_f$ where $M'_f = \operatorname{step}_k[\![e]\!](E)(M')$ and $M''_f = \operatorname{step}_k[\![e]\!](E)(M'')$. The proof goes by structural induction on e and natural induction on k. When k = 0, the argument is trivial. We thus consider the case when k > 0 and do the case analysis on e.

- $e = c^{\ell}$. $M'_f \sqsubseteq M''_f$ follows from the monotonicity of \sqcup .
- $e = x^{\ell}$. Here, $M'_f \doteq M''_f$ follows from the monotonicity of prop (Lemma 28).
- $e = (e_1^{\ell_1} e_2^{\ell_2})^{\ell}$. We do case analysis on the values computed for e_1 and e_2 by $\operatorname{step}_k[\![e]\!](E)(M')$ and $\operatorname{step}_{k+1}[\![e]\!][k](E)(M'')$.
 - $-v'_1 = \perp$. Then $M'_f = M'_1$ and the result holds by Lemma 27.
 - $v_1'' = \bot$. By i.h., $M_1' \doteq M_1''$, hence $v_1' = \bot$ and the argument follows from the previous case.
 - $-v'_1 \neq \perp \wedge v'_1 \notin \mathcal{T}$. By i.h., $M'_1 \doteq M''_1$ and hence $v''_1 \neq \perp \wedge v'' \notin \mathcal{T}$. Therefore, $M''_f = M_\omega$ and the result follows trivially.
 - $-v_1' \in \mathcal{T}, v_1'' \notin \mathcal{T}$. The argument is trivial as $M_f'' = M_{\omega}$.
 - $-v'_1 \in \mathcal{T}, v''_1 \in \mathcal{T}$. By i.h., $v'_1 \sqsubseteq v''_1$. The cases for v'_2 and v''_2 are handled the same way as for v'_1 and v''_1 . We are thus interested in the remaining case where $v'_2, v''_2 \notin \{\perp, \omega\}$ where by i.h. $v'_2 \sqsubseteq v''_2$. The result then follows by the monotonicity of prop (Lemma 28).
- $e = (\lambda x.e)^{\ell}$. The argument, similar to the case for function applications, follows from Lemmas 28, 26, and basic properties of joins \sqcup .

Proof (of Lemma 2).

Proof. Follows directly from Lemma 26 and Lemma 28.

Proof (of Lemma 3).

Proof. Follows directly from Lemma 27 and Lemma 29. \Box

We next prove that our definitions of joins and meets are indeed correct. The definition of a binary join on concrete values is raised to an arbitrary number of values in the expected way where $\sqcup \emptyset \stackrel{\text{def}}{=} \bot$ and $\sqcup \{v\} \stackrel{\text{def}}{=} v$. The meet \sqcap is defined similarly. We also define a depth of a concrete data flow value in the expected way. The depth of constants, errors, and unreachable values is simply 0. The depths of tables is the maximum depth of any of the values stored for any call site, increased by 1.

Definition 3 (Depth of Values).

$$\begin{split} \operatorname{depth}(\bot) \stackrel{\text{\tiny def}}{=} \operatorname{depth}(c) \stackrel{\text{\tiny def}}{=} \operatorname{depth}(\omega) \stackrel{\text{\tiny def}}{=} 0 \\ \operatorname{depth}(T) \stackrel{\text{\tiny def}}{=} 1 + \max(\bigcup_{n_{\mathrm{cs}}} \operatorname{depth}(\pi_1(T(n_{\mathrm{cs}}))) \cup \bigcup_{n_{\mathrm{cs}}} \operatorname{depth}(\pi_1(T(n_{\mathrm{cs}})))) \end{split}$$

Note that **depth** is well-defined since, although there is an infinite number of call sites in a table, the depth of every value stored at every call site in a table is bounded by the table depth, as tables are defined inductively.

Proof (of Lemma 1).

Proof. We carry the proof by case analysis on elements of V and induction on the minimum depth thereof when V consists of tables only.

We first consider the trivial (base) cases where V is not a set of multiple tables.

- $V = \emptyset$. By \sqsubseteq , $lub(V) = \bot$ and $\sqcup(V) = \bot$ by \sqcup .
- $V = \{v\}$. Trivial.

- $c \in V, T \in V$. By \sqsubseteq , lub of $\{c, T\}$ is ω and since ω is the top element, $lub(V) = \omega = \sqcup V$.
- $c_1 \neq c_2, c_1 \in V, c_2 \in V$. As in the previous case.
- $\perp \in V$. By \sqsubseteq , $lub(V) = lub(V/\{\perp\})$ and $\sqcup(V) = \sqcup(V/\{\perp\})$ by \sqcup . The set $V/\{\perp\}$ either falls into one of the above cases or consists of multiple tables, which is the case we show next.

Let $V \subseteq \mathcal{T}$ and |V| > 1. Let d be the minimum depth of any table in V. By \sqsubseteq , lub of V is a table T that for every call site n_{cs} has input (resp. output) that is lub of the corresponding inputs (resp. outputs) of tables in V for n_{cs} . We focus on inputs; the same argument goes for outputs. Let $V_i = \{\pi_1(v(n_{cs})) \mid n_{cs} \in \mathcal{N}_e \land v \in V\}$. By \sqcup , $\pi_1(T(n_{cs})) = \sqcup V_i$. We hence need to show that $\sqcup V_i = lub(V_i)$. The set V_i either falls into one of the trivial cases, from which the argument easily follows, or $V_i \subseteq \mathcal{T} \land |V| > 1$. In the latter case, the minimum depth of any value in V_i is smaller than d by the definition of **depth** and the result follows from the i.h.

The proof of correctness for the meet \sqcap is similar. \square

A.2.1 Additional Properties

We now introduce and prove several interesting properties related to the concrete data flow semantics.

A.2.1.1 Finiteness

We say a value v is finite fin(v) if it is either a ground value or a table that has seen only a finite number of inputs.

$$\begin{aligned} & \operatorname{fin}(v) \stackrel{\text{def}}{=} (v \in \{\bot, \omega\} \cup Cons) \lor (v \in \mathcal{T} \land |\{n_{\mathsf{cs}} \mid n_{\mathsf{cs}} \in v\}| \in \mathbb{N} \land \\ & \forall n_{\mathsf{cs}} \in v. \operatorname{fin}(\pi_1(v(n_{\mathsf{cs}}))) \land \operatorname{fin}(\pi_2(v(n_{\mathsf{cs}})))) \end{aligned}$$

It can be shown that prop does not affect the finiteness of values. The same clearly holds for joins \sqcup .

Lemma 30. Let v_1 , v_2 , v_3 , and v_4 be values. If $fin(v_1)$, $fin(v_2)$, and $\langle v_3, v_4 \rangle = prop(v_1, v_2)$, then $fin(v_3)$ and $fin(v_4)$.

Proof. By structural induction on v_1 and v_2 . The cases when either v_1 or v_2 are not a table are trivial. For tables, the argument follows by i.h. and the simple fact that **prop** does not invent new call site nodes, yet only manipulates the existing ones, whose total number is finite by the lemma assumption.

We say a map M is finite fin(M) if all of its values are finite and the set of reachable nodes is finite as well $|\{n \mid M(n) \neq \bot\}| \in \mathbb{N}$. It can be shown that step preserves the finiteness of maps.

Lemma 31. Let M and M' be maps, E an environment, e^{ℓ} an expression, and k a fuel. If fin(M), (e, E) is well-formed, and step_k $[\![e^{\ell}]\!](E)(M) = M'$, then fin(M').

Proof. The proof goes by natural induction on k and structural induction on e. The argument follows from the Lemma 30. Note that in the case e is a lambda abstraction and $M(E \diamond \ell) = T$, for some table T, the set of n_{cs} in T is finite. Hence, the set \overline{M} is finite and by i.h. each $M' \in \overline{M}$ is finite as well. The result then follows from the basic properties of joins.

A.2.1.2 Computation Range

We define an *expression range* of an expression node $E \diamond \ell$ subject to a program p to be a set of expression nodes whose location is a sublocation of ℓ in p and whose environment is an extension of E

$$\operatorname{erange}_p(E \diamond \ell) \stackrel{\text{\tiny def}}{=} \{ E' \diamond \ell' \mid E \subseteq E' \land \ell' \in Loc(p(\ell)) \}.$$

The corresponding *range* is then defined as

 $\mathsf{range}_p(E \diamond \ell) \stackrel{\text{\tiny def}}{=} \mathsf{erange}_p(E \diamond \ell) \cup \{n \mid n \in \mathsf{rng}(E') \land E' \diamond_- \in \mathsf{erange}_p(E \diamond \ell))\}$

Note that in the above definition we use the $_{-}$ pattern to existentially quantify a certain object instead of explicitly writing down the quantifier. We do this to avoid clutter.

Let domain validity domvalid(ℓ, p, E) of an environment E subject to program p and a location ℓ in p holds iff dom(E) consists of exactly all variables bound above ℓ in p. We use this definition to characterize the range of a computation performed by step.

Lemma 32. Let M and M' be maps, E an environment, e^{ℓ} an expression of a program p, and k a fuel. If domvalid (ℓ, p, E) and $\operatorname{step}_k[\![e^{\ell}]\!](E)(M) = M'$, then $\forall n.M(n) \neq M'(n) \implies n \in \operatorname{range}_p(E \diamond \ell).$

Proof. The proof goes by natural induction on k and structural induction on e. The domain validity assumption together with the uniqueness of locations in p guarantees that environments are always extended, i.e., their contents never get rewritten by environment updates done when analyzing the bodies of lambda abstraction, which is needed in order to ensure that updated nodes are in the erange of the current node.

We now state few properties relating ranges. First, ranges of non-related ex-

pression do not share visited expression nodes.

Lemma 33. Let *E* be an environment, *p* be a program, ℓ_1 and ℓ_2 two locations such that $Loc(p(\ell_1)) \cap Loc(p(\ell_2)) = \emptyset$. Then, $range_p(E \diamond \ell_1) \cap range_p(E \diamond \ell_2) \cap \mathcal{N}_e = \emptyset$.

Proof. Follows directly from $Loc(p(\ell_1)) \cap Loc(p(\ell_2)) = \emptyset$.

Similarly, the ranges of nodes whose environments are non-related do not have any common expression nodes.

Lemma 34. Let E_1 and E_2 be two environments and ℓ a location of a program p. If dom $(E_1) = \text{dom}(E_2)$ and $E_1 \neq E_2$, then $\text{range}_p(E \diamond \ell_1) \cap \text{range}_p(E \diamond \ell_2) \cap \mathcal{N}_e = \emptyset$.

Proof. The argument follows from the fact that the extensions of E_1 and E_2 have to be different as $dom(E_1) = dom(E_2)$ and environments are modeled as functions. \Box

A.2.1.3 Refined Range

We can also give a more refined definition of a range.

$$\begin{split} \operatorname{eranger}_p(E \diamond \ell) \stackrel{\text{\tiny def}}{=} \mathbf{lfp}_{\{E \diamond \ell\}}^{\subseteq} \lambda L. \ L \cup \{E' \diamond \ell' \mid \ell' \in Loc(p(\ell)) \land E \subseteq E' \land \forall n \in \operatorname{rng}(E' \backslash E) \\\\ \exists n' \in L. \ env(n) = env(n') \land loc(n) \in Loc(p(\ell)) \} \end{split}$$

The corresponding *range* is then defined as expected

$$\operatorname{ranger}_p(E \diamond \ell) \stackrel{\text{\tiny def}}{=} \operatorname{eranger}_p(E \diamond \ell) \cup \{n \mid n \in \operatorname{rng}(E') \land E' \diamond_{\text{-}} \in \operatorname{eranger}_p(E \diamond \ell))\}$$

The same result about the range of computation can be shown.

Lemma 35. Let M and M' be maps, E an environment, e^{ℓ} an expression of a program p, and k a fuel. If domvalid (ℓ, p, E) and $\operatorname{step}_k[\![e^{\ell}]\!](E)(M) = M'$, then $\forall n.M(n) \neq M'(n) \implies n \in \operatorname{ranger}_p(E \diamond \ell).$

Proof. Similar to the proof of Lemma 32.

However, something stronger can be said about the disjointness.

Lemma 36. Let *E* be an environment, *p* be a program, ℓ_1 and ℓ_2 two locations such that $Loc(p(\ell_1)) \cap Loc(p(\ell_2)) = \emptyset$. Then, $range_p(E \diamond \ell_1) \cap range_p(E \diamond \ell_2) = rng(E)$.

Proof. Expression nodes of both ranges are clearly different by $Loc(p(\ell_1)) \cap Loc(p(\ell_2)) = \emptyset$. \emptyset . The environment nodes n not in rng(E) must have the associated variable loc(n)from $Loc(p(\ell_1))$ and $Loc(p(\ell_2))$, respectively. The result then again follows from $Loc(p(\ell_1)) \cap Loc(p(\ell_2)) = \emptyset$.

Similarly, the ranges of nodes whose environments are non-related have the same property.

Lemma 37. Let E_1 and E_2 be two environments and ℓ a location of a program p. If dom $(E_1) = dom(E_2)$ and $E_1 \neq E_2$, then range_p $(E \diamond \ell_1) \cap range_p(E \diamond \ell_2) = rng(E_1) \cap rng(E_2)$.

Proof. The environment nodes n in both ranges must have the associated environment env(n) equal to the environment of some expression node in the corresponding range. These environments are by definition the extensions of E_1 and E_2 , respectively, that have to be different as $dom(E_1) = dom(E_2)$ and environments are modeled as functions. Hence, the only shared nodes are those environment nodes that appear in both E_1 and E_2 .

A.2.1.4 Environment Consistency

We next formalize the notion of environment consistency for a map that states that no node can be visited with an environment holding a \perp or ω value.

 $\mathsf{envcons}(M) \iff \forall n, n_x \in \mathsf{rng}(env(n)). \ M(n) \neq \bot \implies M(n_x) \not\in \{\bot, \omega\}$

Additionally, we also define consistency of an environment E subject to a map M

 $\mathsf{envcons}_M(E) \iff \forall n_x \in \mathsf{rng}(E). \ M(n_x) \not\in \{\bot, \omega\}$

A.2.1.5 Table IO Consistency

We next formalize the notion of input output consistency for tables. tiocons(v) $\iff (v \in \{\bot, \omega\} \cup Cons) \lor (v \in \mathcal{T} \land \forall n_{cs} \in v.\pi_1(v(n_{cs})) = \bot \Longrightarrow \pi_2(v(n_{cs})) = \bot$

 $\land \forall n_{cs} \in v. \operatorname{tiocons}(\pi_1(v(n_{cs}))) \land \operatorname{tiocons}(\pi_2(v(n_{cs}))))$

As expected, a map is input-output consistent tiocons(M) if all of its values are IO consistent.

A.2.1.6 Constant Consistency

We next formalize the notion of constant consistency that intuitively states that reachable nodes corresponding to constant and lambda expression must be either constants and tables, respectively, or an error value. Let p be a program. A map M is constant consistent $\operatorname{constcons}_p(M)$ if for every node n, where M(n) = v and $loc(n) = \ell$, it must be

$$v \in \{\omega, \bot, c\} \qquad \text{if } p(\ell) = c$$
$$v \in \{\omega, \bot\} \cup \mathcal{T} \qquad \text{if } p(\ell) = \mu f. \lambda x. e$$

A.2.1.7 Error Consistency

We continue by formalizing the notion of the error absence of values that intuitively states the a given value is not an error and, if it is a table, it does not recursively contain errors.

$$\mathsf{noerr}(v) \iff (v \in \{\bot\} \cup Cons) \lor (v \in \mathcal{T} \land \forall n_{\mathsf{cs}} \in v.$$

 $\operatorname{noerr}(\pi_1(v(n_{cs}))) \wedge \operatorname{noerr}(\pi_2(v(n_{cs}))))$

A map M is error consistent $\operatorname{errcons}(M)$ if either (1) $M = M_{\omega}$ or (2) for every node n it must be that $\operatorname{noerr}(M(n))$.

A.2.1.8 Concrete Subtyping

Next, we introduce the notion of subtyping for our concrete values that resembles the subtyping on types.

$$v_1 <:^{c} v_2 \iff (v_2 = \bot) \lor (v_1 = v_2 \land v_1 \in \{\omega\} \cup Cons) \lor (v_1 \in \mathcal{T} \land v_2 \in \mathcal{T} \land \forall n_{\mathsf{cs}} \in v_2. \ \pi_1(v_2(n_{\mathsf{cs}})) <:^{c} \pi_1(v_1(n_{\mathsf{cs}})) \land \pi_2(v_1(n_{\mathsf{cs}})) <:^{c} \pi_2(v_2(n_{\mathsf{cs}})))$$

A map M is subtype consistent $\operatorname{csub}_p(M)$, subject to a program p, iff for every node n, where $loc(n) = \ell$, it must be that

$$\begin{split} M(env(n)(x)) <:^{c} M(n) & \text{if } p(\ell) = x \text{ and } x \in \operatorname{dom}(env(n)) \\ M(env(n) \diamond \ell_{1}) <:^{c} [env(n) \diamond \ell_{2} \mapsto \langle M(env(n) \diamond \ell_{2}), M(n) \rangle] \\ & \text{if } p(\ell) = (e_{1}^{\ell_{1}} e_{2}^{\ell_{2}}) \text{ and } M(env(n) \diamond \ell_{1}) \in \mathcal{T} \end{split}$$

 $M(env(n) \diamond \ell_1) <:^c M(n)$

if
$$p(\ell) = (e_0^{\ell_0} ? e_1^{\ell_1} : e_2^{\ell_2})$$
 and $M(env(n) \diamond \ell_0) = true$

$$M(env(n)\diamond \ell_2) <: {}^c M(n)$$

$$\begin{split} \mathbf{if} \ p(\ell) &= (e_0^{\ell_0} ? e_1^{\ell_1} : e_2^{\ell_2}) \ \mathbf{and} \ M(env(n) \diamond \ell_0) = false \\ M(n) <:^c M(n_f) \land [n_{\mathsf{cs}} \mapsto \langle M(n_x), M(E_1 \diamond \ell_1) \rangle] <:^c M(n) \\ & \mathbf{if} \ p(\ell) = (\mu f. \lambda x. e_1^{\ell_1}), \ M(n) \in \mathcal{T}, n_{\mathsf{cs}} \in M(n), \end{split}$$

 $n_x = env(n) \diamond n_{cs} \diamond x, \ n_f = env(n) \diamond n_{cs} \diamond f, \ \text{and} \ E_1 = env(n).x : n_x.f : n_f$

We can also show that propagation does not invalidate subtyping.

Lemma 38. Let $v_1 \in \mathcal{V}$, $v_2 \in \mathcal{V}$, $v_3 \in \mathcal{V}$, and $v_4 \in \mathcal{V}$. If $v_1 <:^c v_2$ and $\langle v_3, v_4 \rangle = prop(v_1, v_2)$, then $v_3 <:^c v_4$.

Proof. By structural induction on v_1 and v_2 .

Lemma 39. Let $v_1 \in \mathcal{V}, v_2 \in \mathcal{V}, v_3 \in \mathcal{V}, and v_4 \in \mathcal{V}$. If $v_1 <:^c v_2, v_1 \notin \{\perp, \omega\}$, and $\langle v_3, v_4 \rangle = \operatorname{prop}(v_1, v_2)$, then $v_3 \notin \{\perp, \omega\}$ and $v_4 \notin \{\perp, \omega\}$.

Proof. Immediate from the definition of prop.

We can even say something stronger.

Lemma 40. Let $v_1 \in \mathcal{V}$, $v_2 \in \mathcal{V}$, $v_3 \in \mathcal{V}$, and $v_4 \in \mathcal{V}$. If $v_1 <:^c v_2$, noerr (v_1) , noerr (v_2) , and $\langle v_3, v_4 \rangle = \operatorname{prop}(v_1, v_2)$, then noerr (v_3) and noerr (v_4) .

Proof. By structural induction on v_1 and v_2 . The base cases are trivial. For the induction case where both v_1 and v_2 are tables, the argument follows by the i.h. \Box

A.2.1.9 Dataflow Invariant

We say a map M is data flow consistent subject to a program p if it satisfies the invariant $df_p(M)$ iff

 $fin(M) \wedge csub_p(M) \wedge errcons(M) \wedge tiocons(M) \wedge constcons_p(M) \wedge envcons(M)$

A.2.1.10 Constant Compatibility

We next introduce formally the notion of constant compatibility between two concrete values that intuitively states that the join of values does not introduce error values.

$$\operatorname{ccomp}(v_1, v_2) \iff (v_1 \in \{\bot, \omega\}) \lor (v_2 \in \{\bot, \omega\}) \lor (v_1 \in Cons \land v \in Cons \land v_1 = v_2)$$
$$\lor (v_1 \in \mathcal{T} \land v_2 \in \mathcal{T} \land \forall n_{\mathsf{cs}}. \operatorname{ccomp}(\pi_1(v_2(n_{\mathsf{cs}})), \pi_1(v_1(n_{\mathsf{cs}}))) \land$$
$$\operatorname{ccomp}(\pi_2(v_1(n_{\mathsf{cs}})), \pi_2(v_2(n_{\mathsf{cs}}))))$$

We also raise the definition of constant compatibility to a set of values V in the expected way $\operatorname{ccomp}(V) \iff \forall v_1, v_2 \in V.\operatorname{ccomp}(v_1, v_2)$. Note that constant compatibility is a reflexive relation.

Lemma 41. Let v be a concrete value. Then, ccomp(v, v).

Proof. By structural induction on v.

Next, we say that two maps M_1 and M_2 are constant compatible $\mathsf{ccomp}(M_1, M_2)$ iff $\forall n. \mathsf{ccomp}(M_1(n), M_2(n))$. Given a set of maps \hat{M} , we define its constant compatibility $\mathsf{ccomp}(\hat{M})$ as $\forall M_1 \in \hat{M}, M_2 \in \hat{M}. \mathsf{ccomp}(M_1, M_2)$.

Constant compatibility preserves error consistency of maps over joins.

Lemma 42. Let M_1 and M_2 be two concrete maps such that $ccomp(M_1, M_2)$, errcons (M_1) , and errcons (M_2) . Then, $errcons(M_1 \sqcup M_2)$.

Proof. By the definition of $\mathsf{ccomp}(v_1, v_2)$ on values v_1 and v_2 , the join $v_1 \sqcup v_2$ cannot introduce errors. That is, if $\mathsf{noerr}(v_1)$ and $\mathsf{noerr}(v_2)$, then $\mathsf{noerr}(v_1 \sqcup v_2)$. This can be easily shown by structural induction on v_1 and v_2 . The result then follows directly from the lemma assumptions.

A similar result can be shown in the other direction.

Lemma 43. Let M, M_1 , and M_2 be concrete maps such that $M = M_1 \dot{\sqcup} M_2$. If errcons(M) and $M \neq M_{\omega}$, then ccomp (M_1, M_2) , errcons (M_1) , and errcons (M_2) .

Proof. Similar to the proof of Lemma 42.

Clearly, the above result can be lifted to an arbitrary finite join of maps.

Lemma 44. Let \hat{M} be a finite set of concrete maps. If $\operatorname{errcons}(\sqcup \hat{M})$ and $\sqcup \hat{M} \neq M_{\omega}$, then $\operatorname{ccomp}(\hat{M})$ and $\forall M \in \hat{M}$. $\operatorname{errcons}(M)$.

Proof. By induction on |M|. The argument follows almost directly from Lemma 43.

We next show that constant compatibility is preserved by **prop** of value chains.

Lemma 45. Let $v_1 \in \mathcal{V}, v_2 \in \mathcal{V}, v_3 \in \mathcal{V}, v_1' \in \mathcal{V}, v_2' \in \mathcal{V}, v_2'' \in \mathcal{V}$, and $v_3' \in \mathcal{V}$. If $v_1 <:^c v_2, v_2 <:^c v_3, \langle v_1', v_2' \rangle = \operatorname{prop}(v_1, v_2), and \langle v_2'', v_3' \rangle = \operatorname{prop}(v_2, v_3),$ then $\operatorname{ccomp}(v_2', v_2'')$.

Proof. By structural induction on v_2 . We first cover the base cases.

• $v_2 = \bot$. Then, $v_3 = \bot$ and hence $v''_2 = \bot$ by the def. of prop.

- $v_2 = c$. By $<:^c, v_1 = c$ and $v_3 = c \lor v_2 = \bot$. Therefore, $v'_2 = v''_2 = c$.
- $v_2 = \omega$. By $\langle :^c, v_1 = \omega$ and $v_3 = \omega \lor v_2 = \bot$. Therefore, $v'_2 = v''_2 = \omega$.

Let us now consider the case when $v_2 \in \mathcal{T}$. By $\langle :^c$, it must be that $v_1 \in \mathcal{T}$. If $v_3 = \bot$, then $v_2'' = v_2$ and $v_2' = v_2 \sqsubseteq v_2'$ by Lemma 26. The result then follows trivially. Hence, we are left with the case where $v_3 \in \mathcal{T}$. Let us focus on any call site n_{cs} and let $v_1(n_{cs}) = \langle v_{1i}, v_{1o} \rangle$, $v_2(n_{cs}) = \langle v_{2i}, v_{2o} \rangle$, and $v_3(n_{cs}) = \langle v_{3i}, v_{3o} \rangle$. By the def. of prop, we have that $v_1' \in \mathcal{T}$, $v_2' \in \mathcal{T}$, $v_2'' \in \mathcal{T}$, and $v_3' \in \mathcal{T}$. Let then $v_1'(n_{cs}) = \langle v_{1i}', v_{1o}' \rangle$, $v_2(n_{cs}) = \langle v_{2i}', v_{2o}' \rangle$, and $v_3'(n_{cs}) = \langle v_{3i}', v_{3o}' \rangle$. By the def. of prop, we have to only consider the cases of whether $n_{cs} \in v_2$ and $n_{cs} \in v_3$.

- $n_{cs} \notin v_2, n_{cs} \notin v_3$. By the def. of prop, we have $v'_{2i} = v''_{2i} = v_{2i}$ and $v'_{2o} = v''_{2o} = v_{2o}$. The argument is then straightforward.
- $n_{cs} \notin v_2, n_{cs} \in v_3$. By the def. of prop, we have $v'_{2i} = v_{2i}$ and $v'_{2o} = v_{2o}$. By Lemma 26, we have $v_{2i} \sqsubseteq v''_{2i}$ and $v_{2o} \sqsubseteq v''_{2o}$. The argument then follows easily.
- $n_{cs} \in v_2, n_{cs} \notin v_3$. By the def. of prop, we have $v_{2i}'' = v_{2i}$ and $v_{2o}'' = v_{2o}$. By Lemma 26, we have $v_{2i} \sqsubseteq v_{2i}'$ and $v_{2o} \sqsubseteq v_{2o}'$. The argument then follows easily.
- $n_{cs} \in v_2, n_{cs} \in v_3$. First, by <:^c we have $v_{3i} <:^c v_{2i}$ and $v_{2i} <:^c v_{1i}$. Next, we also have $v_{1o} <:^c v_{2o}$ and $v_{2o} <:^c v_{3o}$. By the def. of prop, we next have $\langle v'_{2i}, v'_{1i} \rangle = \operatorname{prop}(v_{2i}, v_{1i}), \langle v'_{3i}, v''_{2i} \rangle = \operatorname{prop}(v_{3i}, v_{2i}), \langle v'_{1o}, v'_{2o} \rangle = \operatorname{prop}(v_{1o}, v_{2o}),$ and $\langle v''_{2o}, v''_{3o} \rangle = \operatorname{prop}(v_{2o}, v_{3o})$. By the i.h., we then have $\operatorname{ccomp}(v'_{2i}, v''_{2i})$ and $\operatorname{ccomp}(v'_{2o}, v''_{2o})$. The result then follows from the def. of ccomp.

A.2.2 Inductive Properties

Given a fixed program m, we define the family of inductive properties P, indexed by fuels k, as a set of triples consisting of locations in the program, concrete maps, and environments that is closed subject to the concrete transformer step. More precisely, given a program m and a fuel k, we say that $P_k \in \wp(Loc \times \mathcal{M} \times \mathcal{E})$ is an *inductive invariant* (of step) if

$$(\ell, M, E) \in P_k \implies (\ell, \mathsf{step}_k[\![e^\ell]\!](E)(M), E) \in P_k$$

where $e^{\ell} \in m$ and (e, E) is well-formed. We additionally require that all potential recursive calls of step for the subexpressions of e are closed subject to P, as follows:

•
$$e^{\ell} = e_1^{\ell_1} e_2^{\ell_2}$$
.

$$\frac{(\ell, M, E) \in P_{k+1}}{(\ell_1, M, E) \in P_k}$$
 $(\ell, M, E) \in P_{k+1}$ $(\ell_1, M_1, E) \in P_k$

$$M_1 = \operatorname{step}_k \llbracket e_1^{\ell_1} \rrbracket (E)(M) \quad M_1(E \diamond \ell_1) \notin \{\bot, \omega\}$$
 $(\ell_2, M_1, E) \in P_k$

•
$$e^{\ell} = (e_0^{\ell_0} ? e_1^{\ell_1} : e_2^{\ell_2})^{\ell}.$$

$$\frac{(\ell, M, E) \in P_{k+1}}{(\ell_0, M, E) \in P_k}$$

$$(\ell, M, E) \in P_{k+1} \qquad (\ell_0, M_0, E) \in P_k$$

$$M_0 = \operatorname{step}_k \llbracket e_0^{\ell_0} \rrbracket (E)(M) \qquad M_0(E \diamond \ell_0) = true$$

$$(\ell, M, E) \in P_{k+1} \qquad (\ell_0, M_0, E) \in P_k$$

$$M_0 = \operatorname{step}_k \llbracket e_0^{\ell_0} \rrbracket (E)(M) \qquad M_0(E \diamond \ell_0) = false$$

$$(\ell_2, M, E) \in P_k$$

•
$$e^{\ell} = (\mu f.\lambda x.e_1^{\ell_1})^{\ell}.$$

 $M(\ell \diamond E) = T$ $n_{cs} \in T$ $\pi_1(T(n_{cs})) \notin \{\bot, \omega\}$
 $n_x = E \diamond n_{cs} \diamond x$ $n_f = E \diamond n_{cs} \diamond f$ $E_1 = E.x: n_x.f: n_f$
 $\langle T'_x, T' \rangle = \operatorname{prop}([n_{cs} \mapsto \langle M(n_x), M(E_1 \diamond \ell_1) \rangle], T)$
 $\langle T'', T'_f \rangle = \operatorname{prop}(T, M(n_f))$ $\langle v'_x, v'_1 \rangle = T'x(n_{cs})$
 $M_1 = M[E \diamond \ell \mapsto T' \sqcup T'', n_x \mapsto v'_x, n_f \mapsto T'_f, E_1 \diamond \ell_1 \mapsto v'_1]$
 $(\ell, M, E) \in P_{k+1}$
 $(\ell_1, M_1, E_1) \in P_k$

Given a program e^{ℓ} , we define its *strongest inductive invariant* family I as the smallest inductive property of e^{ℓ} such that $(\ell, M_{\perp}, \emptyset) \in I_k$, for every k.

A.3 Path Data Flow Semantics

In this section, we introduce an alternative definition of the concrete data flow semantics that, besides data flow values, also keeps track of paths the data flow values take. This semantics can be thought of as accumulating ghost information of the data flow semantics. We will use this ghost information to show several useful properties of the original data flow semantics.

A.3.1 Semantic Domains

We define paths $p \in Path$ as a sequence of nodes \mathcal{N}^+ . The path data flow values are defined in a similar way as are the data flow values except that path information is maintained as well.

$w \in \mathcal{V}^{w} ::= \perp \mid \omega \mid \langle c, p \rangle \mid \langle T^{w}, p \rangle$	path (data flow) values
$T^{w} \in \mathcal{T}^{w} \stackrel{\text{\tiny def}}{=} Path \to \mathcal{V}^{w} \times \mathcal{V}^{w}$	path tables
$M^{w} \in \mathcal{M}^{w} \stackrel{\text{\tiny def}}{=} \mathcal{N} o \mathcal{V}^{w}$	path execution maps

The set $Cons^{\mathsf{w}} \stackrel{\text{def}}{=} Cons \times Path$ denotes the set of all constant path values that also Boolean path values $Bool^{\mathsf{w}} \stackrel{\text{def}}{=} \{true, false\} \times Path.$

The computational partial ordering on path values is defined similarly to the ordering on the ordinary the data flow values. In addition, we require that corresponding paths are related by the prefix relation \leq .

 $w_1 \sqsubseteq^{\mathsf{w}} w_2 \xleftarrow{}^{\mathrm{def}} w_1 = \bot \lor w_2 = \omega \lor$

$$(w_1 = \langle c, p_1 \rangle \land w_2 = \langle c, p_2 \rangle \land c \in Cons \land p_1, p_2 \in Path \land p_1 \preceq p_2) \lor w_1 = \langle T_1^{\mathsf{w}}, p_1 \rangle \land w_2 = \langle T_2^{\mathsf{w}}, p_2 \rangle \land T_1^{\mathsf{w}}, T_2^{\mathsf{w}} \in \mathcal{T}^{\mathsf{w}} \land p_1, p_2 \in Path \land p_1 \preceq p_2 \land \forall n_{\mathsf{cs}}, T_1^{\mathsf{w}}(n_{\mathsf{cs}}) \sqsubseteq w T_2^{\mathsf{w}}(n_{\mathsf{cs}})$$

The join on the path values is defined as follows.

$$\langle c, p_1 \rangle \sqcup^{\mathsf{w}} \langle c, p \rangle \stackrel{\text{def}}{=} \langle c, p \rangle$$
 if $p_1 \preceq p_2 \lor p_2 \preceq p_1$ and
 $p = p_1 \preceq p_2 ? p_2 : p_1$
 $\langle T_1^{\mathsf{w}}, p_1 \rangle \sqcup^{\mathsf{w}} \langle T_2^{\mathsf{w}}, p_2 \rangle \stackrel{\text{def}}{=} \langle \Lambda n_{\mathsf{cs}}. T_1^{\mathsf{cr}}(n_{\mathsf{cs}}) \sqcup^{\mathsf{w}} T_2^{\mathsf{cr}}(n_{\mathsf{cs}}), p \rangle$ if $p_1 \preceq p_2 \lor p_2 \preceq p_1$ and
 $p = p_1 \preceq p_2 ? p_2 : p_1$
 $\perp \sqcup^{\mathsf{w}} w \stackrel{\text{def}}{=} w \qquad w \sqcup^{\mathsf{w}} \perp \stackrel{\text{def}}{=} w \qquad w_1 \sqcup^{\mathsf{w}} w_2 \stackrel{\text{def}}{=} \omega$ (otherwise)

The meet on path values is defined similarly, as expected.

Lemma 46. The join \sqcup^{w} and \sqcap^{w} on path values are lub and glb operators subject to the \sqsubseteq^{w} ordering.

Proof. Similar to the proof of Lemma 1.

An empty table that maps every call site to the pair of \perp values is denoted by T^{w}_{\perp} . We reuse the notation for denoting singleton tables from the data flow semantics. The same applies for the notation $n_{\mathsf{cs}} \in T^{\mathsf{w}}$ denoting if some non- \perp input has been seen for the call site n_{cs} in a given table T^{w} . A path value map $M^{\mathsf{w}} \in \mathcal{M}^{\mathsf{w}} \stackrel{\text{def}}{=} \mathcal{N} \to \mathcal{V}^{\mathsf{w}}$ is a function from nodes to path values, as expected. Path maps M^{w}_{\perp} and M^{w}_{ω} assign to each node the \perp and ω value, respectively.

We also define a function paths that collects paths of a value

 $\mathsf{paths}(\bot) \stackrel{\text{def}}{=} \emptyset \qquad \mathsf{paths}(\omega) \stackrel{\text{def}}{=} \emptyset \qquad \mathsf{paths}(\langle c, p \rangle) \stackrel{\text{def}}{=} \{p\}$ $\mathsf{paths}(\langle T^{\mathsf{w}}, p \rangle) \stackrel{\text{def}}{=} \{p\} \cup \bigcup_{p' \in T^{\mathsf{w}}} \{p'\} \cup \mathsf{paths}(\pi_1(T^{\mathsf{w}}(p'))) \cup \mathsf{paths}(\pi_2(T^{\mathsf{w}}(p')))$

All paths $paths(M^w)$ of path map M^w are defined as the union of all paths of values in the map, as expected.

A.3.2 Data Propagation

The data propagation $\operatorname{prop}^{\mathsf{w}} : \mathcal{V}^{\mathsf{w}} \times \mathcal{V}^{\mathsf{w}} \to \mathcal{V}^{\mathsf{w}} \times \mathcal{V}^{\mathsf{w}}$ on path values closely resembles the data propagation on the plain data flow values. The main difference is that a meaningful propagation only happens in the case when the paths of the given values match by the \preceq relation. Otherwise, the propagation just sets both values to top ω . As we shall see shortly, such cases never happen during the computation.

$$\begin{aligned} & \operatorname{prop}^{\mathsf{w}}(\langle T_{1}^{\mathsf{w}}, p_{1} \rangle, \langle T_{2}^{\mathsf{w}}, p_{2} \rangle) \stackrel{\operatorname{def}}{=} & \operatorname{if} p_{1} \preceq p_{2} \\ & \operatorname{let} T' = \Lambda p. \\ & \operatorname{if} p \notin T_{2}^{\mathsf{w}} \operatorname{then} \langle T_{1}^{\mathsf{w}}(p), T_{2}^{\mathsf{w}}(p) \rangle \operatorname{else} \\ & \operatorname{let} \langle w_{1i}, w_{1o} \rangle = T_{1}^{\mathsf{w}}(n_{\mathrm{cs}}); \langle w_{2i}, w_{2o} \rangle = T_{2}^{\mathsf{w}}(n_{\mathrm{cs}}) \\ & \langle w_{2i}', w_{1i}' \rangle = \operatorname{prop}^{\mathsf{w}}(w_{2i}, w_{1i}); \langle w_{1o}', w_{2o}' \rangle = \operatorname{prop}^{\mathsf{w}}(w_{1o}, w_{2o}) \\ & \operatorname{in} (\langle w_{1i}', w_{1o}' \rangle, \langle w_{2i}', w_{2o}' \rangle) \\ & \operatorname{in} \langle \Lambda p. \ \pi_{1}(T'(p)), \Lambda p. \ \pi_{2}(T'(p)) \rangle \\ & \operatorname{prop}^{\mathsf{w}}(\langle T^{\mathsf{w}}, p \rangle, \bot) \stackrel{\mathrm{def}}{=} (\langle T^{\mathsf{w}}, p \rangle, \langle T_{\perp}^{\mathsf{w}}, p \rangle) \quad \operatorname{prop}^{\mathsf{w}}(\langle T^{\mathsf{w}}, p \rangle, \omega) \stackrel{\mathrm{def}}{=} \langle \omega, \omega \rangle \\ & \operatorname{prop}^{\mathsf{w}}(\langle c, p_{1} \rangle, \langle c, p_{2} \rangle) \stackrel{\mathrm{def}}{=} (\langle c, p_{1} \rangle, \langle c, p_{2} \rangle) \quad & \operatorname{if} \ p_{1} \preceq p_{2} \\ & \operatorname{prop}^{\mathsf{w}}(\langle c, p_{1} \rangle, \langle c, p_{2} \rangle) \stackrel{\mathrm{def}}{=} \langle \omega, \omega \rangle & \operatorname{if} \ p_{1} \preceq p_{2} \\ & \operatorname{prop}^{\mathsf{w}}(\langle c, p_{1} \rangle, \langle c, p_{2} \rangle) \stackrel{\mathrm{def}}{=} \langle w_{1}, w_{1} \sqcup^{\mathsf{w}} w_{2} \rangle \quad & (\operatorname{otherwise}) \end{aligned}$$

Lemma 47. prop^w is increasing and monotone.

Proof. By structural induction on the arguments of prop^w. The argument for the increasing property is rather straightforward. Regarding monotonicity, arguably the most interesting case is that of constants. Suppose $w_1 = \langle c, p_1 \rangle$, $w_2 = \langle c, p_2 \rangle$,

 $w_3 = \langle c, p_3 \rangle$, and $w_4 = \langle c, p_4 \rangle$. Let $w_1 \sqsubseteq^w w_3$ and $w_2 \sqsubseteq^w w_4$. First, if $p_1 \not\leq p_2$ then $p_3 \not\leq p_4$ since by the definition of \sqsubseteq^w it must be that $p_1 \preceq p_3$ and $p_2 \preceq p_4$. In this case, the argument is trivial. If $p_1 \preceq p_2$ the argument is again simple as the values w_1 and w_2 do not get updated. Assume now that $w_2 = \bot$ and $w_4 = \bot$. The argument is again straightforward. Lastly, assume that only $w_2 = \bot$ (note that we cannot have the case when only $w_4 = \bot$ since we assume for the purpose of showing monotonicity that $w_2 \sqsubseteq^w w_4$). The propagation will result in an updated value $w'_2 = w_1$. If $p_3 \not\sqsubseteq^w p_4$, the argument is easy. Otherwise, since $p_1 \preceq p_3$ and $p_3 \preceq p_4$, then it also must be that $p_1 \preceq p_4$. The same reasoning, together with the induction hypothesis, applies to the case of tables. \Box

It can be also shown that the propagation does not invent new paths. That is, propagation between two values results in two values where each resulting value is either top ω or a new path values that has the same path as the original path value.

Lemma 48. Let $\langle w'_1, w'_2 \rangle = \operatorname{prop}^{\mathsf{w}}(w_1, w_2)$ for some path values w_1, w'_1, w_2 , and w'_2 . Then, if $w_1 = \langle -, p \rangle$ for some path p, then $w'_1 = \omega$ or $w'_2 = \langle -, p \rangle$. The same hold for w_2 and w'_2 .

Proof. The proof goes by structural induction on both w_1 and w_2 .

A.3.3 Transformer

The transformer $\operatorname{step}^{\mathsf{w}} : \mathbb{N} \to \lambda^d \to \mathcal{E} \to \mathcal{M}^{\mathsf{w}} \to \mathcal{M}^{\mathsf{w}}$ of the paths semantics works similarly as the transformer for the original data flow semantics. The only difference is that the path information needs to be properly handled when creating new values. For that purpose, we introduce a function for extending path information in path values with nodes.

$$ext(p,n) \stackrel{\text{def}}{=} p \quad \text{if } top(p) = n \quad ext(p,n) \stackrel{\text{def}}{=} p \cdot n \quad (\text{otherwise})$$
$$ext(\langle T^{\mathsf{w}}, p \rangle, n) \stackrel{\text{def}}{=} \langle T^{\mathsf{w}}, ext(p,n) \rangle \quad ext(\langle c, p \rangle, n) \stackrel{\text{def}}{=} \langle c, ext(p,n) \rangle$$
$$ext(\omega, n) \stackrel{\text{def}}{=} \omega \quad ext(\bot, n) \stackrel{\text{def}}{=} \bot$$

Here, we use **top** to denote the top/final node of a path. Further, \cdot operation concatenates two given paths. Note that we also consider a node as a path. Finally, observe that the operation **ext** on paths is idempotent. Also, let **path** be a partial function on path values that returns a path associated with the value, if such a path exists. Note that path semantics models call sites as paths rather than nodes, but the variable nodes are created only using the top of the call site path. As we shall see later, the actual paths observed during the computation that have unique top nodes.

Lemma 49. step^w is increasing and monotone.

Proof. The increasing property follows from the increasing property of $prop^w$ (Lemma 47). The proof of monotonicity closely follows that of Lemma 28.

A.3.4 Additional Properties

We next introduce several useful properties related to the path semantics.

A.3.4.1 Computation Range

The results regarding the computation range are the same as for the ordinary data flow semantics.

Lemma 50. Let M^{w} and M_0^{w} be path maps, E an environment, e^{ℓ} an expression

 $\operatorname{step}_{0}^{\mathsf{w}} \llbracket e \rrbracket (E)(M^{\mathsf{w}}) \stackrel{\text{def}}{=} M^{\mathsf{w}}$ $\mathsf{step}_{k+1}^{\mathsf{w}}[c^{\ell}](E)(M^{\mathsf{w}}) \stackrel{\text{def}}{=} M^{\mathsf{w}}[E \diamond \ell \to M^{\mathsf{w}}(E \diamond \ell) \sqcup \langle c, E \diamond \ell \rangle]$ $\operatorname{step}_{k+1}^{\mathsf{w}} \llbracket x^{\ell} \rrbracket (E) (M^{\mathsf{w}}) \stackrel{\text{def}}{=}$ let $\langle w'_x, w' \rangle = \operatorname{prop}^{\mathsf{w}}(M^{\mathsf{w}}(E(x)), M^{\mathsf{w}}(E \diamond \ell))$ in $M^{\mathsf{w}}[E(x) \mapsto w'_{x}, E \diamond \ell \mapsto \mathsf{ext}(w', E \diamond \ell)]$ $step_{k+1}^{w} [\![(e_1^{\ell_1} e_2^{\ell_2})^{\ell}]\!](E)(M^{w}) \stackrel{\text{def}}{=}$ let $M_1^{w} = \operatorname{step}_k^{w} \llbracket e_1 \rrbracket (E)(M^{w}); w_1 =_{M_1^{w}} M_1^{w}(E \diamond \ell_1)$ in if $w_1 \notin \mathcal{T}^w$ then M^w_ω else let $M_2^{w} = \operatorname{step}_k^{w} \llbracket e_2 \rrbracket(E)(M_1^{w}); w_2 =_{M_2^{w}} M_2^{w}(E \diamond \ell_2)$ in let $\langle w'_1, T^{\mathsf{w}} \rangle = \mathsf{prop}^{\mathsf{w}}(w_1, [\mathsf{path}(w_1) \mapsto \langle w_2, M_2^{\mathsf{w}}(E \diamond \ell) \rangle])$ $\langle w_2', w' \rangle = T^{\mathsf{w}}(\mathsf{path}(w_1))$ in $M_2^{\mathsf{w}}[E \diamond \ell_1 \mapsto w_1', E \diamond \ell_2 \mapsto w_2', E \diamond \ell \mapsto \mathsf{ext}(w', E \diamond \ell)]$ $\operatorname{step}_{k+1}^{\mathsf{w}} \llbracket (\mu f.\lambda x.e_1^{\ell_1})^{\ell} \rrbracket (E)(M^{\mathsf{w}}) \stackrel{\text{def}}{=}$ let $T_w =_{M^w} \langle T^w_{\perp}, E \diamond \ell \rangle \sqcup^w M^w(E \diamond \ell)$ in let $\overline{M^{\mathsf{w}}} = \Lambda p \in T_w$. if $\pi_1(T^{\mathsf{w}}(n_{\mathsf{cs}})) = \omega$ then M_{ω}^{w} else let $n_x = E \diamond top(p) \diamond x$; $n_f = E \diamond top(p) \diamond f$; $E_1 = E[x \mapsto n_x, f \mapsto n_f]$ $\langle T_x^{\mathsf{w}}, T_0^{\mathsf{w}} \rangle = \mathsf{prop}^{\mathsf{w}}([p \mapsto \langle M^{\mathsf{w}}(n_x), M^{\mathsf{w}}(E_1 \diamond \ell_1) \rangle], T_w)$ $\langle T_2^{\mathsf{w}}, T_f^{\mathsf{w}} \rangle = \mathsf{prop}^{\mathsf{w}}(T_w, M^{\mathsf{w}}(n_f)); \langle w'_x, w'_1 \rangle = T_x^{\mathsf{w}}(p)$ $M_1^{\mathsf{w}} = M^{\mathsf{w}}[E \diamond \ell \mapsto T_1^{\mathsf{w}} \sqcup^{\mathsf{w}} T_2^{\mathsf{w}}, n_f \mapsto \mathsf{ext}(T_f^{\mathsf{w}}, n_f), n_x \mapsto \mathsf{ext}(w'_x, n_x), E_1 \diamond \ell_1 \mapsto w'_1]$ in step_k^w $[\![e_1]\!](E_1)(M_1^w)$ in $M^{\mathsf{w}}[E \diamond \ell \mapsto T^{\mathsf{w}}] \stackrel{\cdot}{\sqcup}^{\mathsf{w}} \stackrel{\cdot}{\coprod}^{\mathsf{w}}_{n \in T^{\mathsf{w}}} \overline{M}(p)$ $\operatorname{step}_{k+1}^{\mathsf{w}} \llbracket (e_0^{\ell_0} ? e_1^{\ell_1} : e_2^{\ell_2})^{\ell} \rrbracket (E)(M^{\mathsf{w}}) \stackrel{\text{def}}{=}$ let $M_0^w = \text{step}_k^w [\![e_0]\!](E)(M^w); w_0 =_{M_0^w} M_0^w(E \diamond \ell_0)$ in if $w_0 \notin Bool^w$ then M^w_{ω} else let $b = if w_0 = \langle true, \rangle$ then 1 else 2 in let $M_b^{\mathsf{w}} = \operatorname{step}_k^{\mathsf{w}} \llbracket e_b \rrbracket(E)(M_0^{\mathsf{w}}); w_b =_{M_b^{\mathsf{w}}} M_b^{\mathsf{w}}(E \diamond \ell_b)$ in let $\langle w'_{b}, w' \rangle = \operatorname{prop}^{\mathsf{w}}(w_{b}, M^{\mathsf{w}}_{b}(E \diamond \ell))$ in $M_{h}^{\mathsf{w}}[E \diamond \ell \mapsto \mathsf{ext}(w', E \diamond \ell), E \diamond \ell_{h} \mapsto w_{h}']$

of a program e' where (e, E) is well-formed, and k a fuel. If domvalid (ℓ, e', E) and $\operatorname{step}_{k}^{\mathsf{w}}[\![e^{\ell}]\!](E)(M^{\mathsf{w}}) = M_{0}^{\mathsf{w}}$ then $\forall n.M^{\mathsf{w}}(n) \neq M_{0}^{\mathsf{w}}(n) \implies n \in \operatorname{ranger}_{e'}(E \diamond \ell).$ *Proof.* Similar to the proof of Lemma 35.

Lemma 51. Let E be an environment, e' be a program, ℓ_1 and ℓ_2 two locations such that $Loc(e'(\ell_1)) \cap Loc(e'(\ell_2)) = \emptyset$. Then, $range_{e'}(E \diamond \ell_1) \cap range_{e'}(E \diamond \ell_2) = rng(E)$.

Proof. Similar to the proof of Lemma 36.

Lemma 52. Let E_1 and E_2 be two environments and ℓ a location of a program e'. If dom $(E_1) = dom(E_2)$ and $E_1 \neq E_2$, then range_{e'} $(E \diamond \ell_1) \cap range_{e'}(E \diamond \ell_2) = rng(E_1) \cap rng(E_2)$.

Proof. Similar to the proof of Lemma 37.

A.3.4.2 Finiteness

We say a value w is finite fin(w) if it is either a ground value or a table that has seen only a finite number of inputs, just as in the case of regular data flow semantics.

$$\begin{aligned} \operatorname{fin}(w) \stackrel{\text{\tiny def}}{=} (w \in \{\bot, \omega\} \cup \operatorname{Cons}^{\mathsf{w}}) \lor (\exists T^{\mathsf{w}}. w = \langle T^{\mathsf{w}}, \rangle \in \mathcal{T}^{\mathsf{w}} \land |\{p \mid p \in T^{\mathsf{w}}\}| \in \mathbb{N} \land \\ \forall p \in T^{\mathsf{w}}. \operatorname{fin}(\pi_1(T^{\mathsf{w}}(p))) \land \operatorname{fin}(\pi_2(T^{\mathsf{w}}(p)))) \end{aligned}$$

As expected, $prop^w$ does not affect the finiteness of values. The same clearly holds for joins \sqcup .

Lemma 53. Let w_1 , w_2 , w_3 , and w_4 be path values. If $fin(w_1)$, $fin(w_2)$, and $\langle w_3, w_4 \rangle = prop^w(w_1, w_2)$, then $fin(w_3)$ and $fin(w_4)$.

Proof. Similar to the proof of Lemma 30.

We say a path map M^{w} is finite $\operatorname{fin}(M^{\mathsf{w}})$ if all of its values are finite and the set of reachable nodes is finite as well $|\{n \mid M^{\mathsf{w}}(n) \neq \bot\}| \in \mathbb{N}$. It can be shown that $\operatorname{step}^{\mathsf{w}}$ preserves the finiteness of maps.

Lemma 54. Let M^{w} and M_0^{w} be maps, E an environment, e^{ℓ} an expression where (e, E) is well-formed, and k a fuel. If $\operatorname{fin}(M^{\mathsf{w}})$ and $\operatorname{step}_k^{\mathsf{w}} \llbracket e^{\ell} \rrbracket (E)(M^{\mathsf{w}}) = M_0^{\mathsf{w}}$, then $\operatorname{fin}(M_0^{\mathsf{w}})$.

Proof. Similar to the proof of Lemma 31.

A.3.4.3 Top of the Path Consistency

We say a path map M^{w} is consistent subject to the top of the paths $\mathsf{tpcons}(M^{\mathsf{w}})$ iff for every node n it is the case that $M^{\mathsf{w}}(n) \in \{\bot, \omega\}$ or $M^{\mathsf{w}}(n) = \langle_, p\rangle$ for some path p where $\mathsf{path}(p) = n$.

The consistency of maps subject to the top of paths is an preserved by the path semantics transformer.

Lemma 55. Let $M^{\mathsf{w}} \in \mathcal{M}^{\mathsf{w}}$ where $\operatorname{tpcons}(M^{\mathsf{w}})$, $M_0^{\mathsf{w}} \in \mathcal{M}^{\mathsf{w}}$, e^{ℓ} an expression, $k \in \mathbb{N}$, and E an environment where (e, E) is well-formed. If $\operatorname{step}_k^{\mathsf{w}} \llbracket e^{\ell} \rrbracket (E)(M^{\mathsf{w}}) = M_0^{\mathsf{w}}$, then $\operatorname{tpcons}(M_0^{\mathsf{w}})$.

Proof. The proof goes by induction on k and e. The argument follows from the basic properties of the join on path values, Lemma 48, and the idempotency of ext operation.

Top of the path consistency can be used to establish an interesting and useful property of the concrete transformer $step^w$. Once path has been inferred for some value in the map, then after applying the transformer, the updated value is either top ω or a new path value but with the same associated path.

Lemma 56. Let $n \in \mathcal{N}$, $M^{\mathsf{w}} \in \mathcal{M}^{\mathsf{w}}$ where $\mathsf{tpcons}(M^{\mathsf{w}})$, $M_0^{\mathsf{w}} \in \mathcal{M}^{\mathsf{w}}$, e^{ℓ} an expression, $k \in \mathbb{N}$, and E an environment where (e, E) is well-formed. Suppose

 $\operatorname{step}_{k}^{\mathsf{w}}\llbracket e^{\ell} \rrbracket(E)(M^{\mathsf{w}}) = M_{0}^{\mathsf{w}}. \text{ If } M^{\mathsf{w}}(n) = \langle -, p \rangle \text{ for some path } p, \text{ then } M^{\mathsf{w}}(n_{0}) = \langle -, p \rangle \text{ or } M_{0}^{\mathsf{w}}(n) = \omega.$

Proof. The proof goes by induction on k and e. The argument follows from the basic properties of the join on path values and Lemma 48. The top of the path consistency property together with Lemma 55 is needed to prove that the calls to ext on existing paths do not create new paths.

A.3.4.4 Call Site Path consistency

We next define the notion of call site path consistency intuitively capturing the idea that a table can be called at call site path that refers to the future points reached by the table, i.e., call site paths must be an extension of the path associated with the table. A path value w is call site path consistent pcscons(w) iff

$$(w \in \{\perp, \omega\} \cup Cons^{\mathsf{w}}) \lor$$

 $\exists T^{\mathsf{w}}, p.w = \langle T^{\mathsf{w}}, p \rangle \land \forall p' \in T^{\mathsf{w}}. p \leq p' \land \mathsf{pcscons}(\pi_1(T^{\mathsf{w}}(p'))) \land \mathsf{pcscons}(\pi_2(T^{\mathsf{w}}(p'))) \\ \text{We raise the definition of call site path consistency to the maps. That is, a path map <math>M^{\mathsf{w}}$ is said to be call site path consistent $\mathsf{pcscons}(M^{\mathsf{w}})$ iff for every node n we have $\mathsf{pcscons}(M^{\mathsf{w}}(n))$.

It can be shown that the data propagation does not invalidate the call site path consistency.

Lemma 57. Let $\langle w'_1, w'_2 \rangle = \mathsf{prop}^{\mathsf{w}}(w_1, w_2)$ for path values w_1, w'_1, w_2 , and w'_2 . If $\mathsf{pcscons}(w_1)$ and $\mathsf{pcscons}(w_2)$, then $\mathsf{pcscons}(w'_1)$ and $\mathsf{pcscons}(w'_2)$.

Proof. The proof goes by structural induction on both w_1 and w_2 . The base cases are trivial and so is the the induction case when both w_1 and w_2 are path tables but the path associated with w_1 is not a prefix of the one associated with w_2 . Otherwise, the argument follows from the i.h. and the fact that any call site path seen in w_2 has to also be an extension of the path associated with w_1 .

The call site path consistency is preserved by the transformer of the path semantics.

Lemma 58. Let $M^{\mathsf{w}} \in \mathcal{M}^{\mathsf{w}}$ where $\mathsf{pcscons}(M^{\mathsf{w}})$, $M_0^{\mathsf{w}} \in \mathcal{M}^{\mathsf{w}}$, e^{ℓ} an expression, $k \in \mathbb{N}$, and E an environment where (e, E) is well-formed. If $\mathsf{step}_k^{\mathsf{w}} \llbracket e^{\ell} \rrbracket (E)(M^{\mathsf{w}}) = M_0^{\mathsf{w}}$, then $\mathsf{pcscons}(M_0^{\mathsf{w}})$.

Proof. The proof goes by induction on k and e. The argument follows from the basic properties of the join on path values, Lemma 48, Lemma 57, and Lemma 56.

A.3.4.5 Root Path Consistency

A path p is called to be root consistent $rootcons_e(p)$, subject to some program e, iff its root node corresponds to a constant or lambda expression of the program.

 $\operatorname{rootcons}_e(p) \iff \exists c.e(\operatorname{loc}(\operatorname{root}(p))) = c \lor \exists f, x, e'.e(\operatorname{loc}(\operatorname{root}(p))) = \mu f.\lambda x.e'$ Note that we use $\operatorname{root}(p)$ to denote the starting/root node of a path p. A map M^w is root consistent if all of its paths in $\operatorname{paths}(M^w)$ are root consistent.

Lemma 59. Let e be a program, $\langle w'_1, w'_2 \rangle = \operatorname{prop}^{\mathsf{w}}(w_1, w_2)$ for path values w_1, w'_1, w_2 , and w'_2 . If $\operatorname{rootcons}_e(w_1)$ and $\operatorname{rootcons}_e(w_2)$, then $\operatorname{rootcons}(w'_1)$ and $\operatorname{rootcons}(w'_2)$. *Proof.* The proof goes by structural induction on both w_1 and w_2 .

The root path consistency is preserved by the transformer of the path semantics. **Lemma 60.** Let $M^{\mathsf{w}} \in \mathcal{M}^{\mathsf{w}}$ where $\mathsf{pcscons}(M^{\mathsf{w}})$, $M_0^{\mathsf{w}} \in \mathcal{M}^{\mathsf{w}}$, e^{ℓ} an expression, $k \in \mathbb{N}$, and E an environment where (e, E) is well-formed. If $\mathsf{step}_k^{\mathsf{w}}[\![e^{\ell}]\!](E)(M^{\mathsf{w}}) = M_0^{\mathsf{w}}$, then $\mathsf{pcscons}(M_0^{\mathsf{w}})$. *Proof.* The proof goes by induction on k and e. The argument follows from the basic properties of the join on path values, Lemma 48, and Lemma 59.

A.3.4.6 Table IO Consistency

We next formalize the notion of input-output consistency for tables that is the same as for the standard data flow semantics.

$$\operatorname{tiocons}(w) \stackrel{\text{def}}{\longleftrightarrow} (w \in \{\bot, \omega\} \cup Cons^{\mathsf{w}}) \lor (\exists T^{\mathsf{w}}. w = \langle T^{\mathsf{w}}, \lrcorner\rangle \land \forall p \in T^{\mathsf{w}}.$$
$$\pi_1(T^{\mathsf{w}}(p)) = \bot \implies \pi_2(T^{\mathsf{w}}(p)) = \bot \land$$
$$\operatorname{tiocons}(\pi_1(T^{\mathsf{w}}(p))) \land \operatorname{tiocons}(\pi_2(T^{\mathsf{w}}(p))))$$

As expected, a map is input-output consistent $tiocons_e(M)$, subject to a program e, iff all of its values are IO consistent and for every node n

$$M^{\mathsf{w}}(env(n)\diamond\ell_1) = \bot \Longrightarrow M^{\mathsf{w}}(env(n)\diamond\ell_2) = \bot \land$$
$$M^{\mathsf{w}}(env(n)\diamond\ell_2) = \bot \Longrightarrow M^{\mathsf{w}}(env(n)\diamond\ell) = \bot$$

if
$$\exists e_1, e_2, \ell_1, \ell_2. e(loc(n)) = e_1^{\ell_1} e_2^{\ell_2}$$

$$\pi_1(M^{\mathsf{w}}(n)(p)) = \bot \Longrightarrow M^{\mathsf{w}}(n_f) = \bot \land$$
$$\pi_1(M^{\mathsf{w}}(n)(p)) = \bot \Longrightarrow M^{\mathsf{w}}(n_x) = \bot \land$$
$$\pi_1(M^{\mathsf{w}}(n_x)) = \bot \Longrightarrow M^{\mathsf{w}}(E_1 \diamond \ell_1) = \bot$$

if $M(n) \in \mathcal{T}^{\mathsf{w}}$ and $\exists x, f, e_1, \ell_1, n_x, n_f, p. \quad p \in M^{\mathsf{w}}(n) \land n_x = env(n) \diamond \mathsf{top}(p) \diamond x$ $\land n_f = env(n) \diamond \mathsf{top}(p) \diamond f \land E_1 = E.xn_x.f:n_f$

The input-output consistency of values is not violated by data propagation.

Lemma 61. Let e be a program, $\langle w'_1, w'_2 \rangle = \operatorname{prop}^{\mathsf{w}}(w_1, w_2)$ for path values w_1, w'_1, w_2 , and w'_2 . If $\operatorname{tiocons}_e(w_1)$ and $\operatorname{tiocons}_e(w_2)$, then $\operatorname{tiocons}(w'_1)$ and $\operatorname{tiocons}(w'_2)$.

Proof. The proof goes by structural induction on both w_1 and w_2 .

The input-output consistency is preserved by **step**^w.

Lemma 62. Let $M^{\mathsf{w}} \in \mathcal{M}^{\mathsf{w}}$ where $\operatorname{tiocons}(M^{\mathsf{w}})$, $M_0^{\mathsf{w}} \in \mathcal{M}^{\mathsf{w}}$, e^{ℓ} an expression, $k \in \mathbb{N}$, and E an environment where (e, E) is well-formed. If $\operatorname{step}_k^{\mathsf{w}} \llbracket e^{\ell} \rrbracket (E)(M^{\mathsf{w}}) = M_0^{\mathsf{w}}$, then $\operatorname{tiocons}(M_0^{\mathsf{w}})$.

Proof. The proof goes by induction on k and e. The argument follows from the fact that the join on path values does not invalidate input output consistency and Lemma 61.

A.3.4.7 Table Input Consistency

A path value w is input consistent $\operatorname{inpcons}_{e}(w)$, subject to some program e, iff the path associated with every input of a table has to come from a function application expression in e.

$$\begin{split} \mathsf{inpcons}_e(w) & \iff (w \in \{\bot, \omega\} \cup \mathit{Cons}^\mathsf{w}) \lor \exists T^\mathsf{w}. w = \langle T^\mathsf{w}, \lrcorner \rangle \land \forall p \in T^\mathsf{w}.\\ \mathsf{inpcons}_e(\pi_1(T^\mathsf{w}(p))) \land \mathsf{inpcons}_e(\pi_2(T^\mathsf{w}(p))) \land \exists e_1, e_2, \ell_2, p'.\\ e_1^{\mathit{loc}(\mathsf{top}(p))} e_2^{\ell_2} \in e \land (\pi_1(T^\mathsf{w}(p)) \in \{\bot, \omega\} \lor \pi_1(T^\mathsf{w}(p)) = \langle \lrcorner, p' \rangle\\ \land \mathsf{top}(p') = env(\mathsf{top}(p)) \diamond \ell_2) \end{split}$$

A map M^{w} is input consistent $\mathsf{inpcons}_{e}(M^{\mathsf{w}})$ iff all of its values are input consistent.

Lemma 63. Let e be a program, $\langle w'_1, w'_2 \rangle = \operatorname{prop}^{\mathsf{w}}(w_1, w_2)$ for path values w_1, w'_1, w_2 , and w'_2 . If $\operatorname{inpcons}_e(w_1)$ and $\operatorname{inpcons}_e(w_2)$, then $\operatorname{inpcons}(w'_1)$ and $\operatorname{inpcons}(w'_2)$.

Proof. The proof goes by structural induction on both w_1 and w_2 and it relies on Lemma 48.

The input consistency of path maps is preserved by step^w.
Lemma 64. Let $M^{\mathsf{w}} \in \mathcal{M}^{\mathsf{w}}$ where $\operatorname{inpcons}(M^{\mathsf{w}})$, $M_0^{\mathsf{w}} \in \mathcal{M}^{\mathsf{w}}$, e^{ℓ} an expression, $k \in \mathbb{N}$, and E an environment where (e, E) is well-formed. If $\operatorname{step}_k^{\mathsf{w}} \llbracket e^{\ell} \rrbracket (E)(M^{\mathsf{w}}) = M_0^{\mathsf{w}}$, then $\operatorname{inpcons}(M_0^{\mathsf{w}})$.

Proof. The proof goes by induction on k and e. The argument follows from the fact that join on paths values does not invalidate input consistency and Lemma 63. \Box

A.3.4.8 Table Output Consistency

A path value w is output consistent $\mathsf{outcons}_e(w)$, subject to some program e, iff the path associated with every input of a table has to come from a function application expression in e.

$$\begin{aligned} \mathsf{outcons}_e(w) &\iff (w \in \{\bot, \omega\} \cup Cons^\mathsf{w}) \lor \exists T^\mathsf{w}, p, f, x, e_1, \ell_1. w = \langle T^\mathsf{w}, p \rangle \land \\ &e(loc(\mathsf{root}(p))) = \mu f. \lambda x. e_1^{\ell_1} \land \forall p' \in T^\mathsf{w}. \mathsf{outcons}_e(\pi_1(T^\mathsf{w}(p'))) \land \\ &\mathsf{outcons}_e(\pi_2(T^\mathsf{w}(p'))) \land (\exists E, p''. \pi_2(T^\mathsf{w}(p')) \in \{\bot, \omega\} \lor \\ &\pi_2(T^\mathsf{w}(p')) = \langle_{-}, p'' \rangle \land \mathsf{top}(p'') = E \diamond \ell_1 \land E = \\ &env(\mathsf{root}(p)).x: env(\mathsf{root}(p)) \diamond \mathsf{top}(p') \diamond x. f: env(\mathsf{root}(p)) \diamond \mathsf{top}(p') \diamond f) \end{aligned}$$

A map M^{w} is output consistent $\mathsf{outcons}_e(M^{\mathsf{w}})$ iff all of its values are output consistent.

Lemma 65. Let e be a program, $\langle w'_1, w'_2 \rangle = \operatorname{prop}^{\mathsf{w}}(w_1, w_2)$ for path values w_1, w'_1, w_2 , and w'_2 . If $\operatorname{outcons}_e(w_1)$ and $\operatorname{outcons}_e(w_2)$, then $\operatorname{outcons}(w'_1)$ and $\operatorname{outcons}(w'_2)$.

Proof. The proof goes by structural induction on both w_1 and w_2 .

The output consistency of path maps is preserved by step^w.

Lemma 66. Let $M^{\mathsf{w}} \in \mathcal{M}^{\mathsf{w}}$ where $\operatorname{outcons}(M^{\mathsf{w}})$, $M_0^{\mathsf{w}} \in \mathcal{M}^{\mathsf{w}}$, e^{ℓ} an expression, $k \in \mathbb{N}$, and E an environment where (e, E) is well-formed. If $\operatorname{step}_k^{\mathsf{w}} \llbracket e^{\ell} \rrbracket (E)(M^{\mathsf{w}}) = M_0^{\mathsf{w}}$, then $\operatorname{outcons}(M_0^{\mathsf{w}})$. *Proof.* The proof goes by induction on k and e. The argument follows from the fact that join on paths values does not invalidate input consistency and Lemma 63. \Box

A.3.4.9 Constant Consistency

We next formalize the notion of constant consistency for path maps that is identical to the notion of constant consistency for regular maps. A path map M^{w} is constant consistent $\mathsf{constcons}_e(M^{\mathsf{w}})$, subject to a program e, iff for every node n, where $M^{\mathsf{w}}(n) = w$ for some path value w and $loc(n) = \ell$ for some location ℓ , it must be

$$w \in \{\omega, \bot, \langle c, n \rangle\} \qquad \qquad \text{if } e(\ell) = c$$

 $w \in \{\omega, \bot\} \cup \{\langle T^{\mathsf{w}}, E \diamond \ell \rangle \mid T^{\mathsf{w}} \in \mathcal{T}^{\mathsf{w}}\} \text{ if } e(\ell) = \mu f \cdot \lambda x \cdot e_1 \text{ for some } x, f, \text{ and } e_1$

A.3.4.10 Concrete Path Subtyping

The notion of subtyping for path values resembles the one on the regular data flow values. Additionally, we also require that a table cannot see a call site path (coming from the future) unless the future table has seen it as well.

$$w_{1} <: {}^{\mathsf{w}}w_{2} \iff (w_{2} = \bot) \lor (w_{1} = \langle c, p_{1} \rangle \land w_{2} = \langle c, p_{2} \rangle \land c \in Cons \land p_{1}, p_{2} \in Path)$$
$$\lor (w_{1} = w_{2} = \omega) \lor (w_{1} = \langle T_{1}^{\mathsf{w}}, p_{1} \rangle \land w_{2} = \langle T_{2}^{\mathsf{w}}, p_{2} \rangle \land T_{1}^{\mathsf{w}}, T_{2}^{\mathsf{w}} \in \mathcal{T}^{\mathsf{w}}$$
$$\land p_{1}, p_{2} \in Path \land p_{1} \preceq p_{2} \land (\forall p \in T_{1}^{\mathsf{w}}, p_{2} \preceq p \Longrightarrow p \in T_{2}^{\mathsf{w}}) \land$$
$$\forall p \in T_{2}^{\mathsf{w}}.\pi_{1}(T_{2}^{\mathsf{w}}(p)) <: {}^{\mathsf{w}}\pi_{1}(T_{1}^{\mathsf{w}}(p)) \land \pi_{2}(T_{1}^{\mathsf{w}}(p)) <: {}^{\mathsf{w}}\pi_{2}(T_{2}^{\mathsf{w}}(p)))$$

Unlike the subtyping relation on the ordinary data flow values, the above subtyping relation for path values is transitive, given certain simple assumptions on the values.

Lemma 67. Let $w_1 \in \mathcal{V}^w$, $w_2 \in \mathcal{V}^w$, and $w_3 \in \mathcal{V}^w$ such that $tiocons(w_1)$, $tiocons(w_2)$, $tiocons(w_3)$, $pcscons(w_1)$, $pcscons(w_2)$, and $pcscons(w_3)$. If $w_1 <: w_2$ and $w_2 <:^{w} w_3$, then $w_1 <:^{w} w_3$.

Proof. By structural induction on w_2 . We first cover the base cases.

• $w_2 = \bot$ $w_3 = \bot$ by def. of $<:^{w}$ $w_1 <:^{w} \bot$ • $w_2 = \omega$ $w_3 = \bot \lor w_3 = \bot$ by def. of $<:^{w}$ $w_1 = \omega$ by def. of $<:^{w}$ $w < :^{w} \omega \land \omega <:^{w} \bot$

$$w_1 = \langle c, p_1 \rangle, p_1 \preceq p_2 \qquad \text{by def. of } <:^{\mathsf{w}}$$
$$w_3 = \bot \lor w_3 = \langle c, p_3 \rangle \land p_2 \preceq p_3 \qquad \text{by def. of } <:^{\mathsf{w}}$$
$$\langle c, p_1 \rangle <:^{\mathsf{w}} \bot \land \langle c, p_1 \rangle <:^{\mathsf{w}} \langle c, p_3 \rangle \qquad \text{by trans. of } \preceq$$

For the induction case, suppose $w_2 = \langle T_2^{\mathsf{w}}, p_2 \rangle$ for some T_2^{w} and p_2 . If $w_3 = \bot$, the argument is trivial. We are thus left with the case where $w_1 = \langle T_1^{\mathsf{w}}, p_1 \rangle$ and $w_3 = \langle T_3^{\mathsf{w}}, p_3 \rangle$ for some $T_1^{\mathsf{w}}, T_3^{\mathsf{w}}, p_1$, and p_3 where $p_1 \preceq p_2$ and $p_2 \preceq p_3$. Hence, $p_1 \preceq p_3$. Let us first argue that $\forall p \in T_1^{\mathsf{w}}. p_3 \preceq p \Longrightarrow p \in T_3^{\mathsf{w}}$. Suppose not, i.e., there exists a path $p \in T_1^{\mathsf{w}}$ such that $p_3 \preceq p$ and $p \notin T_3^{\mathsf{w}}$. By $p_2 \preceq p_3$, we have that $p_2 \preceq p$ and hence $p \in T_2^{\mathsf{w}}$ by $w_1 <: {}^{\mathsf{w}}w_2$. Therefore, it also must be that $p \in T_3^{\mathsf{w}}$ by $w_2 <: {}^{\mathsf{w}}w_3$.

Consider next any call site path p. Let $T_1^{\mathsf{w}}(p) = \langle w_{1i}, w_{1o} \rangle$, $T_2^{\mathsf{w}}(p) = \langle w_{2i}, w_{2o} \rangle$, and $T_3^{\mathsf{w}}(p) = \langle w_{3i}, w_{3o} \rangle$. If $p \notin T_3^{\mathsf{w}}$, the argument is trivial. Hence, we can only focus on the case where $p \in T_3^{\mathsf{w}}$. We branch on the occurrence of p in T_2^{w} . Suppose $p \notin T_2^{\mathsf{w}}$. Hence, $w_{2i} = \bot$ and by $\mathsf{tiocons}(w_2)$ it must be that $w_{2o} = \bot$. Next, it cannot be the case where both $p \in T_1^{\mathsf{w}}$ and $p \in T_3^{\mathsf{w}}$ when $p \notin T_2^{\mathsf{w}}$. To see this, suppose this is not the case, i.e., $p \in T_1^{\mathsf{w}}$. By $\mathsf{pcscons}(w_3)$, we have that $p_3 \preceq p$. Since $p_2 \preceq p_3$, then also $p_2 \preceq p$. Then, $p_2 \in T_2^{\mathsf{w}}$ by $w_1 <:^{\mathsf{w}} w_2$, a contradiction. Hence, $p \notin T_1^{\mathsf{w}}$. By $\mathsf{tiocons}(w_1)$ we have $w_{1i} = w_{1o} = \bot$. The result then follows by $T_1^{\mathsf{w}}(p) = T_2^{\mathsf{w}}(p)$.

Suppose now $p \in T_2^{\mathsf{w}}$. Then, we have $w_{3i} <: {}^{\mathsf{w}}w_{2i}$ and $w_{2i} <: {}^{\mathsf{w}}w_{1i}$ as well as $w_{1o} <: {}^{\mathsf{w}}w_{2o}$ and $w_{2o} <: {}^{\mathsf{w}}w_{3o}$ and the result follows by i.h.

We next prove several results stating how the propagation between values affect the subtyping relation between the values. We start with the expected result that the propagation preserves subtyping between path values.

Lemma 68. Let $w_1 \in \mathcal{V}^w$, $w_2 \in \mathcal{V}^w$, $w_3 \in \mathcal{V}^w$, and $w_4 \in \mathcal{V}^w$. If $w_1 <:^{\mathsf{w}} w_2$ and $\langle w_3, w_4 \rangle = \mathsf{prop}^{\mathsf{w}}(w_1, w_2)$, then $w_3 <:^{\mathsf{w}} w_4$.

Proof. By structural induction on w_1 and w_2 . The proof argument is similar to the one of Lemma 38.

The propagation between two values reserves the subtyping relation not only between the two values, but also on the whole subtyping chain containing these two values.

Lemma 69. Let $w_1 \in \mathcal{V}^w$, $w_2 \in \mathcal{V}^w$, $w_3 \in \mathcal{V}^w$, $w'_1 \in \mathcal{V}^w$, $w'_2 \in \mathcal{V}^w$, and $w'_3 \in \mathcal{V}^w$. Suppose $w_1 <:^w w_2$, $w_2 <:^w w_3$, tiocons (w_1) , tiocons (w_2) , tiocons (w_3) , pcscons (w_1) , pcscons (w_2) , and pcscons (w_3) . If $\langle w'_1, w'_2 \rangle = \text{prop}^w(w_1, w_2)$ then $w'_2 <:^w w_3$. Similarly, if $\langle w'_2, w'_3 \rangle = \text{prop}^w(w_2, w_3)$ then $w_1 <:^w w'_2$. *Proof.* The proof goes by structural induction on w_2 and follows the structure of the proof of Lemma 67. Also, the proof relies on Lemma 61 and Lemma 57 to trigger the induction hypothesis.

It can also be shown that updated versions, by means of propagation, of value lying on the subtyping chain can be joined while retaining the subtyping property of the chain.

Lemma 70. Let $w_1 \in \mathcal{V}^w$, $w_2 \in \mathcal{V}^w$, $w_3 \in \mathcal{V}^w$, $w'_1 \in \mathcal{V}^w$, $w'_2 \in \mathcal{V}^w$, $w''_2 \in \mathcal{V}^w$, and $w''_3 \in \mathcal{V}^w$. Suppose $w_1 <:^w w_2$, $w_2 <:^w w_3$, tiocons (w_1) , tiocons (w_2) , tiocons (w_3) , pcscons (w_1) , pcscons (w_2) , and pcscons (w_3) . Next, suppose $\langle w'_1, w'_2 \rangle = \operatorname{prop}^w(w_1, w_2)$ and $\langle w''_2, w''_3 \rangle = \operatorname{prop}^w(w_2, w_3)$ such that $w'_1 <:^w w'_2 <:^w w_3$ and $w_1 <:^w w''_2 <:^w w''_3$. Then, $w'_1 <:^w w_2 \sqcup^w w''_2 <:^w w'_3$.

Proof. By structural induction on w_1 , w_2 , and w_3 .

We now finally define the subtyping consistency of path maps. A map M^{w} is subtyping consistent $\mathsf{csub}_e(M^{\mathsf{w}})$, subject to a program e, iff for every node n, where $loc(n) = \ell$, it must be that

$$\begin{split} M^{\mathsf{w}}(env(n)\diamond\ell_2) <:^{\mathsf{w}} \pi_1(M^{\mathsf{w}}(env(n)\diamond\ell_1)(p)) \wedge \pi_2(M^{\mathsf{w}}(env(n)\diamond\ell_1)(p)) <:^{\mathsf{w}} M^{\mathsf{w}}(n) \\ & \quad \mathbf{if} \; \exists e_1, e_2, \ell_1, \ell_2, T^{\mathsf{w}}, p. \, e(\ell) = (e_1^{\ell_1} e_2^{\ell_2}) \; \mathbf{and} \; M^{\mathsf{w}}(env(n)\diamond\ell_1) = \langle T^{\mathsf{w}}, p \rangle \\ & \quad \pi_1(M^{\mathsf{w}}(n)(p)) <:^{\mathsf{w}} M^{\mathsf{w}}(n_x) \wedge M^{\mathsf{w}}(E_1\diamond\ell_1) <:^{\mathsf{w}} \pi_2(M^{\mathsf{w}}(n)(p)) \\ & \quad \mathbf{if} \; \exists x, f, e_1, \ell_1, T^{\mathsf{w}}, p, n_x, n_f. \, e(\ell) = (\mu f.\lambda x. e_1^{\ell_1}), \; M^{\mathsf{w}}(n) = \langle T^{\mathsf{w}}, _- \rangle, p \in M^{\mathsf{w}}(n), \\ & \quad n_x = env(n)\diamond\mathsf{top}(p)\diamond x, \; n_f = env(n)\diamond\mathsf{top}(p)\diamond f, \; \mathbf{and} \; E_1 = env(n).x: n_x.f: n_f \\ & \quad \mathrm{and} \; \forall p \in \mathsf{paths}(M^{\mathsf{w}}), \langle n_1, n_2 \rangle \in p. \; M^{\mathsf{w}}(n_1) <:^{\mathsf{w}} M^{\mathsf{w}}(n_2). \end{split}$$

A.3.4.11 Path Consistency

We next introduce a path consistency, a property of maps similar to subtyping but providing more information about paths of values related to subtyping. A path map M^{w} is called path consistent $\mathsf{pvcons}(M^{\mathsf{w}})$, subject to a program e, iff $\forall p \in \mathsf{paths}(M^{\mathsf{w}})$. $M^{\mathsf{w}}(\mathsf{top}(p)) = p \land \forall \langle n_1, n_2 \rangle \in p$. $\mathsf{agree}_{n_2}(M^{\mathsf{w}}(n_1), M^{\mathsf{w}}(n_2))$ where $\mathsf{agree}_n(w_1, w_2) \iff \exists p_1, p_2. p_2 = p_1 \cdot n \land w_1 = \langle ., p_1 \rangle \land w_2 = \langle ., p_2 \rangle.$

That is, the values associated with neighbouring nodes cannot be bottom \perp or top ω and associated paths have to be an extension of each other.

We next prove several interesting properties of path consistent maps. We use prefix(p, i) to denote a prefix of p up to and including the node at position i in p. We also use $n \in i p$ to denote that n appears in p at the position i.

Lemma 71. Let M^{w} be a path map, $p \in \mathsf{paths}(M^{\mathsf{w}})$, and $n \in p$. If $\mathsf{pvcons}(M^{\mathsf{w}})$, then $M^{\mathsf{w}}(n) = \langle -, \mathsf{prefix}(p, i) \rangle$.

Proof. Follows from the definition of pvcons and agree. \Box

If two values agree, the data propagation between these two values is preserved.

Lemma 72. Let e be a program, n a node, and $\langle w_3, w_4 \rangle = \mathsf{prop}^{\mathsf{w}}(w_1, w_2)$ for path values w_1, w_2, w_3 , and w_4 . If $\mathsf{agree}_n(w_1, w_2)$, then $\mathsf{agree}_n(w_3, w_4)$.

Proof. The argument follows from Lemma 48.

A.3.4.12 Path Determinism

We next introduce another useful property of paths taken by data flow values in our semantics. That is, every taken path is deterministic in the sense that for every node in a path there is always a unique predecessor node. For this purpose, we define the predecessor relation $\operatorname{pred}_{M^{\mathsf{w}}}^{e}(n_1, n_2)$ to hold between nodes n_1 and n_2 , subject to a path map M^{w} and program e, iff one of the following six cases hold

1.

$$n_1 \in \mathcal{N}_x \land n_2 \in \mathcal{N}_e \land \exists x. \ e(loc(n_1)) = x \land x \in \mathsf{dom}(env(n_1)) \land n_2 = env(n_1)(x)$$
2.

$$n_{1} \in \mathcal{N}_{e} \wedge n_{2} \in \mathcal{N}_{x} \wedge \exists n, E_{0}, e_{1}, e_{2}, \ell_{1}, \ell_{2}, x, T^{\mathsf{w}}, p. n_{2} = E_{0} \diamond n \diamond x \wedge n_{1} = env(n) \diamond \ell_{2}$$
$$\wedge e_{1}^{loc(n)} e_{2}^{\ell_{2}} \in e \wedge M^{\mathsf{w}}(env(n) \diamond \ell_{2}) \not\in \{\bot, \omega\} \wedge M^{\mathsf{w}}(n) = \langle T^{\mathsf{w}}, p \rangle \wedge env(\operatorname{root}(p)) = E_{0} \wedge e(loc(\operatorname{root}(p))) = \mu_{-}.\lambda x.$$

3.

$$n_{1} \in \mathcal{N}_{e} \land n_{2} \in \mathcal{N}_{x} \land \exists n, E_{0}, e_{1}, e_{2}, \ell_{1}, \ell_{2}, f, T^{\mathsf{w}}, p. n_{2} = E_{0} \diamond n \diamond f \land n_{1} = \mathsf{root}(p)$$
$$\land e_{1}^{loc(n)} e_{2}^{\ell_{2}} \in e \land M^{\mathsf{w}}(env(n) \diamond \ell_{2}) \notin \{\bot, \omega\} \land M^{\mathsf{w}}(n) = \langle T^{\mathsf{w}}, p \rangle \land$$
$$env(\mathsf{root}(p)) = E_{0} \land e(loc(\mathsf{root}(p))) = \mu f.\lambda_{--}$$

4.

$$n_{1} \in \mathcal{N}_{e} \land n_{2} \in \mathcal{N}_{e} \land \exists n, E_{0}, E_{b}, e_{1}, e_{2}, e_{b}, \ell_{1}, \ell_{2}, \ell_{b}, f, x, T^{w}, p.$$

$$env(n_{2}) = env(n) \land n_{1} = E_{b} \diamond \ell_{b} \land e(loc(n_{2})) = e_{1}^{loc(n)} e_{2}^{\ell_{2}} \land M^{w}(n) = \langle T^{w}, p \rangle \land$$

$$M^{w}(env(n) \diamond \ell_{2}) \notin \{\bot, \omega\} \land env(root(p)) = E_{0} \land e(loc(root(p))) = \mu f. \lambda x. e_{b}^{\ell_{b}} \land$$

$$E_{b} = E_{0}.x : E_{0} \diamond n \diamond x. f : E_{0} \diamond n \diamond f$$

5.

$$n_1 \in \mathcal{N}_e \wedge n_2 \in \mathcal{N}_e \wedge \exists e_0, e_1, e_2, \ell_0.$$
$$e(loc(n_2)) = e_0^{\ell_0} ? e_1^{loc(n_1)} : e_2 \wedge M^{\mathsf{w}}(env(n_2) \diamond \ell_0) = \langle true, _\rangle \wedge env(n_1) = env(n_2)$$

$$n_1 \in \mathcal{N}_e \wedge n_2 \in \mathcal{N}_e \wedge \exists e_0, e_1, e_2, \ell_0.$$
$$e(loc(n_2)) = e_0^{\ell_0} ? e_1 : e_2^{loc(n_1)} \wedge M^{\mathsf{w}}(env(n_2) \diamond \ell_0) = \langle false, \lrcorner \rangle \wedge env(n_1) = env(n_2)$$

Note that we again use the _ pattern to quantify out certain mathematical objects to avoid clutter.

The first interesting property of the predecessor relation is that there is no predecessor of a node that corresponds to the constant expression or lambda expression.

Lemma 73. Let n_2 be a node, M^w a path map, and e a program. If $e(loc(n_2))$ is a constant or a lambda expression, then there does not exist a node n_1 such that $\operatorname{pred}_{M^w}^e(n_1, n_2)$.

Proof. Follows from the definition of the pred relation. \Box

A node cannot be its predecessor, i.e., the **pred** relation is not reflexive.

Lemma 74. Let n be a node, M^{w} a path map, and e a program. Then, it cannot be the case that $\operatorname{pred}_{M^{\mathsf{w}}}^{e}(n, n)$.

Proof. Follows from the definition of the pred relation. \Box

Next, the predecessor relation is deterministic.

Lemma 75. Let n, n_1 , and n_2 be nodes, M^w a path map, and e a program. If $\operatorname{pred}_{M^w}^e(n, n_1)$ and $\operatorname{pred}_{M^w}^e(n, n_2)$, then $n_1 = n_2$.

Proof. Follows from the definition of the pred relation, the fact that environments and path maps are modeled as functions, and uniqueness of locations in e.

Although the predecessor relation is deterministic, two different nodes can share their predecessor.

Lemma 76. Let n, n_1 , and n_2 be nodes where $n_1 \neq n_2$, M^w a path map, and e a program. If $\operatorname{pred}^e_{M^w}(n, n_1)$ and $\operatorname{pred}^e_{M^w}(n, n_2)$, then $n \in \mathcal{N}_x$ or $e(\operatorname{loc}(n))$ is a lambda expression.

Proof. Follows from the definition of the pred relation. \Box

We say that path p is deterministic $\det^{e}_{M^{\mathsf{w}}}(p)$, subject to a program e and path map M^{w} , iff $\forall \langle n_{1}, n_{2} \rangle \in p$. $\mathsf{pred}^{e}_{M^{\mathsf{w}}}(n_{1}, n_{2})$. A map M^{w} is deterministic $\det^{e}(M^{\mathsf{w}})$ subject to a program e iff $\forall p \in \mathsf{paths}(M^{\mathsf{w}})$. $\det^{e}_{M^{\mathsf{w}}}(p)$.

A.3.4.13 Error Consistency

Similar to data flow semantics, we introduce the notion of the error absence of values that intuitively states the a given value is not an error and, if it is a table, it does not recursively contain errors.

$$\operatorname{noerr}(w) \iff (w \in \{\bot\} \cup Cons^{\mathsf{w}}) \lor (\exists T^{\mathsf{w}}. w = \langle T^{\mathsf{w}}, \lrcorner\rangle \land \forall p \in T^{\mathsf{w}}.$$
$$\operatorname{noerr}(\pi_1(T^{\mathsf{w}}(p))) \land \operatorname{noerr}(\pi_2(T^{\mathsf{w}}(p))))$$

A path map M^{w} is error consistent $\operatorname{errcons}(M^{\mathsf{w}})$ if either (1) $M^{\mathsf{w}} = M^{\mathsf{w}}_{\omega}$ or (2) for every node *n* it must be that $\operatorname{noerr}(M^{\mathsf{w}}(n))$.

Data propagation does not invalidate the absence of error property of path values, given that these values are related by the concrete subtyping relation on path values.

Lemma 77. Let $w_1 \in \mathcal{V}^w$, $w_2 \in \mathcal{V}^w$, $w_3 \in \mathcal{V}^w$, and $w_4 \in \mathcal{V}^w$. If $w_1 <: w_2$, noerr (w_1) , noerr (w_2) , and $\langle w_3, w_4 \rangle = \text{prop}^w(w_1, w_2)$, then noerr (w_3) and noerr (w_4) . *Proof.* By structural induction on w_1 and w_2 . The proof argument is almost identical to the one of Lemma 40.

A.3.4.14 Environment Consistency

The notion of environment consistency for a path map is identical to the one for regular data flow maps.

 $\operatorname{envcons}(M^{\mathsf{w}}) \stackrel{\text{def}}{\Longrightarrow} \forall n, n_x \in \operatorname{rng}(env(n)). \ M^{\mathsf{w}}(n) \neq \bot \implies M^{\mathsf{w}}(n_x) \notin \{\bot, \omega\}.$

Consistency of an environment E subject to a path map M^w is defined as follows

envcons^w_M(E)
$$\iff \forall n_x \in \operatorname{rng}(E). M^{\mathsf{w}}(n_x) \notin \{\bot, \omega\}$$

A.3.4.15 Consistency

We can now define a general consistency $cons_e(M^w)$ of a path map M^w subject to a program e as follows

$$\begin{aligned} & \operatorname{cons}_e(M^{\mathsf{w}}) & \xleftarrow{}^{\operatorname{def}} \operatorname{tpcons}(M^{\mathsf{w}}) \wedge \operatorname{pcscons}(M^{\mathsf{w}}) \wedge \operatorname{rootcons}_e(M^{\mathsf{w}}) \wedge \operatorname{inpcons}_e(M^{\mathsf{w}}) \wedge \\ & \operatorname{outcons}_e(M^{\mathsf{w}}) \wedge \operatorname{det}_e(M^{\mathsf{w}}) \wedge \operatorname{csub}_e(M^{\mathsf{w}}) \wedge \operatorname{tiocons}_e(M^{\mathsf{w}}) \wedge \operatorname{fin}(M^{\mathsf{w}}) \\ & \wedge \operatorname{constcons}_e(M^{\mathsf{w}}) \wedge \operatorname{errcons}(M^{\mathsf{w}}) \wedge \operatorname{envcons}(M^{\mathsf{w}}) \wedge \operatorname{pvcons}(M^{\mathsf{w}}) \end{aligned}$$

We now prove several interesting properties of general consistency of maps.

Lemma 78. Let M^{w} be a consistent map $\mathsf{cons}_e(M^{\mathsf{w}})$ subject to a program e. Then, $\forall p \in \mathsf{paths}(M^{\mathsf{w}}), n \in p. M^{\mathsf{w}}(n) \notin \{\bot, \omega\}.$

Proof. Follows from the definition of pvcons.

The next result states that the observed paths in a consistent map can be uniquely identified by their top nodes.

Lemma 79. Let M^{w} be a consistent map $\mathsf{cons}_e(M^{\mathsf{w}})$ subject to a program e. Then, $\forall p_1, p_2 \in \mathsf{paths}(M^{\mathsf{w}}). p_1 \neq p_2 \Longrightarrow \mathsf{top}(p_1) \neq \mathsf{top}(p_2).$ *Proof.* By the induction on both $|p_1|$ and $|p_2|$. By $\mathsf{pvcons}(M^{\mathsf{w}})$, prefixes of both p_1 and p_2 are in $\mathsf{paths}(M^{\mathsf{w}})$. The argument then follows from $\mathsf{det}_e(M^{\mathsf{w}})$, $\mathsf{rootcons}_e(M^{\mathsf{w}})$, Lemma 75, and Lemma 73.

Further, a node cannot appear at two different location in any path of a consistent map. That is, individual paths do not constitute cycles.

Lemma 80. Let M^{w} be a consistent map $\mathsf{cons}_e(M^{\mathsf{w}})$ subject to a program e and $p \in \mathsf{paths}(M^{\mathsf{w}})$. For any given node n, suppose $n \in^i p$ and $n \in^j \in p$. Then, it must be that i = j.

Proof. Suppose not, $i \neq j$. Let $p_i = \operatorname{prefix}(p, i)$ and $p_j = \operatorname{prefix}(p, j)$. Clearly, $p_i \neq p_j$. However, by Lemma 71 it must be that $M^{\mathsf{w}}(n) = \langle -, p_i \rangle = \langle -, p_j \rangle$. Hence, $p_i = p_j$. Contradiction.

Moreover, a node is always visited from the same predecessor node.

Lemma 81. Let M^{w} be a consistent map $\mathsf{cons}_e(M^{\mathsf{w}})$ subject to a program e. Also, let n_1 and n_2 be nodes such that $\mathsf{pred}^e_{M^{\mathsf{w}}}(n_1, n_2)$. Then, $M^{\mathsf{w}}(n_2) = \bot \lor M^{\mathsf{w}}(n_2) = \omega \lor \mathsf{agree}_{n_2}(M^{\mathsf{w}}(n_1), M^{\mathsf{w}}(n_2))$.

Proof. Let $M^{\mathsf{w}}(n_2) = \langle , p \rangle$ for some p. First, consider the case $|p| \models 1$.

 $p \in \mathsf{paths}(M^{\mathsf{w}})$ by def. of paths

 $e(loc(n_2))$ is a constant or lambda by def. of rootcons

$$\neg \mathsf{pred}^e_{M^{\mathsf{w}}}(n_1, n_2)$$
 by Lemma 73

Next, consider the case |p| > 1.

$p = p' \cdot n_2$	by $tpcons_e(M^{w})$ for some p'
$M^{w}(n_1) = \langle _, p' \rangle$	by Lemma 71 and Lemma 80

We next show that propagation among values from a consistent map cannot invalidate subtyping on branching points in the subtyping chains.

Lemma 82. Let M^{w} be a consistent map $\mathsf{cons}_e(M^{\mathsf{w}})$ subject to a program e. Also, let $n, n_1, and n_2$ be nodes such that $n_1 \neq n_2, \operatorname{pred}_{M^{\mathsf{w}}}^e(n, n_2), and \operatorname{pred}_{M^{\mathsf{w}}}^e(n, n_2).$ Suppose $M^{\mathsf{w}}(n) <:^{\mathsf{w}} M^{\mathsf{w}}(n_1), M^{\mathsf{w}}(n) <:^{\mathsf{w}} M^{\mathsf{w}}(n_2), and \langle w', w_1' \rangle = \operatorname{prop}_{\mathsf{w}}(M^{\mathsf{w}}(n), M^{\mathsf{w}}(n_1)).$ Then, $w' <:^{\mathsf{w}} M^{\mathsf{w}}(n_2).$

Proof. Let $M^{\mathsf{w}}(n) = w$, $M^{\mathsf{w}}(n_1) = w_1$, and $M^{\mathsf{w}}(n_2) = w_2$. First, observe that if $w_1 = \bot$ then w' = w and the result follows immediately. Otherwise, the proof goes by a case analysis on $M^{\mathsf{w}}(n_2)$.

- $w_2 = \bot$. Trivial, by the def. of $<:^{\mathsf{w}}$.
- $w_2 = \langle c, p_2 \rangle$. Then $w = \langle c, p \rangle$ so that $p \preceq p_2$ and $w_1 = \bot$ or $w_1 = \langle c, p_1 \rangle$ where $p \preceq p_1$. Then, by the def. of prop^w it follows that $w' = \langle c, p \rangle$.

•
$$w_2 = \omega$$
. Then, $w = w' = w'_1 = w_2 = \omega$.

Next, suppose $w = \langle T_2^{\mathsf{w}}, p_2 \rangle$ for some T_2^{w} and p_2 . Then it must be that $w_1 = \langle T_1^{\mathsf{w}}, p_1 \rangle$ for some T_1^{w} and p_1 . Also, it is the case that $w = \langle T^{\mathsf{w}}, p \rangle$ for some T^{w} and p. First, observe that

$$\operatorname{top}(p) = n, \operatorname{top}(p_1) = n_1, \operatorname{top}(p_2) = n_2$$
 by $\operatorname{tpcons}(M^w)$
 $p_1 \neq p_2$

We next show that $p_1 \not\leq p_2$; the case $p_2 \not\leq p_1$ is handled similarly. Suppose $p_1 \leq p_2$. By $\mathsf{pvcons}(M^{\mathsf{w}})$ and $\mathsf{det}_e(M^{\mathsf{w}})$, we know that $p_2 = p \cdot n_2$ and $p_1 = p \cdot n_1$. Hence, there exist *i* and *j* such that $i \neq j$, $n \in^i p_2$, $n \in^j \in p_2$. However, this is a contradiction by Lemma 80. Thus, $p_1 \neq p_2$, $p_1 \not\leq p_2$, and $p_2 \not\leq p_1$. By $\mathsf{pcscons}(M^{\mathsf{w}})$, the set of observed call site paths in T_1^w and T_2^w is different. The result then follows immediately from the definition of prop^w.

We next show that output values in tables must agree, via concrete subtyping, with the value of the lambda body used to initially populate the table.

Lemma 83. Let M^{w} be a consistent map $\mathsf{cons}_e(M^{\mathsf{w}})$ subject to a program e, wa path value, n a node, and p_c as well as p' paths. Suppose $\langle T^{\mathsf{w}}, p \rangle = M^{\mathsf{w}}(n)$, for some T^{w} and p, and let $\pi_2(T^{\mathsf{w}}(p_c)) = w$. If $w = \langle -, p' \rangle$, then $M^{\mathsf{w}}(\mathsf{top}(p')) <:^{\mathsf{w}} w$ and $M^{\mathsf{w}}(\mathsf{top}(p')) = \langle -, p' \rangle$.

Proof.

$$\begin{array}{ll} p,p_c,p'\in \mathsf{paths}(M^{\mathsf{w}}) & \text{by def. of paths} \\ e(loc(\mathsf{root}(p))) = \mu f.\lambda x.e_1^{\ell_1} & \text{for some } f,x,e_1,\ell_1 \text{ by }\mathsf{rootcons}_e(M^{\mathsf{w}}) \\ n_x = env(\mathsf{root}p)\diamond\mathsf{top}(p_c)\diamond x, & \text{assume} \\ n_f = env(\mathsf{root}p)\diamond\mathsf{top}(p_c)\diamond f, \\ E_1 = env(\mathsf{root}(p)).x:n_x.f:n_f & \\ \mathsf{top}(p') = E_1\diamond\ell_1 & \text{by }\mathsf{outcons}_e(M^{\mathsf{w}}) \\ M^{\mathsf{w}}(\mathsf{root}(p)) < :^{\mathsf{w}} M^{\mathsf{w}}(n) & \text{by }\mathsf{csub}_e(M^{\mathsf{w}}), \mathsf{tiocons} M^{\mathsf{w}}, \mathsf{pcscons}_e(M^{\mathsf{w}}), \\ & \\ M^{\mathsf{w}}(\mathsf{root}(p)) = \langle T_r^{\mathsf{w}}, _{-} \rangle & \text{for some } T_r^{\mathsf{w}} \text{ by }\mathsf{def. of } <:^{\mathsf{w}} \\ M^{\mathsf{w}}(E_1\diamond\ell_1) < :^{\mathsf{w}} \pi_2(T_r^{\mathsf{w}}(p_c)) & \text{by }\mathsf{csub}_e(M^{\mathsf{w}}), \mathsf{tiocons} M^{\mathsf{w}}, \mathsf{pcscons}_e(M^{\mathsf{w}}), \\ & \\ M^{\mathsf{w}}(E_1\diamond\ell_1) < :^{\mathsf{w}} w & \text{by }\mathsf{csub}_e(M^{\mathsf{w}}), \mathsf{tiocons} M^{\mathsf{w}}, \mathsf{pcscons}_e(M^{\mathsf{w}}), \\ & \\ M^{\mathsf{w}}(E_1\diamond\ell_1) < :^{\mathsf{w}} w & \text{by }\mathsf{csub}_e(M^{\mathsf{w}}), \mathsf{tiocons} M^{\mathsf{w}}, \mathsf{pcscons}_e(M^{\mathsf{w}}), \\ & \\ M^{\mathsf{w}}(E_1\diamond\ell_1) = \langle -, p' \rangle & \text{by }\mathsf{tpcons}(M^{\mathsf{w}}) \text{ and }\mathsf{Lemma } 79 \end{array}$$

A similar result holds for inputs of tables that are pushed to tables at the execution points corresponding to function application expressions.

Lemma 84. Let M^{w} be a consistent map $\mathsf{cons}_e(M^{\mathsf{w}})$ subject to a program e, wa path value, n a node, and p_c as well as p' paths. Suppose $\langle T^{\mathsf{w}}, p \rangle = M^{\mathsf{w}}(n)$, for some T^{w} and p, and let $\pi_1(T^{\mathsf{w}}(p_c)) = w$. If $w = \langle -, p' \rangle$, then $M^{\mathsf{w}}(\mathsf{top}(p')) <:^{\mathsf{w}} w$ and $M^{\mathsf{w}}(\mathsf{top}(p')) = \langle -, p' \rangle$.

Proof.

 $p, p_c, p' \in \mathsf{paths}(M^{\mathsf{w}})$ by def. of paths by $pcscons(M^w)$ $p \preceq p_c$ $e_1^{loc(\operatorname{top}(p_c))} e_2^{loc(\operatorname{top}(p'))} \in e,$ for some e_1, e_2 by $\mathsf{inpcons}_e(M^{\mathsf{w}})$ $env(top(p_c)) = env(top(p'))$ $M^{\mathsf{w}}(n) <:^{\mathsf{w}} M^{\mathsf{w}}(\mathsf{top}(p_c))$ by $\mathsf{csub}_e(M^{\mathsf{w}})$, $\mathsf{tiocons}M^{\mathsf{w}}$, $\mathsf{pcscons}_e(M^{\mathsf{w}})$, and Lemma 67 $M^{\mathsf{w}}(\mathsf{top}(p_c)) = \langle T_c^{\mathsf{w}}, _{-} \rangle$ by def. of $<:^{\mathsf{w}}$ and $\mathsf{pvcons}(M^{\mathsf{w}})$ $\pi_1(T_c^{\mathsf{w}}(p_c)) <:^{\mathsf{w}} w$ by def. of $<:^{\mathsf{w}}$ $M^{\mathsf{w}}(\mathsf{top}(p')) <:^{\mathsf{w}} \pi_1(T_c^{\mathsf{w}}(p_c))$ by $\operatorname{csub}_e(M^{\mathsf{w}})$ $M^{\mathsf{w}}(\mathsf{top}(p')) <:^{\mathsf{w}} w$ by $\mathsf{csub}_e(M^{\mathsf{w}})$, $\mathsf{tiocons}M^{\mathsf{w}}$, $\mathsf{pcscons}_e(M^{\mathsf{w}})$, and Lemma 67 by $tpcons(M^w)$ and Lemma 79 $M^{\mathsf{w}}(\mathsf{top}(p')) = \langle _, p' \rangle$

We next show that the concrete transformer does not change the structure of the paths, i.e., the **pred** relation.

Lemma 85. Let M^{w} , M_1^{w} , and M_2^{w} be path maps such that $\operatorname{cons}_e(M_1^{\mathsf{w}})$, $\operatorname{cons}_e(M_2^{\mathsf{w}})$, and $\operatorname{cons}_e(M^{\mathsf{w}})$ for some program e. Given an expression $e_0^{\ell} \in e$, fuel k, and two environments E_1 and E_2 such that $E_1 \neq E_2$, $\operatorname{domvalid}(\ell, e, E_1, , \operatorname{domvalid}(\ell, e, E_2)$, and $\operatorname{dom}(E_1) = \operatorname{dom}(E_2)$, assume that $\operatorname{step}_k^{\mathsf{w}}[\![e_0^{\ell}]\!](E_1)(M^{\mathsf{w}}) = M_1^{\mathsf{w}}$ and $\operatorname{step}_k^{\mathsf{w}}[\![e_0^{\ell}]\!](E_2)(M^{\mathsf{w}}) =$ M_2^{w} . Then, $\forall p_1, p_2 \in \operatorname{paths}(M_1^{\mathsf{w}}) \cup \operatorname{paths}(M_2^{\mathsf{w}})$, $i, j, n. n \in i p_1 \land n \in j p_2 \Longrightarrow \operatorname{prefix}(p_1, i) =$ $\operatorname{prefix}(p_2, j)$.

Proof. Let $p' = \operatorname{prefix}(p_1, i)$ and $p'' = \operatorname{prefix}(p_2, j)$. Clearly, $\operatorname{top}(p') = n = \operatorname{top}(p'')$ by the definition of prefix. First, if $p_1 \in \operatorname{paths}(M_1^w)$ and $p_2 \in \operatorname{paths}(M_1^w)$, then by Lemma 71 we have that $p', p'' \in \operatorname{paths}(M_1^w)$ and the result thus follows from Lemma 79. The same argument applies for the case when $p_1 \in \operatorname{paths}(M_2^w)$ and $p_2 \in \operatorname{paths}(M_2^w)$.

We are therefore left with the case when $p_1 \in \mathsf{paths}(M_1^w)$ and $p_2 \in \mathsf{paths}(M_2^w)$ but $p_1 \notin \mathsf{paths}(M_2^w)$ and $p_2 \notin \mathsf{paths}(M_1^w)$. (The same argument that follows also applies for the case where we invert the memberships of p_1 and p_2 .) By Lemma 71, we have that $p' \in \mathsf{paths}(M_1^w)$ and $p'' \in \mathsf{paths}(M_2^w)$. We thus have $\mathsf{det}_{M_1^w}^e(p')$ and $\mathsf{det}_{M_2^w}^e(p'')$. The argument follows by a mutual induction on |p'| and |p''|. If |p'| = |p''| = 1, the argument is trivial. Next, observe that by $\mathsf{rootcons}_e(M_1^w)$ and $\mathsf{rootcons}_e(M_2^w)$ and Lemma 73 it cannot be the case that |p'| > 1 and |p''| = 1 (and similarly |p''| > 1 and |p'| = 1). Hence, both p' and p'' have a predecessor of a top node n; let those predecessor node be n' and n'', respectively. We next show that n' = n''. The argument then inductively follows by Lemma 71. We consider all the possible cases of n by inspecting the definition of pred. Observe that we have $\mathsf{pred}_{M_1^w}^e(n', n)$ and $\mathsf{pred}_{M_2^w}^e(n'', n)$.

- e(loc(n)) = x.
- $x \in \mathsf{dom}(env(n))$ by def. of pred $\forall M^{\mathsf{w}}, n_0. \operatorname{pred}^{e}_{M^{\mathsf{w}}}(n_0, n) \Longrightarrow n_0 = env(n)(x)$ by def. of pred $\mathsf{pred}^e_{M^{\mathsf{w}}_1}(env(n)(x),n),\mathsf{pred}^e_{M^{\mathsf{w}}_2}(env(n)(x),n)$ n' = n'' = env(n)(x)• $n = E \diamond n_c \diamond x$ for some E, n_c , and x. • $n_c \notin \operatorname{ranger}(E_1 \diamond \ell), n_c \notin \operatorname{ranger}(E_2 \diamond \ell)$ assume $M^{\mathsf{w}}(n_c) = M_1^{\mathsf{w}}(n_c) = M_2^{\mathsf{w}}(n_c)$ by Lemma 50 and Lemma 51 n' = n''by def. of pred • $n_c \in \operatorname{ranger}(E_1 \diamond \ell), n_c \in \operatorname{ranger}(E_2 \diamond \ell)$ assume by Lemma 50 and Lemma 51 impossible • $n_c \in \operatorname{ranger}(E_1 \diamond \ell), n_c \notin \operatorname{ranger}(E_2 \diamond \ell)$ assume
 - $$\begin{split} M_2^{\sf w}(n_c) &= M^{\sf w}(n_c) & \text{by Lemma 50 and Lemma 51} \\ M^{\sf w}(n_c) &\neq \bot & \text{by def. of pred} \\ M_1^{\sf w}(n_c) &\neq \omega & \text{by errcons}(M_1^{\sf w}) \\ M^{\sf w}(n_c) &= \langle _, p_c \rangle, & \text{for some } p_c \text{ by Lemma 56} \\ M_2^{\sf w}(n_c) &= \langle _, p_c \rangle \\ n' &= n'' & \text{by def. of pred} \end{split}$$
 - $n_c \not\in \operatorname{ranger}(E_1 \diamond \ell), n_c \in \operatorname{ranger}(E_2 \diamond \ell)$ assume
 - similar to above
- Similar for other cases of n.

We also show that the join of two consistent path maps, obtained by parallel applications of the concrete transformer, is subtype consistent.

Lemma 86. Let M^{w} , M_1^{w} , and M_2^{w} be path maps such that $\operatorname{cons}_e(M_1^{\mathsf{w}})$, $\operatorname{cons}_e(M_2^{\mathsf{w}})$, and $\operatorname{cons}_e(M^{\mathsf{w}})$ for some program e. Assume we are also given an expression $e_0^{\ell} \in e$, fuel k, and two environments E_1 and E_2 such that $E_1 \neq E_2$, $\operatorname{domvalid}(\ell, e, E_1)$, $\operatorname{domvalid}(\ell, e, E_2)$, and $\operatorname{dom}(E_1) = \operatorname{dom}(E_2)$. If $\operatorname{step}_k^{\mathsf{w}}[\![e_0^{\ell}]\!](E_1)(M^{\mathsf{w}}) = M_1^{\mathsf{w}}$ and $\operatorname{step}_k^{\mathsf{w}}[\![e_0^{\ell}]\!](E_2)(M^{\mathsf{w}}) = M_2^{\mathsf{w}}$, then $\operatorname{csub}_e(M_1^{\mathsf{w}} \sqcup^{\mathsf{w}} M_2^{\mathsf{w}})$.

Proof. Let $p \in \mathsf{paths}(M_1^{\mathsf{w}} \dot{\sqcup}^{\mathsf{w}} M_2^{\mathsf{w}})$ and $\langle n_1, n_2 \rangle$. We show that

 $M_1^{\mathsf{w}}(n_1) \sqcup^{\mathsf{w}} M_2^{\mathsf{w}}(n_1) <: {}^{\mathsf{w}} M_2^{\mathsf{w}}(n_1) \sqcup^{\mathsf{w}} M_1^{\mathsf{w}}(n_2)$

For the remainder of the definition of csub_e , the argument follows from Lemma 50 and Lemma 51. Let $w_1 = M_1^{\mathsf{w}}(n_1) \sqcup^{\mathsf{w}} M_2^{\mathsf{w}}(n_1)$ and $w_2 = M_1^{\mathsf{w}}(n_2) \sqcup^{\mathsf{w}} M_2^{\mathsf{w}}(n_2)$.

First, observe that $p \in \mathsf{paths}(M_1^\mathsf{w})$ or $p \in \mathsf{paths}(M_2^\mathsf{w})$ by the definition of $\dot{\sqcup}^\mathsf{w}$. Also, note that $M_1^\mathsf{w}(n_1) \neq \omega, M_2^\mathsf{w}(n_1) \neq \omega, M_2^\mathsf{w}(n_1) \neq \omega$, and $M_1^\mathsf{w}(n_2) \neq \omega$. Suppose not. By $\mathsf{errcons}(M_1^\mathsf{w})$ and $\mathsf{errcons}(M_2^\mathsf{w})$ we have that $M_1^\mathsf{w} = M_\omega^\mathsf{w}$ or $M_2^\mathsf{w} = M_\omega^\mathsf{w}$, in which case $M_1^\mathsf{w} \dot{\sqcup}^\mathsf{w} M_2^\mathsf{w} = M_\omega^\mathsf{w}$ and thus $\mathsf{paths}(M_1^\mathsf{w} \dot{\sqcup}^\mathsf{w} M_2^\mathsf{w}) = \emptyset$, a contradiction. Further, note that $w_1 \neq \bot$ and $w_2 \neq \bot$ by $\mathsf{cons}_e(M_1^\mathsf{w})$, $\mathsf{cons}_e(M_2^\mathsf{w})$, and Lemma 78.

We proceed with the proof by structural induction on w_1 and w_2 where we focus on the cases where $w_1 \neq \bot$ and $w_2 \neq \bot$. Suppose first $w_1 = \omega$. There are two cases possible. First, $M_1^w(n_1) = \langle c_1, p_1 \rangle$ and $M_2^w(n_1) = \langle c_2, p_2 \rangle$ where either $c_1 \neq c_2 \text{ or } p_1 \not\preceq p_2 \land p_2 \not\preceq p_1.$

contradiction

The other case is when $M_1^w(n_1)$ is a constant path value and $M_2^w(n_1)$ is table path value, or vice versa. The argument is then exactly the same. Moreover, the same reasoning steps are sufficient to show that $w_2 \neq \omega$. In conclusion, $w_1 \neq \omega$ and $w_2 \neq \omega$.

Next, suppose that $w_1 = \langle c_1, p_1 \rangle$ and $w_2 = \langle c_2, p_2 \rangle$ where either $c_1 \neq c_2$ or $p_1 \not\preceq p_2 \land p_2 \not\preceq p_1$.

- $(M_1^{\mathsf{w}}(n_1) = \langle c_1, p_1 \rangle \land M_1^{\mathsf{w}}(n_2) = \langle c_1, p_2 \rangle) \lor \qquad \text{by } \operatorname{cons}_e(M_1^{\mathsf{w}}), \operatorname{cons}_e(M_2^{\mathsf{w}}),$ $(M_2^{\mathsf{w}}(n_1) = \langle c_1, p_1 \rangle \land M_2^{\mathsf{w}}(n_2) = \langle c_1, p_2 \rangle) \qquad \text{and Lemma 78}$ $p_1 \preceq p_2 \qquad \text{by } \operatorname{csub}_e(M_1^{\mathsf{w}}) \text{ and } \operatorname{csub}_e(M_2^{\mathsf{w}})$
 - $c_1 = c_2$ by $\operatorname{rootcons}_e(M_1^{\mathsf{w}}), \operatorname{rootcons}_e(M_2^{\mathsf{w}}),$

 $\operatorname{constcons}(M_1^{\mathsf{w}}), \operatorname{constcons}(M_2^{\mathsf{w}}),$

 $\operatorname{csub}_e(M_1^{\mathsf{w}}), \operatorname{csub}_e(M_2^{\mathsf{w}}),$

Lemma 67, and def. of $<:^{\mathsf{w}}$

contradiction

The same reasoning steps apply for the other base cases of $w_1 \in Cons^{\mathsf{w}} \wedge w_2 \in \mathcal{T}^{\mathsf{w}}$ and $w_1 \in \mathcal{T}^{\mathsf{w}} \wedge w_2 \in Cons^{\mathsf{w}}$. When $w_1 \in \mathcal{T}^{\mathsf{w}}$ and $w_2 \in \mathcal{T}^{\mathsf{w}}$, for every call site path p the inputs of w_1 and w_2 , and outputs respectively, at p must agree on the top nodes by $\operatorname{inpcons}_e(M_1^w)$, $\operatorname{inpcons}_e(M_2^w)$, $\operatorname{outcons}_e(M_1^w)$, and $\operatorname{outcons}_e(M_2^w)$. The same reasoning as above then inductively applies by Lemma 84, Lemma 83, and the definition of $<:^w$.

Lemma 87. Let M^{w} , M_1^{w} , and M_2^{w} be path maps such that $\operatorname{cons}_e(M_1^{\mathsf{w}})$, $\operatorname{cons}_e(M_2^{\mathsf{w}})$, and $\operatorname{cons}_e(M^{\mathsf{w}})$ for some program e. Given an expression $e_0^{\ell} \in e$, fuel k, and two environments E_1 and E_2 such that $E_1 \neq E_2$, $\operatorname{domvalid}(\ell, e, E_1)$, $\operatorname{domvalid}(\ell, e, E_2)$, and $\operatorname{dom}(E_1) = \operatorname{dom}(E_2)$, assume that $\operatorname{step}_k^{\mathsf{w}}[\![e_0^{\ell}]\!](E_1)(M^{\mathsf{w}}) = M_1^{\mathsf{w}}$ and $\operatorname{step}_k^{\mathsf{w}}[\![e_0^{\ell}]\!](E_2)(M^{\mathsf{w}}) =$ M_2^{w} . Then, $\operatorname{pvcons}(M_1^{\mathsf{w}} \sqcup^{\mathsf{w}} M_2^{\mathsf{w}})$.

Proof. Similar to the proof of Lemma 86.

Lemma 88. Let M^{w} , M_1^{w} , and M_2^{w} be path maps such that $\operatorname{cons}_e(M_1^{\mathsf{w}})$, $\operatorname{cons}_e(M_2^{\mathsf{w}})$, and $\operatorname{cons}_e(M^{\mathsf{w}})$ for some program e. Given an expression $e_0^{\ell} \in e$, fuel k, and two environments E_1 and E_2 such that $E_1 \neq E_2$, $\operatorname{domvalid}(\ell, e, E_1)$, $\operatorname{domvalid}(\ell, e, E_2)$, and $\operatorname{dom}(E_1) = \operatorname{dom}(E_2)$, assume that $\operatorname{step}_k^{\mathsf{w}}[\![e_0^{\ell}]\!](E_1)(M^{\mathsf{w}}) = M_1^{\mathsf{w}}$ and $\operatorname{step}_k^{\mathsf{w}}[\![e_0^{\ell}]\!](E_2)(M^{\mathsf{w}}) =$ M_2^{w} . Then, $\operatorname{errcons}(M_1^{\mathsf{w}} \sqcup^{\mathsf{w}} M_2^{\mathsf{w}})$.

Proof. Similar to the proof of Lemma 86.

The above results are used to show that the general consistency of path maps is preserved by **step^w**.

Theorem 7. Let $M^{\mathsf{w}}, M_f^{\mathsf{w}} \in \mathcal{M}^{\mathsf{w}}, M_f^{\mathsf{w}} \in \mathcal{M}^{\mathsf{w}}, e^{\ell}$ an expression of a program m, $k \in \mathbb{N}$, and E an environment where (e, E) is well-formed, domvalid (ℓ, m, E) , and $\mathsf{envcons}_M^{\mathsf{w}}(E)$. If $\mathsf{cons}_m(M^{\mathsf{w}})$ and $\mathsf{step}_k^{\mathsf{w}}[\![e^{\ell}]\!](E)(M^{\mathsf{w}}) = M_f^{\mathsf{w}}$, then $\mathsf{cons}_m(M_f^{\mathsf{w}})$ and $\mathsf{envcons}_{M_f^{\mathsf{w}}}(E)$. *Proof.* The proof goes by induction on k and structural induction on e. We need to show that $envcons(M^w)$, $csub_m(M^w_f)$, $pvcons_m(M^w_f)$, $errcons(M^w_f)$, $constcons_m(M^w_f)$, $det_m(M^w_f)$. When k = 0, the argument is trivial. We consider the cases for e when k > 1.

- e = c. Let M^w(E◊ℓ) = w. By constcons_m(M^w_f), we have that w ∈ {⊥, ω, ⟨c, ℓ⟩}. If w = ⟨c, ℓ⟩, then M^w = M^w_f and the argument is trivial. If w = ω, then by errcons(M^w), we have M^w = M^w_ω and hence M^w_f = M^w_ω = M^w, from which the argument directly follows. If w = ⊥, then M^w_f = M^w[E◊ℓ ↦ ⟨c, E◊ℓ⟩], from which det_e(M^w_f), errcons(M^w_f), envcons(M^w_f), and constcons_m(M^w_f) immediately follow. Lastly, pvcons(M^w_f) and csub_m(M^w_f) hold by Lemma 78. Note that envcons_{M^w_f}(E) holds trivially.
- e = x. Let $w_x = M^{\mathsf{w}}(E(x))$ and $w = M^{\mathsf{w}}(E \diamond \ell)$. If $w_x = \omega \lor w = \omega$, then $M^{\mathsf{w}} = M^{\mathsf{w}}_{\omega}$ and by Lemma 47 it follows that $M^{\mathsf{w}}_f = M^{\mathsf{w}}_{\omega}$ and $\mathsf{cons}_m(M^{\mathsf{w}}_f)$ trivially. Also note that $\mathsf{constcons}_m(M^{\mathsf{w}}_f)$ holds trivially.

Next, note that by the def. of pred, we have $\operatorname{pred}_{M^{\mathsf{w}}}^{m}(E(x), E \diamond \ell)$. By Lemma 81, it is the case that $w = \bot$ or $\operatorname{agree}_{E \diamond \ell}(w_x, w)$. In case $w = \bot$, $\operatorname{det}^{m}(M_f^{\mathsf{w}})$ follows by $\operatorname{pred}_{M^{\mathsf{w}}}^{m}(E(x), E \diamond \ell)$ and Lemma 78 whereas $\operatorname{pvcons}(M_f^{\mathsf{w}})$ follows by the def. of $\operatorname{prop}^{\mathsf{w}}$ and ext. In the other case, $\operatorname{det}^{m}(M_f^{\mathsf{w}})$ and $\operatorname{pvcons}(M_f^{\mathsf{w}})$ follow by Lemma 56.

Since $w = \bot$ or $\operatorname{agree}_{E \diamond \ell}(w_x, w)$, by $\operatorname{cons}_m(M^w)$ and the definition of $<:^w$, we have that $w_x <:^w w$. Then by Lemma 77 it follows that $\operatorname{errcons}(M_f^w)$. This lemma and Lemma 47 imply $\operatorname{envcons}(M_f^w)$ and $\operatorname{envcons}_{M_f^w}(E)$.

We now turn to proving that $\operatorname{csub}_m(M_f^w)$. By Lemma 68, we have that $w'_x <: {}^ww'$. For any path at which the pair $\langle E(x), E \diamond \ell \rangle$ lies, the subtyping

across the whole path is not invalidated by Lemma 69 and $\operatorname{cons}_m(M_f^{\mathsf{w}})$. As paths can branch on variable nodes by Lemma 76, the subtyping between E(x) and successive nodes in other paths on which E(x) lies holds by Lemma 82. Next, suppose $e_1^{\ell_1} x^{\ell} \in m$ for some e_1 and ℓ_1 . If $M^{\mathsf{w}}(E \diamond \ell_1) =$ $\langle T^{\mathsf{w}}, p_1 \rangle$, for some $T^{\mathsf{w}} \in \mathcal{T}^{\mathsf{w}}$ and path p_1 , then we need to show that $w' <:^{\mathsf{w}} \pi_2(T^{\mathsf{w}}(p_1))$. We know by the def. of csub_m that $w <:^{\mathsf{w}} \pi_1(T^{\mathsf{w}}(p_1))$ and, from earlier, that $w <:^{\mathsf{w}} w'_x$. Then, the argument follows by Lemma 69. The similar reasoning applies for other places where x can appear in m.

e = e₁^{ℓ₁}e₂<sup>ℓ₂</sub>. Clearly (e₁, E) is well-formed and domvalid(ℓ₁, m, E) so by the i.h., we have that cons_m(M₁^w) and envcons_{M₁^w}(E). Also, (e₂, E) is well-formed and domvalid(ℓ₂, m, E) so we have by the i.h. that cons_m(M₂^w) and envcons_{M₂^w}(E). By Lemma 50 and Lemma 51 we have that w₁ = M₂^w(E ◊ ℓ₁). Hence, we have that w₁<:^w[path(w₁) ↦ ⟨w₂, w⟩], where w = M₂^w(E ◊ ℓ₁), by the def. of csub_m(M₂^w). Then, envcons(M_f^w) and envcons_{M_f^w}(E) hold trivially, errcons(M_f^w) follows by Lemma 77, and constcons(M_f^w) holds by the def. of prop^w. That is, suppose m(ℓ₁) is a lambda expression. Then by Lemma 47, w₁ ∈ T^w ∪ {ω}. However, w₁['] ≠ ω by Lemma 77.
</sup>

We next show that $\operatorname{csub}_m(M_f^w)$. First, let $w_o = \pi_2(T^w(\operatorname{path}(w_1)))$. We know that $w_o <:^w w$. If $w_o = \bot$, then $w = \bot$. A more interesting case is when $w_o = \langle \neg, p_o \rangle$. By $\operatorname{outcons}_m(M_2^w)$, we have that $\operatorname{top}(p_o) = E_b \diamond \ell_b$ where $E_b =$ $E_0.x: E_0 \diamond n_c \diamond x.f: E_0 \diamond n_c \diamond f, E_0 = env(\operatorname{root}(\operatorname{path}(w_1))), E_c = \operatorname{top}(\operatorname{path}(w_1)),$ and $m(\operatorname{loc}(\operatorname{root}(\operatorname{path}(w_1)))) = \mu f.\lambda x.e_b^{\ell_b}$ for some $e_b, f, x, \text{ and } \ell_b$. Hence, $\operatorname{pred}_{M_2^w}^m(E_b \diamond \ell_b, E \diamond \ell)$ by the def. of pred. By Lemma 81, it is the case that $w = \bot$ or $\operatorname{agree}_{E \diamond \ell}(M_2^w(E_b \diamond \ell_b), w)$. We next need to show that $M_f^{\mathsf{w}}(E_b \diamond \ell_b) <: {}^{\mathsf{w}}w'$. Note that $M_f^{\mathsf{w}}(E_b \diamond \ell_b) = M_2^{\mathsf{w}}(E_b \diamond \ell_b)$. We know that $w_o <: {}^{\mathsf{w}}w$ by $\mathsf{csub}_m(M_2^{\mathsf{w}})$ and hence $w'_o <: {}^{\mathsf{w}}w'$ by Lemma 68. By Lemma 83, we have that $M_2^{\mathsf{w}}(E_b \diamond \ell_b) <: {}^{\mathsf{w}}w_o$. Therefore, by Lemma 69 we have $M_2^{\mathsf{w}}(E_b \diamond \ell_b) <: {}^{\mathsf{w}}w'_o$ and thus by Lemma 67 it must be that $M_2^{\mathsf{w}}(E_b \diamond \ell_b) <: {}^{\mathsf{w}}w'$. The same argument applies for the updated input of w'_1 . The remaining argument for $\mathsf{csub}_m(M_f^{\mathsf{w}})$ as well as $\mathsf{det}_m(M_f^{\mathsf{w}})$, $\mathsf{pvcons}(M_f^{\mathsf{w}})$, is the same as in the case of program variables.

• $e = \mu f \lambda x. e_1^{\ell_1}$. First, observe that for any $p_1, p_2 \in T^{\mathsf{w}}$ where $p_1 \neq p_2$, we have that $\mathsf{top}(p_1) \neq \mathsf{top}(p_2)$ by Lemma 79. Hence, for any $p \in T^{\mathsf{w}}$, we have unique n_x, n_f , and E_1 . By $\mathsf{csub}_m(M^{\mathsf{w}})$, we know that $T^{\mathsf{w}} <:^{\mathsf{w}} M^{\mathsf{w}}(n_f)$ and $[p \mapsto \langle M^{\mathsf{w}}(n_x), M^{\mathsf{w}}(E_1 \diamond \ell_1) \rangle] <:^{\mathsf{w}} T^{\mathsf{w}}$. Since $\pi_1(T^{\mathsf{w}}(p)) \neq \bot$, we have by the definition of pred and $\mathsf{inpcons}_m(M^{\mathsf{w}})$ that $\mathsf{pred}_{M^{\mathsf{w}}}^m(\mathsf{top}(\mathsf{path}(\pi_1(T^{\mathsf{w}}(p)))), n_x)$ and by the def. of pred we also have $\mathsf{pred}_{M^{\mathsf{w}}}^m(E \diamond \ell, n_f)$. The argument then goes as usual to show that $\mathsf{cons}_m(M_1^{\mathsf{w}})$. Note that in order to show that $\mathsf{csub}_m(M_1^{\mathsf{w}})$ we also additionally use Lemma 70 to show that the join $T_1^{\mathsf{w}} \sqcup^{\mathsf{w}} T_2^{\mathsf{w}}$ does not invalidate the subtyping on the paths at which n_f appears.

Finally, we switch to proving that the join of all such M_1^w , resulting in M_f^w , is consistent. First, $\operatorname{errcons}(M_f^w)$ follows by Lemma 54 and Lemma 88. Next, $\operatorname{envcons}(M_f^w)$ and $\operatorname{envcons}_{M_f^w}(E)$ simply follows by the def. of $\dot{\sqcup}^w$. Similarly, $\operatorname{det}_m(M_f^w)$ follows from the def. of $\dot{\sqcup}^w$. Finally, $\operatorname{csub}_m(M_f^w)$ and $\operatorname{pvcons}(M_f^w)$ follows by Lemma 54, Lemma 86, and Lemma 87.

• $e = e_0^{\ell_0} ? e_1^{\ell_1} : e_2^{\ell_2}$. Similar to the earlier cases.

A.3.5 Connection to Data Flow Semantics

In this subsection, we make the formal connection between path data flow semantics and the regular data flow semantics without the path information. To that end, we first introduce the notion on uniqueness tpcsuniq(w) of a value w in terms of the uniqueness of paths subject to their top nodes.

$$\begin{aligned} \mathsf{tpcsuniq}(w) & \iff (w \in \{\bot, \omega\} \cup Cons^{\mathsf{w}}) \lor (\exists T^{\mathsf{w}}. w = \langle T^{\mathsf{w}}, \lrcorner\rangle \land (\forall p_1, p_2 \in T^{\mathsf{w}}.\\ \mathsf{top}(p_1) = \mathsf{top}(p_2) \Longrightarrow p_1 = p_2) \land \forall p \in T^{\mathsf{w}}.\\ \mathsf{tpcsuniq}(\pi_1(T^{\mathsf{w}}(p))) \land \mathsf{tpcsuniq}(\pi_2(T^{\mathsf{w}}(p)))) \end{aligned}$$

We first prove the results stating that two paths drawn from a set of root consistent and deterministic paths must have unique top nodes.

Lemma 89. Let P be a set of paths such that for every $p \in P$ we have $\operatorname{rootcons}_e(p)$ and $\det_{M^w}^e(p)$ for some program e and a path map M^w . Then, for any $p_1, p_2 \in P$, $\operatorname{top}(p_1) = \operatorname{top}(p_2)$ implies $p_1 = p_2$.

Proof. Similar to the proof of Lemma 79.

Data propagation does not invalidate the **tpcsuniq** property.

Lemma 90. Let w_1 and w_2 be path values s.t. $\operatorname{tpcsuniq}(w_1)$, $\operatorname{tpcsuniq}(w_2)$, and $P = \operatorname{paths}(w_1) \cup \operatorname{paths}(w_2)$ where for every $p \in P$ we have $\operatorname{rootcons}_e(p)$ and $\operatorname{det}_{M^{w}}^e(p)$ for some program e and a path map M^{w} . Further, let w_3 and w_4 be values such that $\langle w_3, w_4 \rangle = \operatorname{prop}^w(w_1, w_2)$. Then, $\operatorname{tpcsuniq}(w_3)$ and $\operatorname{tpcsuniq}(w_4)$.

Proof. By structural induction on w_1 and w_2 . The base cases are trivial. For the inductive case, let $w_1 = \langle T_1^{\mathsf{w}}, p_1 \rangle$ and $w_2 = \langle T_2^{\mathsf{w}}, p_2 \rangle$ where $p_1 \leq p_2$. Let $p \in T_2^{\mathsf{w}}$ for some p. By Lemma 89, there is no different $p' \in T_2^{\mathsf{w}}$ such that $\mathsf{top}(p) = \mathsf{top}(p')$. Also, for every $p' \in T_1^{\mathsf{w}}$ s.t. $\mathsf{top}(p') = \mathsf{top}(p)$, we have by Lemma 89 that p = p'. In other words, by Lemma 89 the call site paths both in w_1 and w_2 can be uniquely and consistently identified with their top nodes. The argument then follows by i.h.

Next, we present the conversion function strip : $\mathcal{V}^{w} \rightarrow \mathcal{V}$ from path data flow values to regular data flow values that simply strips the path information.

$$\begin{aligned} \operatorname{strip}(\bot) \stackrel{\text{def}}{=} \bot & \operatorname{strip}(\omega) \stackrel{\text{def}}{=} \omega & \operatorname{strip}(\langle c, p \rangle) \stackrel{\text{def}}{=} c \\ \operatorname{strip}(\langle T^{\mathsf{w}}, p \rangle) \stackrel{\text{def}}{=} \Lambda n_{\mathsf{cs}}. \text{ if } \exists p' \in T^{\mathsf{w}}. \operatorname{top}(p') = n_{\mathsf{cs}} & \text{ if } \operatorname{tpcsuniq}(\langle T^{\mathsf{w}}, p \rangle) \\ & \operatorname{then} \langle \operatorname{strip}(\pi_1(T^{\mathsf{w}}(p'))), \operatorname{strip}(\pi_2(T^{\mathsf{w}}(p'))) \rangle \\ & \operatorname{else} \langle \bot, \bot \rangle \end{aligned}$$

The strip function on maps $strip : \mathcal{M}^w \rightharpoonup \mathcal{M}$ is defined in the expected way.

Using the strip function, we next connect the subtyping relation between path values and plain data flow values.

Lemma 91. Let w_1 and w_2 be path values and $P = \mathsf{paths}(w_1) \cup \mathsf{paths}(w_2)$ where for every $p \in P$ we have $\mathsf{rootcons}_e(p)$ and $\mathsf{det}^e_{M^w}(p)$ for some program e and a path map M^w . Further, assume that $\mathsf{tpcsuniq}(w_1)$ and $\mathsf{tpcsuniq}(w_2)$. Then, $w_1 <:^w w_2$ implies $\mathsf{strip}(w_1) <:^c \mathsf{strip}(w_2)$.

Proof. First, note that $tpcsuniq(w_1)$ implies $w_1 \in dom(strip)$. The same holds for w_2 . The proof goes by structural induction on w_1 and w_2 . We first cover the base cases.

w₁ = ⊥, w₁ ∈ V^w. The result immediately follows by the def. of <:^c since strip(⊥) = ⊥.

• $w_1 = \omega$.

$$w_2 = \omega \qquad \text{by def. of } <:^{\mathsf{w}}$$
$$\mathsf{strip}(w_1) = \mathsf{strip}(w_2) = \omega \qquad \text{by def. of strip}$$
$$\bullet w_1 = \langle c, p_1 \rangle, w_2 = \langle c, p_2 \rangle, p_1 \preceq p_2.$$
$$\mathsf{strip}(w_1) = \mathsf{strip}(w_2) = c \qquad \text{by def. of strip}$$

For the inductive case, let $w_1 = \langle T_1^w, p_1 \rangle$ and $w_2 = \langle T_2^w, p_2 \rangle$ where $p_1 \leq p_2$. Let $p \in T_2^w$ for some p. By Lemma 89, there is no different $p' \in T_2^w$ such that top(p) = top(p'). Also, for every $p' \in T_1^w$ s.t. top(p') = top(p), we have by Lemma 89 that p = p'. The argument then follows by i.h.

We next show that the propagation between path values and their stripped counterparts in the data flow semantics yields values that are again related by the strip function.

Lemma 92. Let w_1 and w_2 be path values such that $w_1 <: w_2$ and $P = paths(w_1) \cup paths(w_2)$ where for every $p \in P$ we have $rootcons_e(p)$ and $det^e_{M^w}(p)$ for some program e and a path map M^w . Further, assume that $tpcsuniq(w_1)$, $tpcsuniq(w_2)$, and $\langle w'_1, w'_2 \rangle = prop^w(w_1, w_2)$ for some path values w'_1 and w'_2 . Lastly, let $\langle v'_1, v'_2 \rangle = prop(strip(w_1), strip(w_2))$. Then, $w'_1 \in dom(strip), w'_2 \in dom(strip), strip(w'_1) = v'_1$, and $strip(w'_2) = v'_2$.

Proof. By Lemma 90, we have $\mathsf{tpcsuniq}(w'_1)$ and $\mathsf{tpcsuniq}(w'_2)$. By the def. of strip , it must also be that $w_1 \in \mathsf{dom}(\mathsf{strip})$ and $w_1 \in \mathsf{dom}(\mathsf{strip})$. The proof continues by the structural induction on w_1 and w_2 . We cover the base cases first.

• $w_1 = \bot, w_2 = \bot$.

$$w'_1 = w'_2 = \bot$$
 by def. of prop^w
 $\mathsf{strip}(w_1) = \bot, \mathsf{strip}(w_2) = \bot$ by def. of strip
 $v'_1 = v'_2 = \bot$ by def. of prop

•
$$w_1 = \langle c, p \rangle, w_2 = \bot.$$

 $w'_1 = \langle c, p \rangle, w'_2 = \langle c, p \rangle$ by def. of prop^w
strip $(w_1) = c$, strip $(w_2) = \bot$ by def. of strip
 $v'_1 = c, v'_2 = c$ by def. of prop

$$\begin{split} w_1 &= \langle T^{\mathsf{w}}, p \rangle, w_2 = \bot. \\ w'_1 &= w_1, w'_2 = \langle T^{\mathsf{w}}_{\bot}, p \rangle & \text{by def. of prop}^{\mathsf{w}} \\ v'_1 &= \mathsf{strip}(w_1), v'_2 = T_{\bot} & \text{by def. of prop} \end{split}$$

•
$$w_1 = \langle c, p_1 \rangle, w_2 = \langle c, p_2 \rangle, p_1 \preceq p_2.$$

 $w'_1 = \langle c, p_1 \rangle, w'_2 = \langle c, p_2 \rangle$ by def. of prop^w
strip $(w_1) = c$, strip $(w_2) = c$ by def. of strip
 $v'_1 = c, v'_2 = c$ by def. of prop

• $w_1 = \omega, w_2 = \omega.$

•

$$\begin{split} w_1' &= \omega, \omega & \text{by def. of } \mathsf{prop}^\mathsf{w} \\ \mathsf{strip}(w_1) &= \omega, \mathsf{strip}(w_2) &= \omega & \text{by def. of } \mathsf{strip} \\ v_1' &= \omega, v_2' &= \omega & \text{by def. of } \mathsf{prop} \end{split}$$

For the inductive case, let $w_1 = \langle T_1^{\mathsf{w}}, p_1 \rangle$ and $w_2 = \langle T_2^{\mathsf{w}}, p_2 \rangle$ where $p_1 \leq p_2$. Let $p \in T_2^{\mathsf{w}}$ for some p. By Lemma 89, there is no different $p' \in T_2^{\mathsf{w}}$ such that $\mathsf{top}(p) = \mathsf{top}(p')$. Also, for every $p' \in T_1^{\mathsf{w}}$ s.t. $\mathsf{top}(p') = \mathsf{top}(p)$, we have by Lemma 89 that

p = p'. Hence, the call site paths seen in T_1^w and T_2^w are uniquely identified by their top nodes. The argument then follows by def. of paths, tpcsuniq, and i.h. \Box

Lemma 93. Let w_1 and w_2 be path values such that $P = \mathsf{paths}(w_1) \cup \mathsf{paths}(w_2)$ where for every $p, p_2 \in P$ we have $\mathsf{top}(p_1) = \mathsf{top}(p_2) \Longrightarrow p_1 = p_2$. Further, assume that $\mathsf{tpcsuniq}(w_1)$, $\mathsf{inpcons}_e(w_1)$, $\mathsf{outcons}_e(w_1)$, $\mathsf{tiocons}(w_1)$, $\mathsf{tpcsuniq}(w_2)$, $\mathsf{inpcons}_e(w_2)$, $\mathsf{outcons}_e(w_2)$, and $\mathsf{tiocons}(w_2)$. Further, assume $w_1 \in \mathsf{dom}(\mathsf{strip})$, $w_2 \in \mathsf{dom}(\mathsf{strip})$, $\mathsf{strip}(w_1) = v_1$, and $\mathsf{strip}(w_2) = v_2$. Lastly, we assume that w_1 and w_2 agree on the paths. Then, $w_1 \sqcup^{\mathsf{w}} w_2 \in \mathsf{dom}(\mathsf{strip})$ and $\mathsf{strip}(w_1 \sqcup^{\mathsf{w}} w_2) = v_1 \sqcup v_2$.

Proof. We carry the proof by structural induction on w_1 and w_2 . We first cover the base cases.

• $w_1 = \bot$.

by def. of \sqcup^{w}	$w_1 \sqcup^{w} w_2 = w_2$
by def. of strip	$v_1 = \bot$
by def. of \sqcup	$v_1 \sqcup v_2 = v_2 = \operatorname{strip}(w_2) = \operatorname{strip}(w_1 \sqcup^{w} w_2)$

- $w_2 = \bot$. Similar to the previous case.
- $w_1 = \omega$.

by def. of \sqcup^{w}	$w_1 \sqcup^{w} w_2 = \omega$
by def. of \sqcup	$v_1 = \omega$
by def. of \sqcup	$v_1 \sqcup v_2 = \omega$
by def. of strip	$strip(\omega) = \omega$

	• $w_1 = \langle c_1, p \rangle, w_2 = \langle c_2, p \rangle, c_1 \neq c_2.$
by def. of $\sqcup^{\sf w}$	$w_1 \sqcup^{\sf w} w_2 = \omega$
by def. of strip	$v_1 = c_1, v_2 = c_2$
by def. of \sqcup	$v_1 \sqcup v_2 = \omega$
by def. of strip	$strip(\omega) = \omega$

$$w_1 \sqcup^w w_2 = \langle c, p \rangle$$
 by def. of \sqcup^w
 $v_1 = c, v_2 = c$ by def. of strip
 $v_1 \sqcup v_2 = c$ by def. of \sqcup

$$strip(\langle c, p \rangle) = c$$
 by def. of strip

$$u_1 = \langle c, p \rangle, w_2 = \langle T^{\mathsf{w}}, p \rangle.$$

$$w_1 \sqcup^{\mathsf{w}} w_2 = \omega \qquad \text{by def. of } \sqcup^{\mathsf{w}}$$

$$v_1 = c, v_2 \in \mathcal{T} \qquad \text{by def. of strip}$$

$$v_1 \sqcup v_2 = \omega \qquad \text{by def. of } \sqcup$$

$$\mathsf{strip}(\omega) = \omega \qquad \text{by def. of strip}$$

• Other. Similar for the other cases.

• $w_1 = \langle c, p \rangle, w_2 = \langle c, p \rangle.$

• w

For the inductive case, let $w_1 = \langle T_1^{\mathsf{w}}, p_1 \rangle$ and $w_2 = \langle T_2^{\mathsf{w}}, p_2 \rangle$ where $p_1 \leq p_2$. Let $p \in T_2^{\mathsf{w}}$ for some p. By Lemma 89, there is no different $p' \in T_2^{\mathsf{w}}$ such that $\mathsf{top}(p) = \mathsf{top}(p')$. Also, for every $p' \in T_1^{\mathsf{w}}$ s.t. $\mathsf{top}(p') = \mathsf{top}(p)$, we have by Lemma 89 that p = p'. Hence, the call site paths seen in T_1^{w} and T_2^{w} are uniquely identified by their top nodes.

• $p \in T_1^w, p \in T_2^w$. Let $T_1^w(p) = \langle w_{1i}, w_{1o} \rangle$ and $T_2^w(p) = \langle w_{2i}, w_{2o} \rangle$. If $w_{1i}, w_{2i} \notin \{\perp, \omega\}$, then $\mathsf{path}(w_{2i}) = \mathsf{path}(w_{1i})$ by $\mathsf{path}(w_{2i}), \mathsf{path}(w_{1i}) \in P$ (by def. of

paths), Lemma 89, and $top(path(w_{2i})) = top(path(w_{1i}))$ which follows by the def. of inpcons. The same property holds for the outputs by the same reasoning and def. of outcons. The argument then follows from the i.h.

- $p \in T_1^w, p \notin T_2^w$. By tiocons (T_2^w) , we then have that $T_2^w(p) = \langle \bot, \bot \rangle$. The argument then follows trivially since $\operatorname{strip}(w_2)(\operatorname{top}(p)) = \langle \bot, \bot \rangle$.
- Other cases. Handled similarly.

We can now establish a result connecting the path based data flow semantics and plain data flow semantics. For that purpose, we first state an almost obvious result.

Lemma 94. Let M^{w} be a consistent map $\mathsf{cons}_e(M^{\mathsf{w}})$ subject to a program e. Then, $M^{\mathsf{w}} \in \mathsf{dom}(\mathsf{strip}).$

Proof. The result follows from Lemma 79.

We can now state our main theorem.

Theorem 8. Let e^{ℓ} be an expression of a program m, M^{w} a consistent map $\operatorname{cons}_m(M^{\mathsf{w}})$, E an environment where $\operatorname{domvalid}(\ell, m, E)$, $\operatorname{envcons}_{M^{\mathsf{w}}}(E)$, and (e, E) is well-formed. Further, let M be a concrete map such that $M = \operatorname{strip}(M^{\mathsf{w}})$ and k a step index fuel. If $\operatorname{step}_k^{\mathsf{w}}[e^{\ell}](E)(M^{\mathsf{w}}) = M_f^{\mathsf{w}}$ and $\operatorname{step}_k[e^{\ell}](E)(M) = M_f$, for some maps M_f^{w} and M_f , then $M_f^{\mathsf{w}} \in \operatorname{dom}(\operatorname{strip})$ and $\operatorname{strip}(M_f^{\mathsf{w}}) = M_f$.

Proof. The proof goes by natural induction on k and structural induction on e. When k = 0, the argument is trivial. Otherwise, the proof branches on e.

- e = c. Let $w = M^{\mathsf{w}}(E \diamond \ell)$, $w_f = M_f^{\mathsf{w}}(E \diamond \ell)$, $v = M(E \diamond \ell)$, and $v_f = M_f(E \diamond \ell)$. By $\operatorname{cons}_m(M^{\mathsf{w}})$, $w \in \{\bot, \omega, \langle c, E \diamond \ell \rangle\}$. If $w = \bot$, then $w_f = \langle c, E \diamond \ell \rangle$, $v = \bot$, $v_f = c$, and the result follows from the definition of strip. If $w = \langle c, E \diamond \ell \rangle$, then $w_f = w$ and $v_f = v$, from which the result follows immediately. Lastly, if $w = \omega$, then by $\operatorname{errcons}(M^{\mathsf{w}})$ (implied by $\operatorname{cons}_e(M^{\mathsf{w}})$) we have $M^{\mathsf{w}} = M_{\omega}^{\mathsf{w}}$ and hence $M_f^{\mathsf{w}} = M_{\omega}^{\mathsf{w}}$, $M = M_{\omega}$, as well as $M_f = M_{\omega}$. The result then follows simply from the definition of strip.
- e = x. Let w_x = M^w(E(x)), w = M^w(E◊ℓ), v_x = M(E(x)), and v = M(E◊ℓ). By csub_m(M^w) (implied by cons_m(M^w)) and Lemma 81, we have w_x <:^w w. Further, by cons_m(M^w), we have tpcsuniq(w₁), tpcsuniq(w₂), and that all paths in paths(w₁) ∪ paths(w₂) are deterministic and root consistent. The result then follows from Lemma 92.
- e = e₁^{ℓ₁}e₂^{ℓ₂}. Clearly, we have domvalid(ℓ₁, m, E) and so by i.h. M₁^w ∈ dom(strip) and M₁ = strip(M₁^w). By Theorem 7 it is the case that envcons_{M₁^w}(E) and cons_m(M₁^w). As domvalid(ℓ₂, m, E) trivially, we have by i.h. that M₂^w ∈ dom(strip) and M₂ = strip(M₂^w). Also, again by Theorem 7 we have envcons_{M₂^w}(E) and cons_m(M₂^w).

By Lemmas 50 and 51, we have $w_1 = M_2^{\mathsf{w}}(E \diamond \ell_1)$ and so by $\operatorname{csub}_m(M_2^{\mathsf{w}})$ (implied by $\operatorname{cons}_m(M_2^{\mathsf{w}})$) it must be $w_1 <:^{\mathsf{w}} w'$ where $w' = \langle [\operatorname{path}(w_1) \mapsto \langle w_2, M_2^{\mathsf{w}}(E \diamond \ell_2) \rangle]$, $\operatorname{path}(w_1) \rangle$. Note that by the def. of strip and $M_2 = \operatorname{strip}(M_2^{\mathsf{w}})$, we have $\operatorname{strip}(w') = [\operatorname{top}(\operatorname{path}(w_1)) \mapsto \langle v_2, M_2(E \diamond \ell_2) \rangle]$. Further, by $\operatorname{cons}_m(M^{\mathsf{w}})$, we have $\operatorname{tpcsuniq}(w_1)$, $\operatorname{tpcsuniq}(w')$, and that all paths in $\operatorname{paths}(w_1) \cup \operatorname{paths}(w')$ are deterministic and root consistent. The result then follows from Lemma 92.

•
$$e = \mu f \cdot \lambda x \cdot e_1^{\ell_1}$$
. Let $w = M^w(E \diamond \ell)$ and $v = M(E \diamond \ell)$. By $\operatorname{cons}_m(M^w)$,

 $w \in \{\perp, \omega, \langle T^{\mathsf{w}}, E \diamond \ell \rangle\}$ for some $T^{\mathsf{w}} \in \mathcal{T}^{\mathsf{w}}$. If $w = \bot$, then $T_w = \langle T^{\mathsf{w}}_{wbot}, E \diamond \ell \rangle$, $v = \bot, T = T_{\bot}$. In that case, $M^{\mathsf{w}}_f = M^{\mathsf{w}}[E \diamond \ell \mapsto T_w], M_f = M[E \diamond \ell \mapsto T]$, the result follows from the definition of strip. If $w = \omega$, then by $\operatorname{errcons}(M^{\mathsf{w}})$ (implied by $\operatorname{cons}_e(M^{\mathsf{w}})$) we have $M^{\mathsf{w}} = M^{\mathsf{w}}_{\omega}$ and hence $M^{\mathsf{w}}_f = M^{\mathsf{w}}_{\omega}$ by Lemma 49, $M = M_{\omega}$ and thus $M_f = M_{\omega}$ by Lemma 27. The result then follows simply from the definition of strip. Lastly, if $w = \langle T^{\mathsf{w}}, E \diamond \ell \rangle$, then $T_w = w$ and T = v. The argument then goes as follows.

By Lemma 79 and def. of strip, we have $\{n_{cs} \mid n_{cs} \in T\} = \{top(p) \mid p \in T^{w}\}$. That is, for every $p \in T^{w}$ there is a unique $n_{cs} \in T$ such that $top(p) = n_{cs}$, and vice versa. Then, by $csub_m(M^w)$ (implied by $cons_m(M^w)$) it is that $\langle [p \mapsto \langle M^w(n_x), M^w(E_1 \diamond \ell_1) \rangle], p \rangle <:^w T_w$. Similarly, by $csub_m(M^w)$ and Lemma 81 we have $T_w <:^w M^w(n_f)$. Also, by $cons_m(M^w)$ we have $tpcsuniq(M^w(n_x))$, $tpcsuniq(M^w(E_1 \diamond \ell_1)), tpcsuniq(T_w), tpcsuniq(M^w(n_f)), and all paths in paths<math>(M^w(n_x)) \cup$ paths $(M^w(n_f)) \cup paths(T_w) \cup pathsM^w(E_1 \diamond \ell_1)$ are deterministic and root consistent. Then by Lemma 92 and Lemma 93 we have $M_1^w \in dom(strip)$ and $strip(M_1^w) = M_1$. By Lemmas 47 and 77 as well as $envcons_{M^w}(E)$ we have $envcons_{M_1^w}(E_1)$. Also, $domvalid(\ell_1, m, E_1)$ holds by the definition of E_1 and $domvalid(\ell, m, E)$. Then by the i.h., $step_k^w[e_1]](E_1)(M_1^w) \in dom(strip)$ and $strip(step_k^w[e_1](E_1)(M_1^w)) = step_k[e_1]](E_1)(M_1)$.

Finally, observe that for each two different $M_1^{\mathsf{w}}, M_2^{\mathsf{w}} \in \vec{M^{\mathsf{w}}}, M_1^{\mathsf{w}}$ and M_2^{w} agree on all nodes except possibly those in $\operatorname{rng}(E)$ by Lemma 52. By $\operatorname{envcons}_{M^{\mathsf{w}}}(E)$, for every node n in $\operatorname{rng}(E)$ we have $M^{\mathsf{w}}(n) \notin \{\perp, \omega\}$. Then by Lemma 56 we have $\operatorname{path}(M^{\mathsf{w}}(n)) = \operatorname{path}(M_1^{\mathsf{w}}(n))$, $\operatorname{path}(M^{\mathsf{w}}(n)) = \operatorname{path}(M_2^{\mathsf{w}}(n))$, for every node n in $\operatorname{rng}(E)$, which clearly implies $\operatorname{path}(M_2^{\mathsf{w}}(n)) = \operatorname{path}(M_1^{\mathsf{w}}(n))$. The result then follows from $\operatorname{fin}(M^{\mathsf{w}})$, which implies $|\{\operatorname{top}(p) \mid p \in T^{\mathsf{w}}\}|$ is finite, and Lemma 93.

•
$$e = e_0^{\ell_0} ? e_1^{\ell_1} : e_2^{\ell_2}$$
. Similar to the earlier cases.

We additionally show that consistency of a path map implies that its stripped plain data flow map has the data flow invariant property.

Lemma 95. Let M^{w} be consistent map $\mathsf{cons}_e(M^{\mathsf{w}})$ subject to a program e. Then, $df(\mathsf{strip}(M^{\mathsf{w}})).$

Proof. First, observe that by Lemma 94 we have $M^{\mathsf{w}} \in \mathsf{dom}(\mathsf{strip})$. Next, let $M = \mathsf{strip}(M^{\mathsf{w}})$. Then, $\mathsf{errcons}(M)$ follows by $\mathsf{errcons}(M^{\mathsf{w}})$ (implied by $\mathsf{cons}_e(M^{\mathsf{w}})$) and the definition of strip . The same argument can be given to show $\mathsf{fin}(M)$, $\mathsf{tiocons}(M)$, $\mathsf{constcons}_e(M)$, and $\mathsf{envcons}(M)$. Finally, $\mathsf{csub}_e(M)$ follows from $\mathsf{csub}_e(M^{\mathsf{w}})$, $\mathsf{rootcons}_e(M^{\mathsf{w}})$, $\mathsf{rootcons}_e(M^{\mathsf{w}})$, $\mathsf{rootcons}_e(M^{\mathsf{w}})$, $\mathsf{rootcons}_e(M^{\mathsf{w}})$, and $\mathsf{det}_e(M^{\mathsf{w}})$. To see this in more detail, first observe that the set of reachable (non-⊥) nodes of M^{w} and M is the same by the def. of strip . For every reachable node n that corresponds to a non-leaf expression of e, we have that $M^{\mathsf{w}}(n) = \langle -, p \rangle$ for some path p. By $\mathsf{tpcons}(M^{\mathsf{w}})$, $\mathsf{top}(p) = n$ and by $\mathsf{rootcons}_e(M^{\mathsf{w}})$ as well as Lemma 73 we have that |p| > 1. Then, for the predecessor n' of n in p, we have that $M^{\mathsf{w}}(n') < :^{\mathsf{w}} M^{\mathsf{w}}(n)$ by $\mathsf{csub}_e(M^{\mathsf{w}})$ and $\mathsf{pred}_{M^{\mathsf{w}}}^e(n', n)$ by $\mathsf{det}_e(M^{\mathsf{w}})$. The result then follows from the definition of pred , Lemma 91, and the fact that $\mathsf{cons}_e(M^{\mathsf{w}})$ implies $\mathsf{tpcsunig}(M^{\mathsf{w}}(n')) \land \mathsf{tpcsunig}(M^{\mathsf{w}}(n))$.

A.4 Relational Semantics Proofs

We first prove that our definitions of meets and joins over relational values are correct.

Proof (Proof of Lemma 4).

Proof. The proof follows that of Lemma 1.

We next prove the monotonicity of the concretization function for relational values.

Lemma 96. Let $r_1 \in \mathcal{V}_N^{\mathsf{r}}$, $r_2 \in \mathcal{V}_N^{\mathsf{r}}$, and $r_1 \sqsubseteq^{\mathsf{r}} r_2$. Then, $\gamma_N^{\mathsf{r}}(r_1) \subseteq \gamma_N^{\mathsf{r}}(r_2)$.

Proof. The proof goes by structural induction on r_2 . We first cover the base cases.

- $r_2 = \perp^{\mathsf{r}}$. Then $r_1 = \perp^{\mathsf{r}}$ and the result follows trivially.
- $r_2 = \top^r$. Trivially by the def. of $\gamma_N^r(\top^r)$.
- $r_2 = R_2^r$. Then either $r_1 = \perp^r$ or $r_1 = R_1^r$ s.t. $R_1^r \subseteq R_2^r$. Since the former case is trivial due to $\forall r. \gamma_N^r(\perp^r) \subseteq \gamma_N^r(r)$ by the def. of γ_N^r , we focus on the latter case. Here,

$$R_1^{\mathsf{r}} \subseteq R_2^{\mathsf{r}} \Longrightarrow \{ \langle M, c \rangle \mid D^{\mathsf{r}} \in R_1^{\mathsf{r}} \land D^{\mathsf{r}}(\nu) = c \} \subseteq \{ \langle M, c \rangle \mid D^{\mathsf{r}} \in R_2^{\mathsf{r}} \land D^{\mathsf{r}}(\nu) = c \} \Longrightarrow$$
$$\{ \langle M, c \rangle \mid D^{\mathsf{r}} \in R_1^{\mathsf{r}} \land D^{\mathsf{r}}(\nu) = c \land \forall n \in N.M(n) \in \gamma^{\mathsf{d}}(M(n)) \} \subseteq$$
$$\{ \langle M, c \rangle \mid D^{\mathsf{r}} \in R_2^{\mathsf{r}} \land D^{\mathsf{r}}(\nu) = c \land \forall n \in N.M(n) \in \gamma^{\mathsf{d}}(M(n)) \} \Longrightarrow$$
$$\gamma_N^{\mathsf{r}}(r_1) \subseteq \gamma_N^{\mathsf{r}}(r_2).$$

For the induction step $r = T_2^r$, if $r_1 = \bot^r$ the argument is trivial. We hence focus on the case where $r_1 = T_1^r$. Let $T_1^r(n_{cs}) = \langle r_{1i}, r_{1o} \rangle$ and $T_2^r(n_{cs}) = \langle r_{2i}, r_{2o} \rangle$ for any call site n_{cs} . By $T_1^r \sqsubseteq^r T_2^r$, we have $r_{1i} \sqsubseteq^r r_{2i}$ and $r_{1o} \sqsubseteq^r r_{2o}$. By i.h. on r_{2i} and r_{2o} , it follows $\gamma_N^r(r_{1i}) \subseteq \gamma_N^r(r_{2i})$ and $\gamma_N^r(r_{1o}) \subseteq \gamma_N^r(r_{2o})$, which implies $\gamma_N^r(T_1^r) \subseteq \gamma_N^r(T_2^r)$ by the def. of γ_N^r .

It can be also shown that the concretization of relational values does not impose any constraints on the values for nodes not in the given scope.

Lemma 97. For any $n \notin N$, n is unconstrained in $\gamma_N^{\mathsf{r}}(r_N)$ for any relational value r. That is, $\forall n \notin N, v, v', M.\langle M, v \rangle \in \gamma_N^{\mathsf{r}}(r_N) \Longrightarrow \langle M[n \to v'], v \rangle \in \gamma_N^{\mathsf{r}}(r_N).$

Proof. The proof goes by structural induction on r. The base cases are trivial. The induction case for tables follows from the i.h. and the fact that for any table T^{r} and call site n_{cs} , n_{cs} is not in the scope of $\pi_i(T^{r}(n_{cs}))$.

The next proof shows that the concretization function satisfies the meet-morphism requirement.

Theorem 9 (Complete Meet Morphism on Rel. Values). Let $V^r \in \wp(\mathcal{V}_N^r)$. Then $\gamma_N^r(\sqcap^r V^r) = \bigcap_{r \in V^r} \gamma_N^r(r).$

Proof. We carry the proof by case analysis on elements of V^r and induction on the minimum depth thereof when V^r consists of tables only.

We first consider the trivial (base) cases where V^r is not a set of multiple relational tables.

- $V^r = \emptyset$. Here, $\gamma_N^r(\Box^r \emptyset) = \gamma_N^r(\top^r) = \mathcal{M} \times \mathcal{V} = \bigcap \emptyset$.
- $V^r = \{r\}$. Trivial.
- $R^{\mathsf{r}} \in V^{r}, T^{\mathsf{r}} \in V^{r}$. We have $\sqcap^{\mathsf{r}} \{R^{\mathsf{r}}, T^{\mathsf{r}}\} = \bot^{\mathsf{r}}$ and since \bot^{r} is the bottom element, $\sqcap^{\mathsf{r}} V^{r} = \bot^{\mathsf{r}}$. Similarly, $\gamma_{N}^{\mathsf{r}}(R^{\mathsf{r}}) \cap \gamma_{N}^{\mathsf{r}}(T^{\mathsf{r}}) = \gamma_{N}^{\mathsf{r}}(\bot^{\mathsf{r}})$ and since $\forall r. \gamma_{N}^{\mathsf{r}}(\bot^{\mathsf{r}}) \subseteq \gamma_{N}^{\mathsf{r}}(r)$, it follows $\bigcap_{r \in V^{r}} \gamma_{N}^{\mathsf{r}}(r) = \gamma_{N}^{\mathsf{r}}(\bot^{\mathsf{r}}) = \gamma_{N}^{\mathsf{r}}(\sqcap^{\mathsf{r}} V^{r})$.

• $V^r \subseteq \mathcal{R}_N^r$. Here,

$$= \bigcap_{R^{\mathsf{r}} \in V^{r}} \left(\left\{ \langle M, c \rangle \mid D^{\mathsf{r}} \in R^{\mathsf{r}} \land D^{\mathsf{r}}(\nu) = c \land \forall n \in N.M(n) \in \gamma^{\mathsf{d}}(M(n)) \right\} \cup \gamma_{N}^{\mathsf{r}}(\bot^{\mathsf{r}}) \right)$$

[by uniqueness/non-overlapping of $\gamma^{\rm d}]$

• $\top^{\mathsf{r}} \in V^r$. Here, $\sqcap^{\mathsf{r}} V^r = \sqcap^{\mathsf{r}} (V^r / \{\top^{\mathsf{r}}\})$ and $\bigcap_{r \in V^r} \gamma_N^{\mathsf{r}}(r) = \bigcap_{r \in V^r, r \neq \top^r} \gamma_N^{\mathsf{r}}(r)$ since $\gamma_N^{\mathsf{r}}(\top^{\mathsf{r}}) = \mathcal{M} \times \mathcal{V}$. The set $V^r / \{\top^{\mathsf{r}}\}$ either falls into one of the above cases or consists of multiple tables, which is the case we show next.

Let $V^r \subseteq \mathcal{T}_N^r$ and $|V^r| > 1$. Let d be the minimum depth of any table in V^r .

$$\bigcap_{T^{\mathsf{r}} \in V^{\mathsf{r}}} \gamma_{N}^{\mathsf{r}}(T^{\mathsf{r}}) = \{ \langle M, T \rangle \mid \forall n_{\mathsf{cs}}. \ T(n_{\mathsf{cs}}) = \langle v_{i}, v_{o} \rangle \land \langle M_{i}, v_{i} \rangle \in \bigcap_{T^{\mathsf{r}} \in V^{\mathsf{r}}} \gamma^{\mathsf{r}}(\pi_{1}(T^{\mathsf{r}}(n_{\mathsf{cs}}))) \land M_{i} = \langle M_{o}, v_{o} \rangle \in \bigcap_{T^{\mathsf{r}} \in V^{\mathsf{r}}} \gamma^{\mathsf{r}}(\pi_{2}(T^{\mathsf{r}}(n_{\mathsf{cs}}))) \land M_{i} = \langle m_{\mathsf{cs}} \rangle M \land M_{o} = M_{i}[n_{\mathsf{cs}} \mapsto v_{i}] \} \cup \gamma_{N}^{\mathsf{r}}(\bot^{\mathsf{r}})$$

[by def. of $\gamma_N^{\rm r}$ and $n_{\rm cs}$ not in the scope of $\pi_1(T^{\rm r}(n_{\rm cs}))]$

We can now state the monotonicity of abstraction function for the relational domain.

Lemma 98. α_N^r is monotone.

Proof. The result follows from the properties of Galois connections. \Box

The meet morphism requirement is also satisfied by the concretization function defined over relational maps.

Theorem 10 (Complete Meet Morphism on Rel. Maps). Let $m^r \in \wp(\mathcal{M}^r)$. Then $\dot{\gamma}^r(\dot{\sqcap}^r m^r) = \bigcap_{M^r \in m^r} \dot{\gamma}^r(M^r).$

Proof.

$$\begin{split} \bigcap_{M^{\mathsf{r}}\in m^{\mathsf{r}}} \dot{\gamma}^{\mathsf{r}}(M^{\mathsf{r}}) &= \{ M \mid \forall n. \langle M, M(n) \rangle \in \bigcap_{M^{\mathsf{r}}\in m^{\mathsf{r}}} \gamma_{N}^{\mathsf{r}}(M^{\mathsf{r}}(n)) \} \\ &= \{ M \mid \forall n. \langle M, M(n) \rangle \in \gamma_{N}^{\mathsf{r}}(\sqcap_{M^{\mathsf{r}}\in m^{\mathsf{r}}}^{\mathsf{r}}M^{\mathsf{r}}(n)) \} \qquad \text{[by Lemma 9]} \\ &= \dot{\gamma}^{\mathsf{r}}(\dot{\sqcap}^{\mathsf{r}}m^{\mathsf{r}}) \qquad \text{[by def. of } \dot{\gamma}^{\mathsf{r}} \text{ and } \dot{\sqcap}^{\mathsf{r}}] \end{split}$$

A.4.1 Domain Operations

In this subsection, we state the basic properties and the corresponding proofs of strengthening operations on relational values.

A.4.1.1 Strengthening

We first show that our strengthening operations indeed perform strengthening.

Lemma 99. Let $r_1 \in \mathcal{V}_N^r$, $r \in \mathcal{V}_N^r$, and $n \in \mathcal{N}$. Then $r[n \leftarrow r_1] \sqsubseteq^r r$.

Proof. In the below calculational proof, let $S = \{ \langle M'[n \mapsto v'], v \rangle \mid \langle M', v' \rangle \in \}$
$\gamma_N^{\mathsf{r}}(r_1)\}.$

by prop. of Galois conn.	$\alpha_N^{\mathbf{r}}(\gamma_N^{\mathbf{r}}(r)) \sqsubseteq^{\mathbf{r}} r$
by Lemma 98	$\alpha_N^{r}(\gamma_N^{r}(r) \cap S) \sqsubseteq^{r} \alpha_N^{r}(\gamma_N^{r}(r))$
by transitivity	$\alpha_N^r(\gamma_N^r(r)\cap S)\sqsubseteq^r r$
by def. of $\cdot [\cdot \leftarrow \cdot]$	$r[n \leftarrow r_1] \sqsubseteq^{r} r$

Lemma 100. Let $r \in \mathcal{V}_N^r$, $n_1 \in \mathcal{N}$, and $n_2 \in \mathcal{N}$. Then $r[n_1=n_2] \sqsubseteq^r r$.

Proof. Similar to the proof of Lemma 99.

Lemma 101. Let $r \in \mathcal{V}_N^r$ and $n \in \mathcal{N}$. Then $r[\nu=n_2] \sqsubseteq^r r$.

Proof. Similar to the proof of Lemma 99.

Lemma 102. Let $r \in \mathcal{V}_N^r$ and $M^r \in \mathcal{M}^r$. Then $r[M^r] \sqsubseteq^r r$.

Proof. Follows from the properties of meets and Lemma 99. \Box

We can also show that strengthening operations are idempotent. We show the proof of idempotency for strengthening by a relational value.

Lemma 103. Let $r_1 \in \mathcal{V}_N^r$, $r \in \mathcal{V}_N^r$, and $n \in \mathcal{N}$. Then $r[n \leftarrow r_1] = r[n \leftarrow r_1][n \leftarrow r_1]$.

Proof. In the below calculational proof, let $S = \{ \langle M'[n \mapsto v'], v \rangle \mid \langle M', v' \rangle \in$

 $\gamma_N^{\mathsf{r}}(r_1)$ and $r' = r[n \leftarrow r_1] = \alpha_N^{\mathsf{r}}(\gamma_N^{\mathsf{r}}(r) \cap S).$

 $\begin{aligned} r'[n \leftarrow r_1] &= \alpha_N^{\mathsf{r}}(\gamma_N^{\mathsf{r}}(r') \cap S) & \text{by def. of } \cdot [\cdot \leftarrow \cdot] \\ \gamma_N^{\mathsf{r}}(r) \cap S \subseteq \gamma_N^{\mathsf{r}}(r') & \text{by prop. of Galois conn.} \\ \gamma_N^{\mathsf{r}}(r) \cap S \subseteq \gamma_N^{\mathsf{r}}(r') \cap S & \text{by def. of } \cap \\ \alpha_N^{\mathsf{r}}(\gamma_N^{\mathsf{r}}(r) \cap S) \sqsubseteq^{\mathsf{r}} \alpha_N^{\mathsf{r}}(\gamma_N^{\mathsf{r}}(r') \cap S) & \text{by Lemma 98} \\ r[n \leftarrow r_1] \sqsubseteq^{\mathsf{r}} r'[n \leftarrow r_1] & \text{by def. of } \cdot [\cdot \leftarrow \cdot] \\ r[n \leftarrow r_1] \sqsubseteq^{\mathsf{r}} r[n \leftarrow r_1][n \leftarrow r_1] & \text{by def. of } r' \\ r[n \leftarrow r_1][n \leftarrow r_1] \sqsubseteq^{\mathsf{r}} r[n \leftarrow r_1] & \text{by Lemma 99} \\ r[n \leftarrow r_1][n \leftarrow r_1] = r[n \leftarrow r_1] \end{aligned}$

A.4.1.2 Monotonicity

All of the strengthening operations on relational values are monotone.

Lemma 104. Let $r_1 \in \mathcal{V}_N^r$, $r_2 \in \mathcal{V}_N^r$, $r \in \mathcal{V}_N^r$, and $n \in \mathcal{N}$. Then $r_1 \sqsubseteq^r r_2 \implies$ $r_1[n \leftarrow r] \sqsubseteq^r r_2[n \leftarrow r].$

Proof. In the below calculational proof, let $S = \{ \langle M'[n \mapsto v'], v \rangle \mid \langle M', v' \rangle \in \gamma_N^r(r) \}.$

$$\begin{aligned} r_1 &\sqsubseteq^{\mathsf{r}} r_2 \implies \gamma_N^{\mathsf{r}}(r_1) \subseteq \gamma_N^{\mathsf{r}}(r_2) & \text{by Lemma 96} \\ & \longleftrightarrow \gamma_N^{\mathsf{r}}(r_1) \cap S \subseteq \gamma_N^{\mathsf{r}}(r_2) \cap S \\ & \Longrightarrow \alpha_N^{\mathsf{r}}(\gamma_N^{\mathsf{r}}(r_1) \cap S) \sqsubseteq^{\mathsf{r}} \alpha_N^{\mathsf{r}}(\gamma_N^{\mathsf{r}}(r_2) \cap S) & \text{by Lemma 98} \\ & \longleftrightarrow r_1[n \leftarrow r] \sqsubseteq^{\mathsf{r}} r_2[n \leftarrow r] & \text{by def. of } \cdot[\cdot \leftarrow \cdot] \end{aligned}$$

Lemma 105. Let $r_1 \in \mathcal{V}_N^r$, $r_2 \in \mathcal{V}_N^r$, $r'_1 \in \mathcal{V}_{N_0}^r$, $r'_2 \in \mathcal{V}_{N_0}^r$, and $n \in \mathcal{N}$. Then $r_1 \sqsubseteq^r r_2 \wedge r'_1 \sqsubseteq^r r'_2 \implies r_1[n \leftarrow r'_1] \sqsubseteq^r r_2[n \leftarrow r'_2].$

Proof. In the below calculational proof, let $S_1 = \{ \langle M'[n \mapsto v'], v \rangle \mid \langle M', v' \rangle \in \gamma_N^r(r_1') \}$ and $S_2 = \{ \langle M'[n \mapsto v'], v \rangle \mid \langle M', v' \rangle \in \gamma_N^r(r_2') \}.$

First note that $r'_1 \sqsubseteq r'_2$ implies by Lemma 96 that $\gamma'_N(r'_1) \subseteq \gamma'_N(r'_2)$. Hence, $S_1 \subseteq S_2$.

$$r_{1} \sqsubseteq^{\mathsf{r}} r_{2} \implies \gamma_{N}^{\mathsf{r}}(r_{1}) \subseteq \gamma_{N}^{\mathsf{r}}(r_{2}) \qquad \text{by Lemma 96}$$

$$\iff \gamma_{N}^{\mathsf{r}}(r_{1}) \cap S_{1} \subseteq \gamma_{N}^{\mathsf{r}}(r_{2}) \cap S_{2} \qquad \text{by } S_{1} \subseteq S_{2}$$

$$\implies \alpha_{N}^{\mathsf{r}}(\gamma_{N}^{\mathsf{r}}(r_{1}) \cap S_{1}) \sqsubseteq^{\mathsf{r}} \alpha_{N}^{\mathsf{r}}(\gamma_{N}^{\mathsf{r}}(r_{2}) \cap S_{2}) \qquad \text{by mono. of } \alpha_{N}^{\mathsf{r}}$$

$$\iff r_{1}[n \leftarrow r] \sqsubseteq^{\mathsf{r}} r_{2}[n \leftarrow r] \qquad \text{by def. of } \cdot[\cdot \leftarrow \cdot]$$

Lemma 106. Let $r_1 \in \mathcal{V}_N^r$, $r_2 \in \mathcal{V}_N^r$, $n_1 \in \mathcal{N}$, and $n_2 \in \mathcal{N}$. Then $r_1 \sqsubseteq^r r_2 \implies r_1[n_1=n_2] \sqsubseteq^r r_2[n_1=n_2]$.

Proof. Similar as for Lemma 104.

Lemma 107. Let $r_1 \in \mathcal{V}_N^r$, $r_2 \in \mathcal{V}_N^r$, and $n \in \mathcal{N}$. Then $r_1 \sqsubseteq^r r_2 \implies r_1[\nu=n] \sqsubseteq^r r_2[\nu=n]$.

Proof. Similar as for Lemma 104.

Lemma 108. Let $r_1 \in \mathcal{V}_N^r$, $r_2 \in \mathcal{V}_N^r$, and $M^r \in \mathcal{M}^r$. Then $r_1 \sqsubseteq^r r_2 \implies r_1[M^r] \sqsubseteq^r r_2[M^r]$.

Proof. The result follows from Lemma 104 and the monotonicity of meets. \Box

Lemma 109. Let $r_1 \in \mathcal{V}_N^r$, $r_2 \in \mathcal{V}_N^r$, $M_1^r \in \mathcal{M}^r$, and $M_2^r \in \mathcal{M}^r$. Then $r_1 \sqsubseteq^r$ $r_2 \wedge M_1^r \stackrel{:}{\sqsubseteq}^r M_2^r \implies r_1[M_1^r] \sqsubseteq^r r_2[M_2^r].$

Proof. The result follows from Lemma 105 and the monotonicity of meets. \Box

Lemma 110. Let $r_1 \in \mathcal{V}_N^r$, $r_2 \in \mathcal{V}_N^r$, and N' be a scope. Then $r_1 \sqsubseteq^r r_2 \implies r_1[N'] \sqsubseteq^r r_2[N']$.

Proof. Similar as for Lemma 104.

A.4.1.3 Structural Strengthening

We also here provide a structural definition of strengthening and argue its soundness. First, whenever $n \notin N$, then $r_{\hat{N}} \lceil n \leftarrow r \rceil \stackrel{\text{def}}{=} r$. Otherwise,

$$\begin{aligned} r \mid n \leftarrow \perp^{\mathsf{r}} \mid & \stackrel{\text{def}}{=} \perp^{\mathsf{r}} \\ r \lceil n \leftarrow r' \rceil & \stackrel{\text{def}}{=} r & \text{if } r \in \{\perp^{\mathsf{r}}, \top^{\mathsf{r}}\} \text{ or } r' \in \{\top^{\mathsf{r}}\} \cup \mathcal{T}^{\mathsf{r}} \\ R_{1}^{\mathsf{r}} \lceil n \leftarrow R_{2}^{\mathsf{r}} \rceil & \stackrel{\text{def}}{=} R_{1}^{\mathsf{r}} [n \leftarrow R_{2}^{\mathsf{r}}] \end{aligned}$$

 $T^{\mathsf{r}}\lceil n \leftarrow r \rceil \stackrel{\text{def}}{=} \Lambda n_{\mathsf{cs}}.\langle \pi_1(T^{\mathsf{r}}(n_{\mathsf{cs}})) \lceil n \leftarrow r \rceil, \pi_2(T^{\mathsf{r}}(n_{\mathsf{cs}})) \lceil n \leftarrow \exists n_{\mathsf{cs}}.r \rceil \rangle \qquad r \in \mathcal{R}_N^{\mathsf{r}}$ It can be shown that the structural strengthening is an overapproximation of the ordinary one.

Lemma 111. Let $r \in \mathcal{V}^r$, $r' \in \mathcal{V}^r$, and $n \in \mathcal{N}$. Then, $r[n \leftarrow r'] \sqsubseteq^r r[n \leftarrow r']$.

Proof. When $n \notin N$, the argument is trivial by Lemma 99. Otherwise, if $r' = \perp^r$, the result follows by the fact monotonicity of α_N^r and by the definition of γ_N^r . The proof is carried by structural induction on r. First, for the cases where $r' \in \{\top^r\} \cup \mathcal{T}^r$, the result follows immediately by Lemma 99. For the rest of the proof, we thus assume that $r' \in \mathcal{R}_N^r$.

For the base cases when $r \in \{\perp^{\mathsf{r}}, \top^{\mathsf{r}}\}$ the result again follows easily from Lemma 99. For the base case $r \in \mathcal{R}_N^{\mathsf{r}}$, the argument is trivial. For the induction case when $r \in \mathcal{T}_N^{\mathsf{r}}$, we have the following:

$$\{\langle M, T \rangle \mid \forall n_{cs}. T(n_{cs}) = \langle v_i, v_o \rangle \land T^{\mathsf{r}}(n_{cs}) = \langle r_i, r_o \rangle \land \langle M_i, v_i \rangle \in \gamma^{\mathsf{r}}(r_i)$$

$$\land \langle M_o, v_o \rangle \in \gamma^{\mathsf{r}}(r_o) \land M_i =_{\backslash \{n_{cs}\}} M \land M_o = M_i[n_{cs} \mapsto v_i]\} \cap \mathcal{M}$$

where $\mathcal{M} = \{\langle M'[n \mapsto v'], v \rangle \mid \langle M', v' \rangle \in \gamma^{\mathsf{r}}(r')\}$
$$= \{\langle M, T \rangle \mid \forall n_{cs}. T(n_{cs}) = \langle v_i, v_o \rangle \land T^{\mathsf{r}}(n_{cs}) = \langle r_i, r_o \rangle \land \langle M_i, v_i \rangle \in \gamma^{\mathsf{r}}(r_i)$$

$$\land \langle M_o, v_o \rangle \in \gamma^{\mathsf{r}}(r_o) \land M_i =_{\backslash \{n_{cs}\}} M \land M_o = M_i[n_{cs} \mapsto v_i] \land \langle M, T \rangle \in \mathcal{M}\}$$

$$= \{\langle M, T \rangle \mid \forall n_{cs}. T(n_{cs}) = \langle v_i, v_o \rangle \land T^{\mathsf{r}}(n_{cs}) = \langle r_i, r_o \rangle \land \langle M_i, v_i \rangle \in (\gamma^{\mathsf{r}}(r_i) \cap \mathcal{M})$$

$$\land \langle M_o, v_o \rangle \in \gamma^{\mathsf{r}}(r_o) \land M_i =_{\backslash \{n_{cs}\}} M \land M_o = M_i[n_{cs} \mapsto v_i] \land M \in \mathcal{M}\}$$

by n_{cs} not in the scope of r_i and def. of \mathcal{M} (values)

$$= \{ \langle M, T \rangle \mid \forall n_{cs}. \ T(n_{cs}) = \langle v_i, v_o \rangle \land T^{\mathsf{r}}(n_{cs}) = \langle r_i, r_o \rangle \land \langle M_i, v_i \rangle \in (\gamma^{\mathsf{r}}(r_i) \cap \mathcal{M}) \\ \land \langle M_o, v_o \rangle \in (\gamma^{\mathsf{r}}(r_o) \cap \mathcal{M}^{n_{cs}}) \land M_i =_{\backslash \{n_{cs}\}} M \land M_o = M_i[n_{cs} \mapsto v_i] \} \\$$
where $\mathcal{M}^{n_{cs}} = \{ \langle M[n_{cs} \mapsto v'], v \rangle \mid \langle M, v \rangle \in \mathcal{M} \land v' \in \mathcal{V} \}$

$$\subseteq \{ \langle M, T \rangle \mid \forall n_{cs}. \ T(n_{cs}) = \langle v_i, v_o \rangle \land T^{\mathsf{r}}(n_{cs}) = \langle r_i, r_o \rangle \land \langle M_i, v_i \rangle \in \gamma^{\mathsf{r}}(r_i[n \leftarrow r']) \\ \land \langle M_o, v_o \rangle \in \gamma^{\mathsf{r}}(r_o[n \leftarrow \exists n_{cs}. r']) \land M_i =_{\backslash \{n_{cs}\}} M \land M_o = M_i[n_{cs} \mapsto v_i] \}$$
by the def. of $\cdot [\cdot \leftarrow \cdot]$, def. of $\exists \cdot . \cdot$, and Proposition 1

$$= \{ \langle M, T \rangle \mid \forall n_{cs}. \ T(n_{cs}) = \langle v_i, v_o \rangle \land T^{\mathsf{r}}(n_{cs}) = \langle r_i, r_o \rangle \land \langle M_i, v_i \rangle \in \gamma^{\mathsf{r}}(r_i \lceil n \leftarrow r' \rceil) \\ \land \langle M_o, v_o \rangle \in \gamma^{\mathsf{r}}(r_o \lceil n \leftarrow \exists n_{cs}. r' \rceil) \land M_i =_{\backslash \{n_{cs}\}} M \land M_o = M_i [n_{cs} \mapsto v_i] \}$$

by i.h. and Proposition 1

= M'We then have that $r[n \leftarrow r'] \sqsubseteq^{\mathsf{r}} \alpha^{\mathsf{r}}(M' \cup \gamma_N^{\mathsf{r}}(\bot^{\mathsf{r}})) = \alpha_N^{\mathsf{r}}(\gamma_N^{\mathsf{r}}(r\lceil n \leftarrow r'\rceil)) \sqsubseteq^{\mathsf{r}} r\lceil n \leftarrow r'\rceil$ by Proposition 1.

A.4.2 Propagation and Transformer

Proof (of Lemma 6).

Proof. For the increasing property, the proof goes by structural induction on both relational arguments of $prop^r$ and it closely follows the proof of Lemma 26. Regarding monotonicity, the proof goes by structural induction on the relational arguments of $prop^r$, closely follows the proof of Lemma 28, and relies on Lemma 105 (introduced later).

Proof (of Lemma 7).

Proof. The increasing property of $\text{Step}_{k+1}^{\mathsf{r}} \llbracket \cdot \rrbracket$ is obvious. Regarding monotonicity, the proof goes by structural induction on program expressions and it closely follows the proof of Lemma 29. The proof relies on Lemma 6 as well as Lemma 109, 107, 110, and 106.

A.4.2.1 Range of Computation

Lemma 112. Let M^{r} and M_1^{r} be relational maps, E an environment, e^{ℓ} an expression of a program p, and k a fuel. If domvalid (ℓ, p, E) and $\operatorname{Step}_k^{\mathsf{r}}[\![e^{\ell}]\!](E)(M^{\mathsf{r}}) = M_1^{\mathsf{r}}$, then $\forall n.M^{\mathsf{r}}(n) \neq M_1^{\mathsf{r}}(n) \implies n \in \operatorname{range}_p(E \diamond \ell)$.

Proof. The same as in the proof of Lemma 32.

The same result about the range of computation can be shown.

Lemma 113. Let M^{r} and M_1^{r} be relational maps, E an environment, e^{ℓ} an expression of a program p, and k a fuel. If $\mathsf{domvalid}(\ell, p, E)$ and $\mathsf{Step}_k^{\mathsf{r}}[\![e^{\ell}]\!](E)(M) = M_1^{\mathsf{r}}$, then $\forall n.M^{\mathsf{r}}(n) \neq M_1^{\mathsf{r}}(n) \implies n \in \mathsf{ranger}_p(E \diamond \ell)$.

Proof. Similar to the proof of Lemma 35.

A.4.3 Soundness

We first provide several useful results used to argue soundness of relational semantics.

Lemma 114. Let $M \in \mathcal{M}$, $v \in \mathcal{V}$, $n \in \mathcal{N}$, and $r \in \mathcal{V}_N^r$. Suppose $\langle M, v \rangle \in \gamma_N^r(r)$ and $M(n) \in \mathcal{T}$. Then, for any $T \in \mathcal{T}$, $\langle M[n \mapsto T], v \rangle \in \gamma_N^r(r)$.

Proof. When $n \notin N$, the argument follows by Lemma 97. If $r \in \{\perp^r, \top^r\}$, the argument is trivial. For other cases, the result follows by the definition of γ_N^r , specifically γ^d ($\gamma^d(\mathsf{F}) = \mathcal{T}$).

Lemma 115. Let $M \in \mathcal{M}$, $M' \in \mathcal{M}$, $v \in \mathcal{V}$, and $r \in \mathcal{V}_N^r$. Suppose $M \doteq M'$, $\langle M, v \rangle \in \gamma_N^r(r)$, and $\forall n \in N$. $M(n) \neq \omega \implies M'(n) \neq \omega$. Then $\langle M', v \rangle \in \gamma_N^r(r)$.

Proof. The argument goes by structural induction on r. By Lemma 97, the proof analysis can focus only on the nodes N of M'. The cases when $r \in \{\perp^r, \top^r\}$ are trivial. The other cases follow by the definition of γ^r . For instance, suppose $n \in N$ and M(n) = c. Then, M'(n) = c. Otherwise, suppose $M(n) \in \mathcal{T}$. The argument then follows from Lemma 114.

A.4.3.1 Data Propagation

We start proving the soundness of the relational semantics by formalizing and showing that relational propagation abstracts the concrete one. We begin by introducing the notion of scope consistency that closely resembles the environment consistency.

$$\operatorname{scopecons}_N(M) \iff \forall n \in N. M(n) \notin \{\bot, \omega\}$$

Lemma 116. Let v_1 , v_2 , v_3 , and v_4 be values, M_1 and M_2 maps, $r_1 \in \mathcal{V}_N^r$, $r_2 \in \mathcal{V}_N^r$, $r_3 \in \mathcal{V}_N^r$, and $r_4 \in \mathcal{V}_N^r$. Let $v_1 <:^c v_2$, $\mathsf{noerr}(v_1)$, $\mathsf{noerr}(v_2)$, $\langle M_1, v_1 \rangle \in \gamma_N^r(r_1)$, $\langle M_2, v_2 \rangle \in \gamma_N^r(r_2)$, $M_1 \doteq M_2$, $\mathsf{scopecons}_N(M_1)$, and $\mathsf{scopecons}_N(M_2)$. If $\langle v_3, v_4 \rangle = \mathsf{prop}(v_1, v_2)$ and $\langle r_3, r_4 \rangle = \mathsf{prop}^r(r_1, r_2)$, then $\langle M_1, v_3 \rangle \in \gamma_N^r(r_3)$ and $\langle M_2, v_4 \rangle \in \gamma_N^r(r_4)$.

Proof. By structural induction on both v_1 and v_2 .

• $v_1 = \bot, v_2 = \bot.$ $v_3 = \bot, v_4 = \bot$ by def. of prop $\forall r. \mathcal{M} \times \{\bot\} \in \gamma_N^r(r)$ by def. of γ_N^r

	• $v_1 = c, v_2 = \bot$.
by def. of prop	$v_3 = c, v_4 = c$
by $scopecons_{M_2}(N)$ and Lemma 115	$\langle M_2, v_4 \rangle \in \gamma_N^{r}(r_1)$
by def. of γ_N^{r}	$r_1 \in \mathcal{R}_N^r \cup \{\top^r\}$
by def. of prop ^r	$r_3 = r_1, r_1 \sqsubseteq^{r} r_4$

• $v_1 = T, v_2 = \bot, r_2 \notin \mathcal{T}_N^r$. $v_3 = v_1, v_4 = T_\bot$ by def. of prop $r_1 \in \mathcal{T}_N^r \cup \{\top^r\}$ by def. of γ_N^r $r_3 = r_1 \lor r_3 = \top^r, r_4 = T_{\bot^r}^r \lor r_4 = \top^r$ by def. of prop^r $\langle M_2, T_\bot \rangle \in \gamma_N^r(r_4)$ by def. of γ_N^r

	$v_1 = T, v_2 = \bot, r_2 \in \mathcal{T}_N^r.$
by def. of prop	$v_3 = v_1, v_4 = T_\perp$
by def. of γ_N^r	$r_1 \in \mathcal{T}_N^{r} \cup \{\top^{r}\}$
by def. of $prop^r$ and Lemma 6	$r_3=r_1\vee r_3=\top^{r},r_2\sqsubseteq^{r}r_4$
by def. of γ_N^r	$\gamma_N^{\mathbf{r}}(T_{\perp^{\mathbf{r}}}^{\mathbf{r}}) \subseteq \gamma_N^{\mathbf{r}}(r_4)$
by def. of γ_N^r	$\langle M_2, T_\perp\rangle \in \gamma_N^{\rm r}(T_{\perp^{\rm r}}^{\rm r})$
	$\langle M_2, T_\perp \rangle \in \gamma_N^{r}(r_4)$

•
$$v_1 = c, v_2 = c$$
.
 $v_3 = c, v_4 = c$ by def. of prop
 $r_1 \sqsubseteq^r r_3, r_2 \sqsubseteq^r r_4$ by def. of prop^r and Lemma 6

For the induction case, let $v_1 = T_1$ and $v_2 = T_2$. By definition of γ_N^r , it must be that $r_1 \in \mathcal{T}_N^r \cup \{\top^r\}$ and $r_2 \in \mathcal{T}_N^r \cup \{\top^r\}$. If $r_1 = \top^r$ or $r_2 = \top^r$, the argument is trivial as $r_3 = r_4 = \top^r$ by the def. of prop^r. We thus focus on the case when $r_1 = T_1^r$ and $r_2 = T_2^r$. First, observe that $n_{cs} \in T_2 \implies n_{cs} \in T_2^r$ by the def. of γ_N^r . If $n_{cs} \in T_2^r$ but $n_{cs} \notin T_2$, the argument follows from Lemma 6.

We are left with the case where $n_{cs} \in T_2^r$ and $n_{cs} \in T_2$. Let $\langle v_{1i}, v_{1o} \rangle = T_1(n_{cs})$, $\langle v_{2i}, v_{2o} \rangle = T_1(n_{cs}), \ \langle r_{1i}, r_{1o} \rangle = T_1^r(n_{cs}), \text{ and } \langle r_{2i}, r_{2o} \rangle = T_2^r(n_{cs})$. By definition of $\langle :^c$, we have that $v_{2i} \langle :^c v_{1i}$ and $v_{1o} \langle :^c v_{2o}$.

Suppose first that $\langle v'_{2i}, v'_{1i} \rangle = \operatorname{prop}(v_{2i}, v_{1i})$ and $\langle r'_{2i}, r'_{1i} \rangle = \operatorname{prop}^{\mathsf{r}}(r_{2i}, r_{1i})$. We have $\operatorname{noerr}(v_{1i})$, $\operatorname{noerr}(v_{2i})$, $\operatorname{noerr}(v_{1o})$, $\operatorname{noerr}(v_{2o})$ by the definition of noerr. Also, it trivially follows that $\operatorname{scopecons}_{M_1}(N \setminus \{n_{\mathsf{cs}}\})$ and $\operatorname{scopecons}_{M_2}(N \setminus \{n_{\mathsf{cs}}\})$. By definition of γ^{r} , the fact that n_{cs} is not in the scope of $v_{2i}, v_{1i}, r_{2i}, r_{1i}$, and Lemma 97, we have $\langle M_1, v_{1i} \rangle \in \gamma^{\mathsf{r}}_{N \setminus \{n_{\mathsf{cs}}\}}(r_{1i})$ and $\langle M_2, v_{2i} \rangle \in \gamma^{\mathsf{r}}_{N \setminus \{n_{\mathsf{cs}}\}}(r_{2i})$. It then follows from induction hypothesis that $\langle M_1, v'_{1i} \rangle \in \gamma^{\mathsf{r}}_{N \setminus \{n_{\mathsf{cs}}\}}(r'_{1i})$ and $\langle M_2, v'_{2i} \rangle \in \gamma^{\mathsf{r}}_{N \setminus \{n_{\mathsf{cs}}\}}(r'_{2i})$.

Next, $\langle M_1[n_{cs} \mapsto v_{1i}], v_{1o} \rangle \in \gamma_{N \cup \{n_{cs}\}}^r(r_{1o})$ and $\langle M_2[n_{cs} \mapsto v_{2i}], v_{2o} \rangle \in \gamma_{N \cup \{n_{cs}\}}^r(r_{2o})$. Since $n_{cs} \in v_2$, we know $v_{2i} \neq \bot$. Also, we know by $\operatorname{noerr}(v_{2i})$ that $v_{2i} \neq \omega$. Hence, if $v_{2i} = c$, then by $v_{2i} <: {}^c v_{1i}$ we have that $v_{1i} = c$. Further, if $v_{2i} \in \mathcal{T}$, then $\langle M_1[n_{cs} \mapsto v_{2i}], v_{1o} \rangle$ by Lemma 114. It also holds that $\langle M_1[n_{cs} \mapsto v_{2i}], v_{1o} \rangle \in$ $\gamma_{N \cup \{n_{cs}\}}^r(r_{1o})$ and $\langle M_2[n_{cs} \mapsto v_{2i}], v_{2o} \rangle \in \gamma_{N \cup \{n_{cs}\}}^r(r_{2o})$. By definition of strengthening, one then also has $\langle M_1[n_{cs} \mapsto v_{2i}], v_{1o} \rangle \in \gamma_{N \cup \{n_{cs}\}}^r(r_{1o}[n_{cs} \leftarrow r_{2i}])$ and $\langle M_2[n_{cs} \mapsto v_{2i}], v_{2o} \rangle \in \gamma_{N \cup \{n_{cs}\}}^r(r_{2o}[n_{cs} \leftarrow r_{2i}])$.

Let $\langle v'_{1o}, v'_{2o} \rangle = \operatorname{prop}(v_{1o}, v_{2o})$ and $\langle r'_{1o}, r'_{2o} \rangle = \operatorname{prop}^{\mathsf{r}}(r_{1o}[n_{\mathsf{cs}} \leftarrow r_{2i}], r_{2o}[n_{\mathsf{cs}} \leftarrow r_{2i}])$. As $M_1[n_{\mathsf{cs}} \mapsto v_{2i}] \stackrel{:}{\sqsubseteq}^{\mathsf{r}} M_2[n_{\mathsf{cs}} \mapsto v_{2i}]$, $\operatorname{scopecons}_{M_1[n_{\mathsf{cs}} \mapsto v_{2i}]}(N \cup \{n_{\mathsf{cs}}\}))$, and $\operatorname{scopecons}_{M_2[n_{\mathsf{cs}} \mapsto v_{2i}]}(N \cup \{n_{\mathsf{cs}}\}))$, we have by induction hypothesis that $\langle M_1[n_{\mathsf{cs}} \mapsto v_{2i}], v'_{1o} \rangle \in \gamma_{N \cup \{n_{\mathsf{cs}}\}}^{\mathsf{r}}(r'_{1o})$ and $\langle M_2[n_{\mathsf{cs}} \mapsto v_{2i}], v'_{2o} \rangle \in \gamma_{N \cup \{n_{\mathsf{cs}}\}}^{\mathsf{r}}(r'_{2o})$.

By Lemma 38 and Lemma 39, we have $v'_{2i} <: v'_{1i}$ and, if $v_{2i} = c$, $v_{2i} = v'_{2i} = v_{1i} = v'_{1i} = c$. In that case, $\langle M_1[n_{cs} \mapsto v'_{1i}], v'_{1o} \rangle \in \gamma^r_{N \cup \{n_{cs}\}}(r'_{1o})$ and $\langle M_1[n_{cs} \mapsto v'_{2i}], v'_{2o} \rangle \in \gamma^r_{N \cup \{n_{cs}\}}(r'_{2o})$. In the case $v'_{2i} \in \mathcal{T}$, then also $v'_{1i} \in \mathcal{T}$ and the same result follows by Lemma 114.

A.4.3.2 Joins

We next show several results showing the soundness of the relational join subject to to the concrete one.

Lemma 117. Let v_1 , v_2 , be values, M a concrete map, $r_1 \in \mathcal{V}_N^r$, and $r_2 \in \mathcal{V}_N^r$. If $\operatorname{noerr}(v_1)$, $\operatorname{noerr}(v_2)$, $\operatorname{ccomp}(v_1, v_2)$, $\langle M, v_1 \rangle \in \gamma_N^r(r_1)$, $\langle M, v_2 \rangle \in \gamma_N^r(r_2)$, then $\langle M, v_1 \sqcup v_2 \rangle \in \gamma_N^r(r_1 \sqcup^r r_2)$.

Proof. By structural induction on both v_1 and v_2 . Let $v = v_1 \sqcup v_2$ and $r = r_1 \sqcup^r r_2$.

• $v_1 = \perp$ or $v_1 = \perp$. Trivial, as v = c and the argument then follows by monotonicity of γ_N^r .

• $v_1 = c, v_2 = c$. Again, by monotonicity of γ_N^r .

For the induction case, let $v_1 = T_1$ and $v_2 = T_2$. By definition of γ_N^r , it must be that $r_1 \in \mathcal{T}_N^r \cup \{\top^r\}$ and $r_2 \in \mathcal{T}_N^r \cup \{\top^r\}$. If $r_1 = \top^r$ or $r_2 = \top^r$, the argument is trivial as $r_3 = \top^r$ by the def. of \sqcup^r . We thus focus on the case when $r_1 = T_1^r$ and $r_2 = T_2^r$. Consider any call site n_{cs} . Let $\langle v_{1i}, v_{1o} \rangle = T_1(n_{cs}), \langle v_{2i}, v_{2o} \rangle = T_1(n_{cs}),$ $\langle r_{1i}, r_{1o} \rangle = T_1^r(n_{cs}),$ and $\langle r_{2i}, r_{2o} \rangle = T_2^r(n_{cs})$. Let us first consider the case for inputs. We have $\langle M, v_{1i} \rangle \in \gamma_N^r(r_{1i}), \langle M, v_{2i} \rangle \in \gamma_N^r(r_{2i}),$ noerr $(v_{2i}),$ noerr $(v_{1i}),$ and $\operatorname{ccomp}(v_{1i}, v_{2i})$. By i.h., we then have $\langle M, v_{1i} \sqcup v_{2i} \rangle \in \gamma^r(r_{1i} \sqcup^r r_{2i})$.

Next, we have that $\langle M[n_{cs} \mapsto v_{1i}], v_{1o} \rangle \in \gamma_N^r(r_{1o}), \langle M[n_{cs} \mapsto v_{2i}], v_{2o} \rangle \in \gamma_N^r(r_{2o}),$ **noerr** (v_{2o}) , **noerr** (v_{1o}) , and **ccomp** (v_{1o}, v_{2o}) . Now, if $v_{1i} = \bot$, then $v_{1o} = \bot$ by the def. of γ^r and **noerr** (v_{1o}) . Hence, $\langle M[n_{cs} \mapsto v_{2i}], v_{1o} \rangle \in \gamma_N^r(r_{1o})$ as well again by the definition of γ_N^r . The same holds in the other direction. For the remaining cases, if $v_{1i} = c$ then by $\text{ccomp}(v_{1i}, v_{2i})$ we have $v_{1i} = v_{2i}$, hence clearly $\langle M[n_{cs} \mapsto v_{2i}], v_{1o} \rangle \in \gamma_N^r(r_{1o})$. Lastly, if $v_{1i} \in \mathcal{T}$ then by $\text{ccomp}(v_{1i}, v_{2i})$ we have $v_{2i} \in \mathcal{T}$, and by Lemma 114 $\langle M[n_{cs} \mapsto v_{2i}], v_{1o} \rangle \in \gamma_N^r(r_{1o})$. By the induction hypothesis, we then have $\langle M[n_{cs} \mapsto v_{2i}], v_{2o} \sqcup v_{1o} \rangle \in \gamma_N^r(r_{1o} \sqcup^r r_{2o})$. Again, if $v_{2i} \in \mathcal{T}$ then by $\text{ccomp}(v_{1i}, v_{2i})$ and Lemma 114 $\langle M[n_{cs} \mapsto v_{2i} \sqcup v_{1i}], v_{2o} \sqcup v_{1o} \rangle \in \gamma_N^r(r_{1o} \sqcup^r r_{2o})$. Similar reasoning applies for the cases when $v_{2i} \in \{\bot\} \cup Cons$.

Lemma 118. Let $K \geq 2$, v_1 to v_K be values where $\operatorname{ccomp}(\{v_1, \ldots, v_K\})$ and $\forall 1 \leq i \leq K$. $\operatorname{noerr}(v_i)$, r_1 to r_K relational values in \mathcal{V}_N^r , and $\forall 1 \leq i \leq K$. $\langle M, v_i \rangle \in \gamma_N^r(r_i)$. Then, $\langle M, \sqcup_{1 \leq i \leq K} v_i \rangle \in \gamma_N^r(\sqcup_{1 \leq i \leq K}^r r_i)$.

Proof. The argument goes by induction on K. The result follows from the Lemma 117 and the fact that join \sqcup does not invalidate the constant compatibility of values. \Box

Lemma 119. Let $K \geq 2$, $\hat{M} = \{M_1, \ldots, M_K\}$ be a set of concrete maps where

 $\operatorname{ccomp}(\hat{M}) \ and \ \forall M \in \hat{M}.\operatorname{errcons}(M), \ \hat{M}_r = \{M_1^{\mathsf{r}}, \dots, M_K^{\mathsf{r}}\} \ a \ set \ of \ relational maps, \ and \ \forall 1 \leq i \leq K. \ M_i \in \dot{\gamma}^{\mathsf{r}}(M_i^{\mathsf{r}}). \ Then, \ \dot{\sqcup}\hat{M} \in \dot{\gamma}^{\mathsf{r}}(\dot{\sqcup}^{\mathsf{r}} \ \hat{M}_r).$

Proof. First, let

$$M_u = \dot{\sqcup}\hat{M} = \Lambda n. \sqcup \{M(n) \mid M \in \hat{M}\}$$

and

$$M_u^{\mathsf{r}} = \dot{\sqcup}^{\mathsf{r}} \, \hat{M} = \Lambda n. \, \sqcup^{\mathsf{r}} \{ M^{\mathsf{r}}(n) \mid M^{\mathsf{r}} \in \hat{M}_r \}$$

First, suppose there exists a map $M_i \in \hat{M}$ s.t. exists a node n where it is not the case that $\operatorname{noerr}(M_i(n))$. By $\operatorname{errcons}(M_i)$, it must be that $M_i = M_{\omega}$, in which case $M_u = M_{\omega}$. Also, it then must be that $M_i^r = M_{\top r}^r$ and clearly $M_u^r = M_{\top r}^r$. The result then trivially follows by the definition of $\dot{\gamma}^r$.

Otherwise, we have the following. As a reminder, by the definition of $\dot{\gamma}^{\mathsf{r}}$ one has that $M \in \dot{\gamma}^{\mathsf{r}}(M^{\mathsf{r}})$ implies that $\forall n. \langle M, M(n) \rangle \in \gamma^{\mathsf{r}}(M^{\mathsf{r}}(n))$. Also, by Lemma 42 and $M_u \neq M_\omega$, we have $\operatorname{errcons}(M_u)$. Hence,

$$\forall 1 \leq i \leq K. \, \forall n. \langle M_i, M_i(n) \rangle \in \gamma^{\mathsf{r}}(M_i^{\mathsf{r}}(n))$$
 by def. of $\dot{\gamma}^{\mathsf{r}}$

$$\forall 1 \le i \le K. \, \forall n. \langle M_u, M_i(n) \rangle \in \gamma^{\mathsf{r}}(M_i^{\mathsf{r}}(n))$$
 by Lemma 115

$$\forall n. \langle M_u, \sqcup \{ M_i(n) \mid M_i \in \hat{M} \} \rangle \in \gamma^{\mathsf{r}} (\sqcup^{\mathsf{r}} \{ M_i^{\mathsf{r}}(n) \mid M_i^{\mathsf{r}} \in \hat{M}_r \})$$
 by Lemma 118

$$\forall n. \langle M_u, M_u(n) \rangle \in \gamma^{\mathsf{r}} (M_u^{\mathsf{r}}(n))$$
 by def. of $\dot{\sqcup}^{\mathsf{r}}$ and $\dot{\sqcup}$

$$M_u \in \dot{\gamma}^{\mathsf{r}} (M_u^{\mathsf{r}})$$
 by def. of $\dot{\gamma}^{\mathsf{r}}$

by def. of
$$\gamma$$

A.4.3.3 Constants

We now prove the soundness of relational abstraction for constants.

Lemma 120. Let $M \in \mathcal{M}$, N a scope, and $c \in Cons$. If scopecons_M(N), then

 $\langle M, c \rangle \in \gamma_N^{\mathsf{r}}(c^{\mathsf{r}}).$

Proof. By definition, we have $c_N^r \{ D^r \in \mathcal{D}_N^r \mid D^r(\nu) = c \}$. By γ_N^r , we also have $\gamma_N^r(c^r) = \gamma_N^r(\{ D^r \in \mathcal{D}_N^r \mid D^r(\nu) = c \}) =$ $\{ \langle M, c \rangle \mid D^r \in \mathcal{D}_N^r \land D^r(\nu) = c \land \forall n \in N. \ M(n) \in \gamma^d(D^r(n)) \} \cup \gamma^r(\bot^r) =$ $\{ \langle M, c \rangle \mid D^r(\nu) = c \land \forall n \in N. \ M(n) \notin \{ \bot, \omega \} \} \cup \gamma^r(\bot^r)$

The result then follows from the definition of scopecons.

A.4.3.4 Strengthening

We now prove the soundness of several strengthening operations.

Lemma 121. Let $M \in \mathcal{M}$, $v_1 \in \mathcal{V}$, $n \in \mathcal{N}$, $r \in \mathcal{V}_N^r$, and $r' \in \mathcal{V}^r$. If $\langle M, v_1 \rangle \in \gamma_N^r(r)$ and $\langle M, M(n) \rangle \in \gamma^r(r')$, then $\langle M, v_1 \rangle \in \gamma_N^r(r[n \leftarrow r'])$.

Proof.

$$\gamma_{N}^{r}(r_{N}[n \leftarrow r']) = \gamma_{N}^{r}(\alpha_{N}^{r}(\gamma^{r}(r_{N}) \cap \{\langle M'[n \mapsto v'], v \rangle \mid \langle M', v' \rangle \in \gamma^{r}(r')\})) \supseteq$$

$$\gamma^{r}(r_{N}) \cap \{\langle M'[n \mapsto v'], v \rangle \mid \langle M', v' \rangle \in \gamma^{r}(r')\} \supseteq$$

$$\{\langle M, v_{1} \rangle\} \cap \{\langle M'[n \mapsto v'], v \rangle \mid \langle M, v' \rangle \in \{\langle M, M(n) \rangle\}\} =$$

$$\{\langle M, v_{1} \rangle\} \cap \{\langle M[n \mapsto M(n)], v \rangle \mid v \in \mathcal{V}\} =$$

$$\{\langle M, v_{1} \rangle\} \cap \{\langle M, v_{1} \rangle\} \cap \{\langle M, v_{1} \rangle \mid v \in \mathcal{V}\} =$$

$$\{\langle M, v_{1} \rangle\}$$

We can generalize the above result.

Lemma 122. Let $M \in \mathcal{M}$, $v \in \mathcal{V}$, N a scope, $M^{\mathsf{r}} \in \mathcal{M}^{\mathsf{r}}$, $r \in \mathcal{V}_{N}^{\mathsf{r}}$. If $\langle M, v \rangle \in \gamma_{N}^{\mathsf{r}}(r)$ and $M \in \dot{\gamma}^{\mathsf{r}}(M^{\mathsf{r}})$, then $\langle M, v \rangle \in r[M^{\mathsf{r}}]$.

Proof. Consider any $n \in N$. By $M \in \dot{\gamma}^{\mathsf{r}}(M^{\mathsf{r}})$, we have that $\langle M, M(n) \rangle \in \gamma^{\mathsf{r}}(M^{\mathsf{r}}(n))$. Then by Lemma 121, $\langle M, v \rangle \in \gamma^{\mathsf{r}}_{N}(r[n \leftarrow M^{\mathsf{r}}(n)])$. Hence, $\langle M, v \rangle \in \bigcap_{n \in N} \gamma^{\mathsf{r}}_{N}(r[n \leftarrow M^{\mathsf{r}}(n)])$. Then by Lemma 9 we have $\langle M, v \rangle \in \gamma^{\mathsf{r}}_{N}(\sqcap_{n \in N}^{\mathsf{r}}r[n \leftarrow M^{\mathsf{r}}(n)])$. \Box

Next, we present a results showing that updates to a relational map with sound information yields a sound map.

Lemma 123. Let $M \in \mathcal{M}$, $M^{\mathsf{r}} \in \mathcal{M}^{\mathsf{r}}$, $n \in \mathcal{N}$, $v \in \mathcal{V}$, and $r \in \mathcal{V}^{\mathsf{r}}$. If $\langle M, v \rangle \in \gamma^{\mathsf{r}}(r)$ and $M \in \dot{\gamma}^{\mathsf{r}}(M^{\mathsf{r}})$, then $M \in \dot{\gamma}^{\mathsf{r}}(M^{\mathsf{r}}[n \mapsto r])$.

Proof. Clearly, as
$$\forall n'. \langle M, M(n') \rangle \in \gamma^{\mathsf{r}}(M^{\mathsf{r}}[n \mapsto r](n')).$$

Our rescoping operations are also sound.

Lemma 124. Let $M \in \mathcal{M}$, $v_1 \in \mathcal{V}$, $r \in \mathcal{V}_N^r$, and $N' \subseteq \wp(\mathcal{N})$ such that scopecons_M(N'). If $\langle M, v_1 \rangle \in \gamma_N^r(r)$, then $\langle M, v_1 \rangle \in \gamma_{N \cup N'}^r(r[N \cup N'])$

Proof.

$$r_{N}[N \cup N'] = \alpha_{N \cup N'}^{\mathsf{r}}(\gamma^{\mathsf{r}}(r_{N}) \cap \{\langle M, v \rangle \mid \forall n \in N' \setminus N.M(n) \notin \{\bot, \omega\}\})$$
$$\exists^{\mathsf{r}} \alpha_{N \cup N'}^{\mathsf{r}}(\{\langle M, v_{1} \rangle)\} \cap \{\langle M, v \rangle \mid \forall n \in N' \setminus N.M(n) \notin \{\bot, \omega\}\})$$
$$= \alpha_{N \cup N'}^{\mathsf{r}}(\{\langle M, v_{1} \rangle)\})$$

Lemma 125. Let $M \in \mathcal{M}$, $v_1 \in \mathcal{V}$, $r \in \mathcal{V}_N^r$, and $N' \subseteq N$. If $\langle M, v_1 \rangle \in \gamma_N^r(r)$, then $\langle M, v_1 \rangle \in \gamma_{N'}^r(r[N'])$

Proof. Straightforward as $N' \setminus N = \emptyset$.

Strengthening of a relational value with the information not in the scope is sound as well.

A.4.3.5 Maps

We next show the soundness of updates to the relational maps subject to their concretizations.

Lemma 126. Let $M \in \mathcal{M}$, $M^{\mathsf{r}} \in \mathcal{M}^{\mathsf{r}}$, $n \in \mathcal{N}$, $v \in \mathcal{V}$, and $r \in \mathcal{V}^{\mathsf{r}}$. If $\langle M, v \rangle \in \gamma^{\mathsf{r}}(r)$, $M \in \dot{\gamma}^{\mathsf{r}}(M^{\mathsf{r}})$, $M(n) \sqsubseteq v$, and $v \neq \omega$, then $M[n \mapsto v] \in \dot{\gamma}^{\mathsf{r}}(M^{\mathsf{r}}[n \mapsto r])$.

Proof. Suppose not. Let $M_0 = M[n \mapsto v]$ and $M_0^r = M^r[n \mapsto r]$. There must be a node n' such that $\langle M_0, M_0(n') \rangle \notin \gamma^r(M_0^r(n'))$. We branch on several cases:

- $n \notin N_{n'}$. If n = n', the argument is trivial. Otherwise we know that $\langle M, M(n') \rangle \in \gamma^{\mathsf{r}}(M^{\mathsf{r}}(n')), M^{\mathsf{r}}(n') = M_0^{\mathsf{r}}(n')$, and $M(n') = M_0(n')$. Then by Lemma 97 also $\langle M_0, M_0(n') \rangle \in \gamma^{\mathsf{r}}(M_0^{\mathsf{r}}(n'))$.
- $n \in N_{n'}$. Again, we know that $\langle M, M(n') \rangle \in \gamma^{\mathsf{r}}(M^{\mathsf{r}}(n')), M^{\mathsf{r}}(n') = M_0^{\mathsf{r}}(n'),$ and $M(n') = M_0(n')$. Since $M \doteq M', M =_{\{n\}} M'$, and $M'(n) \neq \omega$, then by Lemma 115 $\langle M_0, M_0(n') \rangle \in \gamma^{\mathsf{r}}(M^{\mathsf{r}}(n'))$ and hence $\langle M_0, M_0(n') \rangle \in \gamma^{\mathsf{r}}(M_0^{\mathsf{r}}(n'))$.

		I	
		I	
_		4	

A.4.3.6 Transformer

We next show that the relational transformer abstracts the concrete transformer given certain invariants on the concrete maps.

Lemma 127. First, assume that P is an inductive invariant family of a program m where $\forall k, \ell, E.(\ell, M, E) \in P_k \Longrightarrow df_m(M) \land \operatorname{envcons}_M(E)$. Now, let e^{ℓ} be an expression of $m, M_f \in \mathcal{M}, M \in \mathcal{M}, k \in \mathbb{N}$, and E an environment where (e, E) is well-formed, domvalid (ℓ, m, E) , and $(\ell, M, E) \in P_k$. Also, let $M_f^r \in \mathcal{M}^r$ and $M^r \in$ \mathcal{M}^{r} such that $M \in \dot{\gamma}^{\mathsf{r}}(M^{\mathsf{r}})$. If $\operatorname{step}_{k}[\![e^{\ell}]\!](E)(M) = M_{f}$ and $\operatorname{Step}_{k}^{\mathsf{r}}[\![e^{\ell}]\!](E)(M^{\mathsf{r}}) = M_{f}^{\mathsf{r}}$, then $M_{f} \in \dot{\gamma}^{\mathsf{r}}(M_{f}^{\mathsf{r}})$.

Proof. We carry the proof by induction on k and e. The case when k = 0 is trivial. We note that $\operatorname{Step}_{k}^{\mathsf{r}}[\![e^{\ell}]\!](E)(M^{\mathsf{r}}) = M^{\mathsf{r}} \dot{\sqcup}^{\mathsf{r}} \operatorname{step}_{k}^{\mathsf{r}}[\![e^{\ell}]\!](E)(M^{\mathsf{r}})$, hence we focus on $\operatorname{step}_{k}^{\mathsf{r}}[\![e^{\ell}]\!]$ and consider all the cases for e. Also observe that df(M) and $\operatorname{envcons}_{M}(E)$ by P and $(\ell, M, E) \in P_{k}$.

 $e = c^{\ell}$. Let $v = M(E \diamond \ell)$ and $r = M^{\mathsf{r}}(E \diamond \ell)$. By $df_m(M)$, we have that $\mathsf{constcons}_m(M)$ which implies $v \in \{\perp, \omega, c\}$. If $v = \omega$ then by $df_n(M)$ we have $\mathsf{errcons}(M)$ and thus $M = M_{\omega}$. Then it must be that $M^{\mathsf{r}} = M^{\mathsf{r}}_{\top \mathsf{r}}$ and the argument is trivial.

Otherwise, $v \in \{\perp, c\}$. By $envcons_M(E)$ we have $scopecons_M(N)$ where $N = N_{E\diamond\ell}$. Then by Lemma 120 it follows $\langle M, c \rangle \in \gamma_N^{\mathsf{r}}(c^{\mathsf{r}})$ and by Lemma 122 it must be that $\langle M, c \rangle \in \gamma_N^{\mathsf{r}}(c^{\mathsf{r}}[M^{\mathsf{r}}])$. Since noerr(c), noerr(v), ccomp(v, c), and $\langle M, v \rangle \in \gamma^{\mathsf{r}}(r)$, by Lemma 117 it must be that $\langle M, v \sqcup c \rangle \in \gamma_N^{\mathsf{r}}(r \sqcup^{\mathsf{r}} c^{\mathsf{r}}[M^{\mathsf{r}}])$. The result then follows by Lemma 126.

 $e = x^{\ell}$. Let $M(E \diamond \ell) = v$, $M(E(x)) = v_x$, $M^{\mathsf{r}}(E \diamond \ell) = r$, and $M^{\mathsf{r}}(E(x)) = r_x$. If $v_x = \omega$ or $v = \omega$, then by $\operatorname{errcons}(M)$ it must be that $M = M_{\omega}$ and hence $M^{\mathsf{r}} = M^{\mathsf{r}}_{\mathsf{T}^{\mathsf{r}}}$. In this case, the argument is trivial.

Otherwise, we know that $\langle M, v \rangle \in \gamma^{\mathsf{r}}(r)$, and $\langle M, v_x \rangle \in \gamma^{\mathsf{r}}(r_x)$. Let r_0 be r_x implicitly rescoped with $N_{E(x)} \cup N_{E \diamond \ell}$. Let also r_1 be r be rescoped with $N_{E(x)} \cup N_{E \diamond \ell}$. By $\mathsf{envcons}_M(E)$ and $\mathsf{envcons}(M)$ and $\mathsf{Lemma 124}$, we have $\langle M, v_x \rangle \in \gamma^{\mathsf{r}}(r_0)$. $\langle M, v \rangle \in \gamma^{\mathsf{r}}(r_1)$. As E(x) was not initially in the scope of E(x), then also $\langle M, v_x \rangle \in \gamma^{\mathsf{r}}(r_0[E(x)=\nu])$.

Next, by $df_m(M)$ we have that $v_x <: {}^cv$. As by $\operatorname{errcons}(M)$ we have $\operatorname{noerr}(v)$, $\operatorname{noerr}(v_x)$, and $\operatorname{scopecons}_M(N)$ (via $\operatorname{envcons}_M(E)$), then by Lemma 116 we have that $\langle M, v'_x \rangle \in \gamma^{\mathsf{r}}(r'_x)$ and $\langle M, v' \rangle \in \gamma^{\mathsf{r}}(r')$. As r'_x is rescoped back to $N_{E(x)}$ as r_2 and r' is rescoped back to $N_{E \diamond \ell}$ as r_3 , by Lemma 125 we have $\langle M, v'_x \rangle \in \gamma^{\mathsf{r}}(r_2)$ and $\langle M, v' \rangle \in \gamma^{\mathsf{r}}(r_3)$. By Lemma 122, we have $\langle M, v' \rangle \in \gamma^{\mathsf{r}}(r_3[M^{\mathsf{r}}])$.

Let $M_1 = M[E(x) \mapsto v'_x]$, $M_2 = M_1[E \diamond \ell \mapsto v']$, $M_1^r = M^r[E(x) \mapsto r_2]$, and $M_2^r = M_1^r[E \diamond \ell \mapsto r_3]$. By Lemma 126, $M_1 \in \gamma_N^r(M_1^r)$. By the increasing property of prop (Lemma 6), we have $M \doteq M_1$. Also, $M =_{\backslash \{E(x)\}} M_1$ and $v'_x \neq \omega$ by $v_x <: {}^c v$ and Lemma 39. Hence, $\operatorname{envcons}_{M_1}(E)$ and $\operatorname{scopecons}_{M_1}(N)$. Then by Lemma 115 we have $\langle M_1, v' \rangle \in \gamma_N^r(r_3)$. Finally, we then have by Lemma 126 that $M_2 \in \dot{\gamma}^r(M_2^r)$.

 $e = (e_1^{\ell_1} e_2^{\ell_2})^{\ell}$. By i.h., $M_1 \in \dot{\gamma}^r(M_1^r)$. By the inductiveness of P, we have $df_m(M_1)$ and $\mathsf{envcons}_{M_1}(E)$. If $v_1 = \bot$, then $M_f = M_1$ and $M_1^r \doteq^r M_f^r$ by the obvious increasing property of $\mathsf{Step}_{k+1}^r[\![\cdot]\!]$ and prop^r (Lemma 6); the result thus follows by the monotonicity of $\dot{\gamma}^r$. If $v_1 = \omega$, then by $\mathsf{errcons}(M_1)$ it must be that $M_1 = M_\omega$ and hence $M^r = M_{\mathsf{T}^r}^r$, in which case the result follows immediately. If $v_1 \notin \mathcal{T}$, then by the definition of γ^r also $r_1 \notin \mathcal{T}_N^r$. In that case, $M_f = M_\omega$ and $M^r = M_{\mathsf{T}^r}^r$ and the argument is straightforward. If $v_1 \in \mathcal{T}$, but $r_1 \notin \mathcal{T}_N^r$, then $r_1 = \mathsf{T}^r$ and hence $M_f^r = M_{\mathsf{T}^r}^r$. Otherwise, by i.h. $M_2 \in \dot{\gamma}^r(M_2^r)$. Again by the inductive property of P, we have $df_m(M_2)$ and $\mathsf{envcons}_{M_2}(E)$. The cases when $v_2 \in \{\bot, \omega\}$ are handled identically as for v_1 .

By domvalid(ℓ, m, E), domvalid(ℓ_1, m, E), and domvalid(ℓ_2, m, E), Lemma 32, and Lemma 33, we have that $v_1 = M_2(E \diamond \ell_1)$. Similarly, by Lemma 112 it also holds $r_1 = M_2^r(E \diamond \ell_1)$. Let $v = M_2(E \diamond \ell)$ and $r = M_2^r(E \diamond \ell)$. Then by $M_2 \in \dot{\gamma}^r(M_2^r)$, we have $\langle M_2, v_1 \rangle \in \gamma^r(r_1)$, $\langle M_2, v_2 \rangle \in \gamma^r(r_2)$, and $\langle M_2, v \rangle \in \gamma^r(r)$. As r is implicitly rescoped with $N_{E \diamond \ell} \cup \{E \diamond \ell_1\}$ to r_0 , then by Lemma 124, $\langle M_2, v_2 \rangle \in \gamma^r(r_0)$.

If any of v_1, v_2 , or v_3 is ω , then $M_2 = M_{\omega}$ and the argument goes as usual. Oth-

erwise, we have $\operatorname{noerr}(v_1)$, $\operatorname{noerr}(v_2)$, $\operatorname{noerr}(v_3)$ and by $df_m(M_2)$ also $v_1 <: {}^c[E \diamond \ell_1 \mapsto \langle v_2, v_2 \rangle]$. Since $\langle M_2, [E \diamond \ell_1 \mapsto \langle v_2, v_2 \rangle] \rangle \in \gamma_N^r([E \diamond \ell_1 \mapsto \langle r_2, r_0 \rangle])$ by def. of γ^r and $\operatorname{envcons}_{M_2}(E)$, we have by Lemma 116 that $\langle M_2, v_1' \rangle \in \gamma_N^r(r_1')$ and $\langle M_2, [E \diamond \ell_1 \mapsto \langle v_2', v_1' \rangle] \rangle \in \gamma_N^r([E \diamond \ell_1 \mapsto \langle r_2', r_0' \rangle])$ and hence by def. of γ^r we have $\langle M_2, v_1' \rangle \in \gamma^r(r_1')$, $\langle M_2, v_2' \rangle \in \gamma^r(r_2')$, and $\langle M_2, v_1' \rangle \in \gamma^r(r_0')$. As r_0' is implicitly rescoped back with $N_{E \diamond \ell}$ to r_0'' , by Lemma 125 we have $\langle M_2, v_1' \rangle \in \gamma^r(r_0')$. The argument for storing the results to maps M_2 and M_2' then follows by Lemma 39, 115, and 126 as earlier.

 $(e_0^{\ell_0} ? e_1^{\ell_1} : e_2^{\ell_2})^{\ell}$. By i.h., we have that $M_0 \in \gamma_N^r(M_0^r)$. By the inductiveness of P, we have $df_m(M_0)$ and $envcons_{M_0}(E)$. The cases when $v_0 \in \{\perp, \omega\}$ are handled as expected. If $v_0 \in Bool$ and $r_0 \notin Bool^r$, then $r_0 = \top^r$ by the definition of γ^r and $M_f^r = M_{\top^r}^r$.

Otherwise, the argument is as follows. Suppose $v_0 = true$; the argument for the case when $v_0 = false$ is identical. Since $envcons_{M_0}(E)$ implies $scopecons_{M_0}(N_{E \diamond \ell})$, by Lemma 120 we have that $\langle M_0, true \rangle \in \gamma^r(true^r)$. Since $\langle M_0, true \rangle \in \gamma^r(r_0)$ by $M_0 \in \gamma_N^r(M_0^r)$, we have that $\langle M_0, true \rangle \in (\gamma^r(true^r) \cap \gamma^r(r_0))$. Then by Lemma 9, it also holds $\langle M_0, true \rangle \in \gamma^r(true^r \sqcap^r r_0)$. Now, let $n' \in N_{E \diamond \ell_0}$. By $M_0 \in \dot{\gamma}^r(M_0^r)$, we have that $\langle M_0, M_0(n') \rangle \in \gamma^r(M_0^r(n'))$. As $\langle M_0, true \rangle = \langle M_0, M_0(E \diamond \ell_0) \rangle$, we have that $\langle M_0, M_0(n') \rangle \in \gamma^r(M_0^r(n')[E \diamond \ell_0 \leftarrow true^r \sqcap^r r_0])$ by Lemma 121. Hence, by Lemma 123 it is the case that $M_0 \in \dot{\gamma}^r(M_0^r[n' \mapsto M_0^r(n')[E \diamond \ell_0 \leftarrow true^r \sqcap^r r_0])$. Further, by Lemma 10 it is also the case that $M_0 \in \dot{\gamma}^r(M_0^r[E \diamond \ell_0 \leftarrow true^r \sqcap^r r_0])$.

It follows then by i.h. that $M_1 \in \gamma_N^r(M_1^r)$ and by the inductiveness of P, we have $df_m(M_1)$ and $envcons_{M_1}(E)$. The argument continues as usual by Lemma 39, 115, and 126 as before.

 $e = (\mu f.\lambda x.e_1^{\ell_1})^{\ell}$. Let $v = M(E \diamond \ell)$. By $df_m(M)$ it follows that $\operatorname{constcons}_m(M)$ which implies that $v \in \{\bot, \omega\} \cup \mathcal{T}$. If $v = \bot$, then $T = T_{\bot}$ and $M_f = M[E \diamond \ell \mapsto T_{\bot}]$. By the definition of γ_N^r , then either (1) $r = \bot^r$ in which case $M_f^r = M^r[E \diamond \ell \mapsto T_{\bot^r}]$, (2) $r \in \mathcal{R}_N^r$ in which case $M_f^r = M_{\top^r}^r$, (3) $r \in \mathcal{T}^r$ and the result follows by the obvious increasing property of $\operatorname{Step}_{k+1}^r[\cdot]$ and prop^r (Lemma 6), or (4) $r = \top^r$ in which case $M_f^r = M_{\top^r}^r$. If $v = \omega$ then by $\operatorname{ercons}(M)$ it must be that $M = M_f = M_{\omega}$ and hence $M^r = M_f^r = M_{\top^r}^r$. If $v \in \mathcal{T}$ but $r \notin \mathcal{T}^r$, then $r = \top^r$ and hence $M_f^r = M_{\top^r}^r$.

We are left with the case where $T \in \mathcal{T}$ and $T^{\mathsf{r}} \in \mathcal{T}^{\mathsf{r}}$ where v = T and $r = T^{\mathsf{r}}$. By γ^{r} , $n_{\mathsf{cs}} \in T$ implies $n_{\mathsf{cs}} \in T^{\mathsf{r}}$. For the case where $n_{\mathsf{cs}} \notin T$ and $n_{\mathsf{cs}} \notin T^{\mathsf{r}}$, the argument is trivial. When $n_{\mathsf{cs}} \notin T$ but $n_{\mathsf{cs}} \in T^{\mathsf{r}}$, the result follows by the obvious increasing property of $\mathsf{Step}_{k+1}^{\mathsf{r}} \llbracket \cdot \rrbracket$ and $\mathsf{prop}^{\mathsf{r}}$. We are thus left with the case when $n_{\mathsf{cs}} \in T$ and $n_{\mathsf{cs}} \in T^{\mathsf{r}}$. The case when $\pi_1(T(n_{\mathsf{cs}})) \in \{\bot, \omega\}$ are handled as expected.

We thus continue assuming $\pi_1(T(n_{cs})) \notin \{\perp, \omega\}$. First, we have by $df_m(M)$ that $[n_{cs} \mapsto \langle M(n_x), M(E_1 \diamond \ell_1) \rangle] <:^{c}T$ and $T <:^{c}M(n_f)$. By the def. of $<:^{c}, M(n_x) \notin \{\perp, \omega\}$. Next, $\langle M, M(E_1 \diamond \ell_1) \rangle \in \gamma^r(M^r(E_1 \diamond \ell_1))$ by $M \in \gamma_N^r(M^r)$. Since $M^r(E_1 \diamond \ell_1)$ is first implicitly rescoped with $N_{E \diamond \ell_1} \cup \{n_{cs}\}$ to r_0 and $M(n_x) \notin \{\perp, \omega\}$, by Lemma 124 we have $\langle M, M(E_1 \diamond \ell_1) \rangle \in \gamma^r(r_0)$. As $n_{cs} \notin N_{E \diamond \ell_1}$, then it must be that $\langle M, M(E_1 \diamond \ell_1) \rangle \in \gamma^r(r_0[n_{cs}=n_x])$. Then, r_0 is implicitly rescoped to the same scope but without n_x and then by Lemma 125 we have $\langle M, M(E_1 \diamond \ell_1) \rangle \in \gamma^r(r_1)$. We hence have $[n_{cs} \mapsto \langle M(n_x), M(E_1 \diamond \ell_1) \rangle] \in \gamma^r[n_{cs} \mapsto \langle M^r(n_x), r_1 \rangle, T \in \gamma^r(T^r),$ and $M(n_f) \in \gamma_N^r(M^r(n_f))$. As $\operatorname{scopecons}_M(N_{E \diamond \ell})$ by $\operatorname{envcons}_M(E)$, $\operatorname{noerr}(T)$, $\operatorname{noerr}(M(n_x))$, $\operatorname{noerr}(M(n_f))$, and $\operatorname{noerr}(M(E_1 \diamond \ell_1))$, by Lemma 116 we have that $\langle M, T' \rangle \in \gamma_N^r(T_1^r), \langle M, T'' \rangle \in \gamma_N^r(T_2^r), \langle M, T_x' \rangle \in \gamma_N^r(T_x^r),$ and $\langle M, T_f' \rangle \in \gamma^r(T_f^r)$. By Lemma 39, we have $\operatorname{noerr}(T'')$ and $\operatorname{noerr}(T')$. $\mathsf{ccomp}(T',T'')$ and thus $\langle M,T'\sqcup T''\rangle \in \gamma^{\mathsf{r}}(T_1^{\mathsf{r}}\sqcup^{\mathsf{r}}T_2^{\mathsf{r}})$ by Lemma 117.

Note that by Lemma 40 it must be that $\operatorname{noerr}(T'_x)$ and hence $v'_x \neq \omega$. By this and $\operatorname{envcons}_M(E)$ it must be that $\operatorname{envcons}_{M_1}(E_1)$. It follows then by i.h. that $M_1 \in \gamma_N^r(M_1^r)$ and by inductiveness of P we have $df_m(M_1)$ and $\operatorname{envcons}_{M_1}(E_1)$. The argument continues as usual by Lemma 39, 115, and 126 as before. Also, the soundness argument for implicit rescoping and strengthening of r'_1 goes as usual and relies on the fact that n_x is not in the scope of r'_1 .

Remember that for every $n_{cs} \in \operatorname{dom}(\bar{M})$ we have $n_{cs} \in \operatorname{dom}(\bar{M}^r)$ and $\bar{M}(n_{cs}) \in \dot{\gamma}^r((\bar{M}^r)(n_{cs}))$. Given fin(M) implied by $df_m(M)$, we have $|\operatorname{dom}(\bar{M})| = K$ where $K \in \mathbb{N}$. By the inductiveness of P, we have that $df_m(M_f)$ and hence $\operatorname{errcons}(M_f)$. If $M_f = M_\omega$ then there exists n_{cs} such that $\bar{M}(n_{cs}) = M_\omega$ in which case $\bar{M}^r(n_{cs}) = M_{\mathsf{T}^r}$ and hence $M_f^r = M_{\mathsf{T}^r}^r$. Otherwise, by the definition of $\dot{\sqcup}$ we have that $\forall M \in \overline{M}$. $M \neq M_\omega$. By Lemma 44 we have $\operatorname{ccomp}(\overline{M})$ and $\forall M \in \overline{M}$. $\operatorname{errcons}(M)$. The result then follows from Lemma 118.

We can now prove that relational semantics abstract the concrete one.

Proof (of Theorem 1).

Proof. We assume a fixed program e^{ℓ_m} . We first define a set $P_k \in \wp(Loc \times \mathcal{M} \times \mathcal{E})$ for every fuel k as the smallest set satisfying the following. Initially, $(\ell_m, M_{\perp}, \emptyset) \in$ P_k . Second, $(\ell, M, E) \in P_k$ if (1) $\ell \in Loc(e)$, (2) $(e(\ell), E)$ is well-formed, (3) domvalid (ℓ, e, E) , (4) there exists a path map \mathcal{M}^w such that $\operatorname{cons}_m(\mathcal{M}^w)$, $\mathcal{M}^w \in$ dom(strip), $\mathcal{M} = \operatorname{strip}(\mathcal{M}^w)$, and $\operatorname{envcons}_{\mathcal{M}^w}(E)$. Lastly, for every $(\ell, M, E) \in P_k$, if $\operatorname{step}_k^w[\![e(\ell)]\!](E)(\mathcal{M}^w) = \mathcal{M}_1^w$, $\operatorname{cons}_m(\mathcal{M}^w)$, $\operatorname{cons}_m(\mathcal{M}_1^w)$, $\mathcal{M}^w \in \operatorname{dom}(\operatorname{strip})$, $\mathcal{M}_1 = \operatorname{strip}(\mathcal{M}^w)$, and $\mathcal{M}_1 = \operatorname{strip}(\mathcal{M}_1^w)$, then $(\ell, \mathcal{M}_1, E) \in P_k$.

By Theorem 7, Theorem 8, $\operatorname{cons}_m(M^{\mathsf{w}}_{\perp}), M^{\mathsf{w}}_{\perp} \in \operatorname{dom}(\operatorname{strip})$, and $M_{\perp} = \operatorname{strip}(M^{\mathsf{w}}_{\perp})$,

we have that P is an inductive invariant family of m. By Lemma 95 and the definition of envcons for path maps, we have that $(\ell, M, E) \in P_k$, for every k, implies $df_m(M)$ and $envcons_M(E)$. Since $(\ell_m, M_\perp, \emptyset) \in P_k$, we have that the strongest inductive invariant I of m is subsumed by P, i.e, $S_k \subseteq P_k$. The result then follows from the definition of $\mathbf{C}[\![\cdot]\!]$ and Lemma 127. \Box

A.5 Collapsed Semantics Proofs

We first prove that our definitions of collapsed meet and join are correct. A depth of a collapsed value is defined in the similar way as for the concrete and relational values. We hence reuse the function depth to denote the depth of collapsed values.

Proof (of Lemma 8).

Proof. Similar to the proof of Lemma 4, the proof is carried by case analysis on elements of U and induction on the minimum depth thereof when U consists of collapsed tables only.

The next proof shows that the concretization function for collapsed values is a complete meet-morphism.

Theorem 11 (Complete Meet Morphism of $\gamma_{\hat{N}}^{cr}$). Let $U \in \wp(\mathcal{V}_{\hat{N}}^{cr})$ and let N be a concrete scope. Then $\gamma_{\delta}^{cr}(\sqcap^{cr}U) = \sqcap_{u \in U}^{r} \gamma_{\delta}^{cr}(u)$ for $\delta \in \Delta(\hat{N}, N)$.

Proof. We carry the proof by case analysis on elements of U and induction on the minimum depth thereof when U consists of tables only.

We first consider the trivial (base) cases where U is not a set of multiple collapsed tables.

- $U = \emptyset$. Here, $\gamma_{\delta}^{\mathsf{cr}}(\sqcap^{\mathsf{cr}}\emptyset) = \gamma_{\delta}^{\mathsf{cr}}(\top^{\mathsf{cr}}) = \top^{\mathsf{cr}} = \sqcap^{\mathsf{r}}\emptyset$.
- $U = \{u\}$. Trivial.
- $R^{\mathsf{cr}} \in U, T^{\mathsf{cr}} \in U$. We have $\sqcap^{\mathsf{cr}} \{ R^{\mathsf{cr}}, T^{\mathsf{cr}} \} = \bot^{\mathsf{cr}}$ and since \bot^{cr} is the bottom element, $\sqcap^{\mathsf{cr}} U = \bot^{\mathsf{cr}}$. Similarly, $\gamma_{\delta}^{\mathsf{cr}}(R^{\mathsf{cr}}) \sqcap^{\mathsf{r}} \gamma_{\delta}^{\mathsf{cr}}(T^{\mathsf{cr}}) = \bot^{\mathsf{r}} = \gamma_{\hat{N}}^{\mathsf{cr}}(\bot^{\mathsf{cr}})$ so by the definition of \sqcap^{r} we have $\sqcap_{u \in U}^{\mathsf{r}} \gamma_{\delta}^{\mathsf{cr}}(u) = \bot^{\mathsf{r}}$.

• $U \subseteq \mathcal{R}_{\hat{N}}^{\mathsf{cr}}$. Here,

• $\top^{\mathsf{cr}} \in U$. Here, $\sqcap^{\mathsf{cr}} U = \sqcap^{\mathsf{cr}} (U/\{\top^{\mathsf{cr}}\})$ and $\sqcap^{\mathsf{r}}_{u \in U} \gamma^{\mathsf{cr}}_{\delta}(u) = \sqcap^{\mathsf{r}}_{u \in U, u \neq \top^{\mathsf{cr}}} \gamma^{\mathsf{cr}}_{\delta}(u)$ since $\gamma^{\mathsf{cr}}_{\hat{N}}(\top^{\mathsf{cr}}) = \top^{\mathsf{r}}$. The set $U/\{\top^{\mathsf{cr}}\}$ either falls into one of the above cases or consists of multiple tables, which is the case we show next.

Let $U \subseteq \mathcal{T}_{\hat{N}}^{\mathsf{cr}}$ and |U| > 1. Let d be the minimum depth of any table in U.

We can now state the monotonicity of functions constituting the Galois connection between collapsed and relational values. **Lemma 128.** For any abstract scope \hat{N} , concrete scope N, and $\delta \in \Delta(\hat{N}, N)$, both α_{δ}^{cr} and γ_{δ}^{cr} are monotone.

Proof. The argument follows from the fact that α_{δ}^{cr} and γ_{δ}^{cr} constitute a Galois connection (Prop.1), as evidenced by Theorem 11.

The meet morphism requirement is also satisfied by the concretization function defined over collapsed maps.

Theorem 12 (Complete Meet Morphism of $\dot{\gamma}^{cr}$). Let $m^{cr} \in \wp(\mathcal{M}^{cr})$. Then $\dot{\gamma}^{cr}(\dot{\sqcap}^{cr}m^{cr}) = \dot{\sqcap}^{r}{}_{M^{cr}\in m^{cr}}\dot{\gamma}^{cr}(M^{cr})$.

Proof.

$$\dot{\gamma}^{\mathsf{cr}}(\dot{\sqcap}^{\mathsf{cr}}m^{\mathsf{cr}}) = \Lambda n. \ \sqcap^{\mathsf{r}} \left\{ \gamma_{\delta}^{\mathsf{cr}}(\sqcap^{\mathsf{cr}}\{M^{\mathsf{cr}}(\rho(n)) \mid M^{\mathsf{cr}} \in m^{\mathsf{cr}}\} \mid \delta \in \Delta(N_{\rho(n)}, \hat{N}_n) \right\}$$

$$[\text{by def. of } \sqcap^{\mathsf{cr}} \text{ and } \dot{\gamma}^{\mathsf{cr}}]$$

$$= \Lambda n. \ \sqcap^{\mathsf{r}} \{ \sqcap^{\mathsf{r}} \{ \gamma_{\delta}^{\mathsf{cr}}(M^{\mathsf{cr}}(\rho(n))) \mid M^{\mathsf{cr}} \in m^{\mathsf{cr}} \} \mid \delta \in \Delta(N_{\rho(n)}, \hat{N}_n) \}$$

[by Theorem 11]

$$= \Lambda n. \ \sqcap^{\mathsf{r}} \left\{ \sqcap^{\mathsf{r}} \left\{ \gamma_{\delta}^{\mathsf{cr}}(M^{\mathsf{cr}}(\rho(n))) \mid \delta \in \Delta(N_{\rho(n)}, \hat{N}_n) \right\} \mid M^{\mathsf{cr}} \in m^{\mathsf{cr}} \right\}$$

[by assoc. of meets (\pi^{cr})]

г		Т.
		н
		I
		н
		н

Moreover, α_{δ}^{cr} is surjective and γ_{δ}^{cr} is injective for any $\delta \in \Delta(\hat{N}, N)$ given a concrete scope N. Hence, the Galois connection on collapsed values is in fact a Galois insertion.

Lemma 129. Let $\delta \in \Delta(\hat{N}, N)$ for some abstract scope \hat{N} and concrete scope N. Then, α_{δ}^{cr} is surjective and γ_{δ}^{cr} is injective. Proof. We show the injectivity of γ_{δ}^{cr} . The surjective nature of α_{δ}^{cr} then follows from Proposition 2. Suppose γ_{δ}^{cr} is not surjective. Then, there exist $u_1 \in \mathcal{V}_{\hat{N}}^{cr}$ and $u_2 \in \mathcal{V}_{\hat{N}}^{cr}$ s.t. $u_1 \neq u_2$ (modulo variable renaming) and $\gamma_{\delta}^{cr}(u_1) = \gamma_{\delta}^{cr}(u_2)$. We arrive at contradiction by structural induction on u_1 and u_2 . Let $\gamma_{\delta}^{cr}(u_1) = r_1$ and $\gamma_{\delta}^{cr}(u_2) = r_2$.

Base cases.

- $u_1 = \perp^{\mathsf{cr}}$ or $u_2 = \perp^{\mathsf{cr}}$. Assume $u_1 = \perp^{\mathsf{cr}}$ w.l.g. Then, $r_1 = \perp^{\mathsf{r}}$. By def. of $\gamma_{\delta}^{\mathsf{cr}}$, $r_2 = \perp^{\mathsf{r}}$ iff $u_2 = \perp^{\mathsf{cr}}$, contradiction.
- $u_1 = \top^{\mathsf{cr}}$ or $u_2 = \top^{\mathsf{cr}}$. Assume $u_1 = \top^{\mathsf{cr}}$ w.l.g. Then, $r_1 = \top^{\mathsf{r}}$. By def. of $\gamma_{\delta}^{\mathsf{cr}}$, $r_2 = \top^{\mathsf{r}}$ iff $u_2 = \top^{\mathsf{cr}}$, contradiction.
- $u_1 \in \mathcal{R}_{\hat{N}}^{\mathsf{cr}}, u_2 \in \mathcal{T}_{\hat{N}}^{\mathsf{cr}}$ or $u_1 \in \mathcal{T}_{\hat{N}}^{\mathsf{cr}}, u_1 \in \mathcal{R}_{\hat{N}}^{\mathsf{cr}}$. Assume $u_1 = \mathcal{R}_{\hat{N}}^{\mathsf{cr}}$ and $u_2 \in \mathcal{T}_{\hat{N}}^{\mathsf{cr}}$ w.l.g. Then, $r_1 \in \mathcal{R}_N^{\mathsf{r}}$ and $r_2 \in \mathcal{T}_N^{\mathsf{r}}$. Hence, $r_1 \neq r_2$.
- $u_1 \in \mathcal{R}_{\hat{N}}^{\mathsf{cr}}, u_2 \in \mathcal{R}_{\hat{N}}^{\mathsf{cr}}$. By def. of $\gamma_{\delta}^{\mathsf{cr}}$ and δ being a function it follows $r_1 \neq r_2$. Essentially, $\gamma_{\delta}^{\mathsf{cr}}$ simply *renames* dependency variables in dep. vectors using δ .

Induction case. Let $u_1 = \langle z, u_{1i}, u_{1o} \rangle$ and $u_2 = \langle z, u_{2i}, u_{2o} \rangle$. By $u_1 \neq u_2$, we have $u_{1i} \neq u_{2i}$ or $u_{1o} \neq u_{2o}$. The result then follows from the i.h.

A.5.1 Domain Operations

In this subsection, we state the basic properties and the corresponding proofs of strengthening operations on collapsed values.

A.5.1.1 Soundness, Strengthening, Idempotency, and Monotonicity

We start by showing soundness of the domain operations of the collapsed semantics subject to the relational semantics. We first show the proof for strengthening by imposing equality between two abstract nodes.

Lemma 130. Let $\delta_1 \in \Delta(\hat{N}, N)$ for some abstract scope \hat{N} and concrete scope N, $r \in \mathcal{V}_N^r$, $u \in \mathcal{V}_{\hat{N}}^{\mathsf{cr}}$, $\hat{n}_1 \in \hat{\mathcal{N}}$, and $\hat{n}_2 \in \hat{\mathcal{N}}$. If $r \sqsubseteq^\mathsf{r} \gamma_{\delta_1}^{\mathsf{cr}}(u)$, then $r[\delta_1(\hat{n}_1) = \delta_1(\hat{n}_2)] \sqsubseteq^\mathsf{r} \gamma_{\delta_1}^{\mathsf{cr}}(u[\hat{n}_1 = \hat{n}_2])$.

Proof. We recall $u[\hat{n}_1 = \hat{n}_2] = \bigsqcup^{\mathsf{cr}} \{ \alpha_{\delta}^{\mathsf{cr}}(\gamma_{\delta}^{\mathsf{cr}}(u)[\delta(\hat{n}_1) = \delta(\hat{n}_2)]) \mid N \subseteq \mathcal{N} \land \delta \in \Delta(\hat{N}, N) \}$ for a given concrete scope N.

$$r[\delta_{1}(\hat{n}_{1})=\delta_{1}(\hat{n}_{2})] \sqsubseteq^{\mathsf{r}} \gamma_{\delta}^{\mathsf{cr}}(u)[\delta_{1}(\hat{n}_{1})=\delta_{1}(\hat{n}_{2})] \qquad \text{by Lemma 100}$$

$$\alpha_{\delta_{1}}^{\mathsf{cr}}(r[\delta_{1}(\hat{n}_{1})=\delta_{1}(\hat{n}_{2})]) \sqsubseteq^{\mathsf{r}} \alpha_{\delta_{1}}^{\mathsf{cr}}(\gamma_{\delta}^{\mathsf{cr}}(u)[\delta_{1}(\hat{n}_{1})=\delta_{1}(\hat{n}_{2})]) \qquad \text{by Lemma 128}$$

$$\alpha_{\delta_{1}}^{\mathsf{cr}}(r[\delta_{1}(\hat{n}_{1})=\delta_{1}(\hat{n}_{2})]) \sqsubseteq^{\mathsf{r}} u[\delta_{1}(\hat{n}_{1})=\delta_{1}(\hat{n}_{2})] \qquad \text{by def. of } u[\delta_{1}(\hat{n}_{1})=\delta_{1}(\hat{n}_{2})]$$

$$r[\delta_{1}(\hat{n}_{1})=\delta_{1}(\hat{n}_{2})] \sqsubseteq^{\mathsf{r}} \gamma_{\delta_{1}}^{\mathsf{cr}}(u[\delta_{1}(\hat{n}_{1})=\delta_{1}(\hat{n}_{2})]) \qquad \text{by Prop. 1}$$

The above result can be made more general.

Lemma 131. Strengthening operations on collapsed values are sound.

Proof. The proof closely follows that of Lemma 130: the proof relies on the properties of monotonicity of abstraction and concretization functions and monotonicity of the corresponding strengthening operation defined over relational values. \Box

We next show that the strengthening operations on collapsed values indeed perform strengthening.

Lemma 132. Let $u_1 \in \mathcal{V}_{\hat{N}}^{\mathsf{cr}}$, $u \in \mathcal{V}_{\hat{N}}^{\mathsf{cr}}$, $\hat{n}_1 \in \hat{\mathcal{N}}$ and $\hat{n}_2 \in \hat{\mathcal{N}}$. Then $u[\hat{n}_1 = \hat{n}_2] \sqsubseteq_{\hat{N}}^{\mathsf{cr}} u$. *Proof.* Again, recall $u[\hat{n}_1 = \hat{n}_2] = \bigsqcup^{\mathsf{cr}} \{ \alpha_{\delta}^{\mathsf{cr}}(\gamma_{\delta}^{\mathsf{cr}}(u)[\delta(\hat{n}_1) = \delta(\hat{n}_2)]) \mid N \subseteq \mathcal{N} \land \delta \in \Delta(\hat{N}, N) \}$ for a given concrete scope N. Let $\delta_1 \in \Delta(\hat{N}, N)$ for any concrete scope

$$\gamma_{\delta_1}^{\mathsf{cr}}(u)[\delta_1(\hat{n}_1) = \delta_1(\hat{n}_2)] \sqsubseteq^{\mathsf{r}} \gamma_{\delta_1}^{\mathsf{cr}}(u) \qquad \text{by Lemma 100}$$
$$\alpha_{\delta_1}^{\mathsf{cr}}(\gamma_{\delta_1}^{\mathsf{cr}}(u)[\delta_1(\hat{n}_1) = \delta_1(\hat{n}_2)]) \sqsubseteq^{\mathsf{r}} u \qquad \text{by prop. of Galois conn.}$$

The result then follows from the glb property of the join \sqcup^{cr} operator (Theorem 8).

The above result and the proof can be applied to all strengthening operations on collapsed values.

Lemma 133. Strengthening operations on collapsed values are indeed strengthening.

Proof. The argument closely follows that of Lemma 132.

As expected, strengthening operations are also monotone.

Lemma 134. Strengthening operation on collapsed values are monotone.

Proof. The argument follows directly from the monotonicity of strengthening operations on relational values and monotonicity of the join operator \Box^{cr} on collapsed values.

In order to prove the idempotency of a strengthening of a collapsed value with another collapsed value, we first prove the following result.

Lemma 135. Let \hat{N} and \hat{N}_1 be abstract scopes, N and N_1 concrete scopes, $\hat{n} \in \hat{\mathcal{N}}$, $u \in \mathcal{V}_{\hat{N}}^{\mathsf{cr}}, \, u_1 \in \mathcal{V}_{\hat{N}_1}^{\mathsf{r}}, \, r \in \mathcal{V}_N^{\mathsf{r}}, \, \delta \in \Delta(\hat{N}, N), \, and \, \delta_1 \in \Delta(\hat{N}_1, N_1). \ Then, \, \alpha_{\delta}^{\mathsf{cr}}(r[\delta(\hat{n}) \leftarrow \delta_1 + \delta_2 +$ $\gamma_{\delta_1}^{\mathsf{cr}}(u_1)]) \sqsubseteq^{\mathsf{cr}} \alpha_{\delta}^{\mathsf{cr}}(r) [\hat{n} \leftarrow u_1].$

N.

Proof. We recall

 $\alpha_{\delta}^{\mathsf{cr}}(r)[\hat{n} \leftarrow u_1] = \bigsqcup^{\mathsf{cr}} \{ \alpha_{\delta}^{\mathsf{cr}}(\gamma_{\delta}^{\mathsf{cr}}(\alpha_{\delta}^{\mathsf{cr}}(r))[\delta_1(\hat{n}) \leftarrow \gamma_{\delta_1}^{\mathsf{cr}}(u_1)]) \mid \delta \in \Delta(\hat{N}, N), \delta_1 \in \Delta(\hat{N}_1, N_1) \}$ By basic properties of Galois connections, monotonicity of joins and strengthening operations on relational values (Theorem 104), $\alpha_{\delta}^{\mathsf{cr}}(r)[\hat{n} \leftarrow u_1]$ is larger than

$$u' = \bigsqcup^{\mathsf{cr}} \{ \alpha_{\delta}^{\mathsf{cr}}(r[\delta_1(\hat{n}) \leftarrow \gamma_{\delta_1}^{\mathsf{cr}}(u_1)]) \mid \delta \in \Delta(\hat{N}, N), \delta_1 \in \Delta(\hat{N}_1, N_1) \}$$

By definition of u' , we have $\alpha_{\delta}^{\mathsf{cr}}(r[\delta(\hat{n}) \leftarrow \gamma_{\delta_1}^{\mathsf{cr}}(u_1)]) \sqsubseteq^{\mathsf{cr}} u' \sqsubseteq^{\mathsf{cr}} \alpha_{\delta}^{\mathsf{cr}}(r)[\hat{n} \leftarrow u_1].$

We can now prove the idempotency of strengthening using a value.

Lemma 136. Let \hat{N} and \hat{N}_1 be abstract scopes, N and N_1 concrete scopes, $\hat{n} \in \hat{\mathcal{N}}$, $u \in \mathcal{V}_{\hat{N}}^{\mathsf{cr}}, u_1 \in \mathcal{V}_{\hat{N}_1}^{\mathsf{cr}}$. Then, $u[\hat{n} \leftarrow u_1] = u[\hat{n} \leftarrow u_1][\hat{n} \leftarrow u_1]$.

Proof. We first recall

 $u[\hat{n} \leftarrow u_1] = \bigsqcup^{\mathsf{cr}} \{ \alpha_{\delta}^{\mathsf{cr}}(\gamma_{\delta}^{\mathsf{cr}}(u)[\delta_1(\hat{n}) \leftarrow \gamma_{\delta_1}^{\mathsf{cr}}(u_1)]) \mid \delta \in \Delta(\hat{N}, N), \delta_1 \in \Delta(\hat{N}_1, N_1) \}$

By Theorem 133, we know $u[\hat{n} \leftarrow u_1][\hat{n} \leftarrow u_1] \sqsubseteq^{\mathsf{r}} u[\hat{n} \leftarrow u_1]$. We thus need to show $u[\hat{n} \leftarrow u_1] \sqsubseteq^{\mathsf{r}} u[\hat{n} \leftarrow u_1][\hat{n} \leftarrow u_1]$.

$$\begin{split} u[\hat{n} \leftarrow u_{1}] = \\ & [by \ def \ of \ \cdot [\cdot \leftarrow \cdot]] \\ & \bigsqcup^{\mathsf{cr}} \{ \ \alpha^{\mathsf{cr}}_{\delta}(\gamma^{\mathsf{cr}}_{\delta}(u)[\delta_{1}(\hat{n}) \leftarrow \gamma^{\mathsf{cr}}_{\delta_{1}}(u_{1})]) \ | \ \delta \in \Delta(\hat{N}, N), \delta_{1} \in \Delta(\hat{N}_{1}, N_{1}) \} = \\ & [by \ Lemma \ 103] \\ & \bigsqcup^{\mathsf{cr}} \{ \ \alpha^{\mathsf{cr}}_{\delta}(\gamma^{\mathsf{cr}}_{\delta}(u)[\delta_{1}(\hat{n}) \leftarrow \gamma^{\mathsf{cr}}_{\delta_{1}}(u_{1})]) \ | \ \delta \in \Delta(\hat{N}, N), \delta_{1} \in \Delta(\hat{N}_{1}, N_{1}) \} \sqsubseteq^{\mathsf{r}} \\ & [by \ Lemma \ 135] \\ & \bigsqcup^{\mathsf{cr}} \{ \ \alpha^{\mathsf{cr}}_{\delta}(\gamma^{\mathsf{cr}}_{\delta}(u)[\delta_{1}(\hat{n}) \leftarrow \gamma^{\mathsf{cr}}_{\delta_{1}}(u_{1})])[\hat{n} \leftarrow u_{1}] \ | \ \delta \in \Delta(\hat{N}, N), \delta_{1} \in \Delta(\hat{N}_{1}, N_{1}) \} \sqsubseteq^{\mathsf{r}} \\ & [by \ glb \ property \ of \ \sqcup^{\mathsf{cr}}] \\ & (\bigsqcup^{\mathsf{cr}} \{ \ \alpha^{\mathsf{cr}}_{\delta}(\gamma^{\mathsf{cr}}_{\delta}(u)[\delta_{1}(\hat{n}) \leftarrow \gamma^{\mathsf{cr}}_{\delta_{1}}(u_{1})]) \ | \ \delta \in \Delta(\hat{N}, N), \delta_{1} \in \Delta(\hat{N}_{1}, N_{1}) \})[\hat{n} \leftarrow u_{1}] = \\ & [by \ def \ of \ \cdot [\leftarrow \cdot]] \\ & u[\hat{n} \leftarrow u_{1}][\hat{n} \leftarrow u_{1}] \end{split}$$

A.5.1.2 Structural Strengthening

We also here provide a structural definition of strengthening on collapsed values and argue its soundness. First, whenever $z \notin \hat{N}$, then $u_{\hat{N}} \lceil z \leftarrow u \rceil \stackrel{\text{def}}{=} u$. Otherwise,

$$\begin{split} u\lceil \hat{n} \leftarrow \bot^{\mathsf{cr}} \rceil & \stackrel{\text{def}}{=} \bot^{\mathsf{cr}} \\ u\lceil \hat{n} \leftarrow u' \rceil & \stackrel{\text{def}}{=} u & \text{if } u \in \{\bot^{\mathsf{cr}}, \top^{\mathsf{cr}}\} \text{ or } u' \in \{\top^{\mathsf{cr}}\} \cup \mathcal{T}^{\mathsf{cr}} \\ R_1^{\mathsf{cr}}\lceil \hat{n} \leftarrow R_2^{\mathsf{cr}} \rceil & \stackrel{\text{def}}{=} R_1^{\mathsf{cr}}[\hat{n} \leftarrow R_2^{\mathsf{cr}}] \\ \langle z, u_i, u_o \rangle \lceil \hat{n} \leftarrow u \rceil & \stackrel{\text{def}}{=} \langle z, u_i \lceil \hat{n} \leftarrow u \rceil, u_o \lceil n \leftarrow \exists z. u \rceil \rangle \quad u \in \mathcal{R}^{\mathsf{cr}} \\ \text{It can be shown that the structural strengthening is sound.} \end{split}$$

Lemma 137. Structural strengthening is sound. Let $u \in \mathcal{V}_{\hat{N}}^{cr}$, $u' \in \mathcal{V}_{\hat{N}_1}^{cr}$, $\hat{n} \in \hat{\mathcal{N}}$,

 $r \in \mathcal{V}_N^{\mathsf{r}}, r' \in \mathcal{V}_{N_1}^{\mathsf{r}}, \delta \in \Delta(\hat{N}, N), and \delta_1 \in \Delta(\hat{N}_1, N_1).$ If $r \sqsubseteq^{\mathsf{r}} \gamma_{\delta}^{\mathsf{cr}}(u)$ and $r' \sqsubseteq^{\mathsf{r}} \gamma_{\delta_1}^{\mathsf{cr}}(u')$, then $r[\delta(\hat{n}) \leftarrow r'] \sqsubseteq^{\mathsf{r}} u[\hat{n} \leftarrow u'].$

Proof. The proof is carried by structural induction on u. The proof in fact shows that structural strengthening on collapsed values overapproximates the structural strengthening on relational values. The argument then follows from Lemma 111 and Lemma 131.

A.5.2 Propagation and Transformer

Proof (of Lemma 10).

Proof. For the increasing property, the proof goes by structural induction on both collapsed value arguments of $prop^{cr}$ and it closely follows the proof of Theorem 26. Regarding monotonicity, the proof again goes by structural induction on the collapsed arguments of $prop^{cr}$, closely follows the proof of Theorem 28, and relies on Lemma 134 (introduced later).

We next show that propagation in collapsed semantics is an abstraction of the propagation in relational semantics.

Theorem 13. Let $u_1 \in \mathcal{V}_{\hat{N}}^{cr}$, $u_2 \in \mathcal{V}_{\hat{N}}^{cr}$, \hat{N} an abstract scope, N a concrete scope, and $\delta \in \Delta(\hat{N}, N)$. Further, let $\langle r'_1, r'_2 \rangle = \operatorname{prop}^r(r_1, r_2)$ where $r_1 \sqsubseteq^r \gamma_{\delta}^{cr}(u_1)$ and $r_2 \sqsubseteq^r \gamma_{\delta}^{cr}(u_2)$. Lastly, let $\langle u'_1, u'_2 \rangle = \operatorname{prop}^{cr}(u_1, u_2)$. Then, $r'_1 \sqsubseteq^r \gamma_{\delta}^{cr}(u'_1)$ and $r'_2 \sqsubseteq^r \gamma_{\delta}^{cr}(u'_2)$.

Proof. The proof is carried by structural induction on u_1 and u_2 . Base cases. • $u_1 = \bot^{\mathsf{cr}}$.

$$\begin{split} u'_1 &= \bot^{\mathsf{cr}}, u'_2 = u_2 & \text{by prop}^{\mathsf{cr}} \\ r_1 &= \bot^{\mathsf{r}} & \text{by } \gamma^{\mathsf{cr}}_{\delta} \\ r'_1 &= \bot^{\mathsf{r}}, r'_2 = r_2 & \text{by prop}^{\mathsf{r}} \\ \gamma^{\mathsf{cr}}_{\delta}(u'_1) &= \bot^{\mathsf{r}}, \gamma^{\mathsf{cr}}_{\delta}(u'_2) = r'_2 & \text{by } \gamma^{\mathsf{cr}}_{\delta} \\ r'_1 &\sqsubseteq^{\mathsf{r}} \gamma^{\mathsf{cr}}_{\delta}(u'_1), r'_2 &\sqsubseteq^{\mathsf{r}} \gamma^{\mathsf{cr}}_{\delta}(u'_2) \end{split}$$

•
$$u_1 \in \mathcal{R}_{\hat{N}}^{cr}$$
 and $u_2 = \bot^{cr}$.
 $u'_1 = u_1, u'_2 = u_1$ by prop^{cr}
 $r'_2 = \bot^r$ by γ_{δ}^{cr}
 $r'_1 = r_1, r'_2 = r_1$ by prop^r
 $r'_1 \sqsubseteq^r \gamma_{\delta}^{cr}(u'_1), r'_2 \sqsubseteq^r \gamma_{\delta}^{cr}(u'_2)$

•
$$u_1 \in \mathcal{T}_{\tilde{N}}^{cr}$$
 and $u_2 = \bot^{cr}$.
 $u'_1 = u_1, u'_2 = T_{\bot^{cr}}^{cr}$ by prop^{cr}
 $r'_1 = r_1, r'_2 = T_{\bot^r}^r \lor r'_2 = \bot^r$ by prop^r
 $T_{\bot^r}^r = \gamma_{\delta}^{cr}(T_{\bot^{cr}}^{cr})$ by γ_{δ}^{cr}
 $r'_1 \sqsubseteq^r \gamma_{\delta}^{cr}(u'_1), r'_2 \sqsubseteq^r \gamma_{\delta}^{cr}(u'_2)$

• $u_1 = \top^{\operatorname{cr}}$.

$$u'_{1} = \top^{cr}, u'_{2} = \top^{cr} \qquad \text{by prop}^{cr}$$
$$\gamma^{cr}_{\delta}(u'_{1}) = \top^{cr}, \gamma^{cr}_{\delta}(u'_{2}) = \top^{cr} \qquad \text{by } \gamma^{cr}_{\delta}$$
$$r'_{1} \sqsubseteq^{r} \gamma^{cr}_{\delta}(u'_{1}), r'_{2} \sqsubseteq^{r} \gamma^{cr}_{\delta}(u'_{2})$$

• $u_1 \in \mathcal{R}_{\hat{N}}^{\mathsf{cr}}$ and $u_2 = \top^{\mathsf{cr}}$.

(

 $\begin{aligned} u_1' &= u_1, u_2' = u_2 & \text{by prop}^{\mathsf{cr}} \\ r_1' &= r_1 & \text{by prop}^{\mathsf{r}} \\ r_1' &\sqsubseteq^{\mathsf{r}} \gamma_{\delta}^{\mathsf{cr}}(u_1'), r_2' &\sqsubseteq^{\mathsf{r}} \gamma_{\delta}^{\mathsf{cr}}(u_2') \end{aligned}$

•
$$u_1 \in \mathcal{T}_{\hat{N}}^{\mathsf{cr}}$$
 and $u_2 = \top^{\mathsf{cr}}$.
 $u'_1 = \top^{\mathsf{cr}}, u'_2 = u_2 = \top^{\mathsf{cr}}$ by $\mathsf{prop}^{\mathsf{cr}}$
 $\gamma_{\delta}^{\mathsf{cr}}(u'_1) = \top^{\mathsf{r}}, \gamma_{\delta}^{\mathsf{cr}}(u'_2) = \top^{\mathsf{r}}$ by $\gamma_{\delta}^{\mathsf{cr}}$
 $r'_1 \sqsubseteq^{\mathsf{r}} \gamma_{\delta}^{\mathsf{cr}}(u'_1), r'_2 \sqsubseteq^{\mathsf{r}} \gamma_{\delta}^{\mathsf{cr}}(u'_2)$

•
$$u_1 \in \mathcal{T}_{\hat{N}}^{\mathsf{cr}}$$
 and $u_2 \in \mathcal{R}_{\hat{N}}^{\mathsf{cr}}$.
 $u'_1 = u_1, u'_2 = \top^{\mathsf{cr}}$ by $\mathsf{prop}^{\mathsf{cr}}$
 $r'_1 = r_1$ by $\mathsf{prop}^{\mathsf{r}}$
 $r'_1 \sqsubseteq^{\mathsf{r}} \gamma_{\delta}^{\mathsf{cr}}(u'_1), r'_2 \sqsubseteq^{\mathsf{r}} \gamma_{\delta}^{\mathsf{cr}}(u'_2)$

• $u_1 \in \mathcal{R}_{\hat{N}}^{\mathsf{cr}}$ and $u_2 \in \mathcal{T}_{\hat{N}}^{\mathsf{cr}}$. Similar to the previous case.

For the induction case, let $u_1 = \langle z, u_{1i}, u_{1o} \rangle$ and $u_1 = \langle z, u_{2i}, u_{2o} \rangle$. By γ_{δ}^{cr} , the only interesting case is when $r_1 \in \mathcal{T}_N^r$ and $r_2 \in \mathcal{T}_N^r$ per Theorem 10. For any n_{cs} , let $r_1(n_{cs}) = \langle r_{1i}, r_{1o} \rangle$ and $r_2(n_{cs}) = \langle r_{2i}, r_{2o} \rangle$. By $r_1 \sqsubseteq^r \gamma_{\delta}^{cr}(u_1)$ and $r_2 \sqsubseteq^r \gamma_{\delta}^{cr}(u_2)$ we have $r_{1i} \sqsubseteq^r \gamma_{\delta}^{cr}(u_{1i})$, $r_{1o} \sqsubseteq^r \gamma_{\delta}^{cr}(u_{1o})$, $r_{2i} \sqsubseteq^r \gamma_{\delta}^{cr}(u_{2i})$, and $r_{2o} \sqsubseteq^r$ $\gamma_{\delta[z \mapsto n_{cs}]}^{cr}(u_{2o})$. By i.h., we have $r'_{1i} \sqsubseteq^r \gamma_{\delta}^{cr}(u'_{1i})$ and $r'_{2i} \sqsubseteq^r \gamma_{\delta}^{cr}(u'_{2i})$. By monotonicity (Lemmas 99 and 134) and Lemma 135, $r_{1o}[n_{cs} \leftarrow r_{2i}] \sqsubseteq^r \gamma_{\delta[z \mapsto n_{cs}]}^{cr}(u_{1o}[z \leftarrow u_{2i}])$ and $r_{2o}[n_{cs} \leftarrow r_{2i}] \sqsubseteq^r \gamma_{\delta}^{cr}(u'_{1o})$ and $r'_{2o} \sqsubseteq^r \gamma_{\delta}^{cr}(u'_{2o})$. Since the resulting relational tables r'_1 and r'_2 can be expressed as respective joins \sqcup^r of singleton relational tables $[n_{cs} \mapsto \langle r'_{1i}, r'_{1o} \rangle]$ and $[n_{cs} \mapsto \langle r'_{2i}, r'_{1o} \rangle]$ for every call site n_{cs} , the result follows from the fact that the join \sqcup^r over relational values is lub (Lemma 4).

Proof (of Lemma 11).

Proof. The increasing property of $\mathsf{Step}_{k+1}^{\mathsf{cr}} \llbracket \cdot \rrbracket$ is obvious. The proof goes by structural induction on program expressions and it closely follows the proof of Theorem 29. The proof relies on Theorem 10 as well as Lemma 134.

Theorem 14. For all programs e, $\mathbf{C}^{\mathsf{r}}\llbracket e \rrbracket \doteq^{\mathsf{r}} \dot{\gamma}^{\mathsf{cr}}(\mathbf{C}\llbracket e \rrbracket)$.

Proof. First, observe that $M_{\perp^r} \stackrel{:}{\models} \stackrel{r}{\gamma^{cr}} (M_{\perp^{cr}}^{cr})$. The remainder of the proof is carried by structural induction on e showing that if $M^r \stackrel{:}{\models} \stackrel{r}{\gamma^{cr}} (M^{cr})$ then $M_1^r \stackrel{:}{\models} \stackrel{r}{\gamma^{cr}} (M_1^{cr})$ where $M_1^r = \operatorname{step}_{k+1}^r \llbracket e \rrbracket (E)(M^r)$, $M_1^{cr} = \operatorname{step}_{k+1}^{cr} \llbracket e \rrbracket (\rho(E))(M^{cr})$, and ρ function is raised to environments in the expected way. We focus on the lambda abstraction rule; other rules are even more trivial. For any n_{cs} , let $M_{n_{cs}}^r = \overline{M^r}(n_{cs})$. It follows in an almost straightforward fashion by Theorems 131 and 13 that $M_{n_{cs}}^r \stackrel{:}{\models} \stackrel{r}{\gamma^{cr}} (M_1^{cr})$. Since $\stackrel{.}{\sqcup}^r$ is a glb (by def. of $\stackrel{.}{\amalg}^r$ and Theorem 4), $\stackrel{.}{\amalg}_{n_{cs} \in T^r} \overline{M^r}(n_{cs}) \stackrel{.}{\models} \stackrel{r}{\gamma^r} (M_1^r)$.

We now switch to proving Lemma 12 that states that only a finite number of nodes is reachable by the collapsed semantics. Given a program e and a location ℓ in e, let $scope_vars_e(l)$ denote the sequence of program variable pairs (f, x) that are (1) bound by some lambda abstraction strictly above l in the abstract syntax tree (AST) of e and (2) ordered top-bottom based on their binding appearance in e. As an example, given a program $e = \mu f_1 . \lambda x_1 . (\mu f_2 . \lambda x_2 . e^{\ell})^{\ell_1}$, we have for instance $scope_vars_e(\ell) = \langle (f_1, x_1), (f_2, x_2) \rangle$ and $scope_vars_e(\ell_1) = \langle (f_1, x_1) \rangle$. The function $scope_vars_e$ can be simply computed based on the abstract syntax of e. Given an environment \hat{E} and pair of variables (f, x), let $\hat{E}[f, x] = \hat{E}.x:\hat{E} \diamond x.f:\hat{E} \diamond f$. Given a sequence s of variable pairs $(f_1, x_1) \cdots (f_j, x_j)$, let $aenv(s) = \emptyset[f_1, x_1] \cdots [f_j, x_j]$. We can thus define an abstract environment for a location ℓ in a program $e \in \lambda^d$ as
$$\begin{split} &\mathsf{aenv}_e(\ell) \stackrel{\text{def}}{=} \mathsf{aenv}(\mathsf{scope_vars}_e(\ell)). \text{ Given a program } e, \text{ we denote its abstract range} \\ &\mathsf{arange}(e) = \{ \hat{E} \diamond \ell \mid \ell \in Loc(e) \land s = \mathsf{scope_vars}_e(l) \land \hat{E} = \mathsf{aenv}(s) \}. \end{split}$$

Proof (of Lemma 12).

Proof. The proof follows from a straightforward inductive argument on e. Either $\mathbf{C}^{\mathsf{cr}}\llbracket e \rrbracket(k)(\hat{n}) = M_{\top^{\mathsf{cr}}}^{\mathsf{cr}}$ or the nodes updated by $\mathsf{step}_k^{\mathsf{cr}}\llbracket e \rrbracket$ belong to the set $\mathsf{arange}(e) \cup \{\hat{n} \mid \hat{n}' \in \mathsf{arange}(e) \land \hat{n} \in \mathsf{rng}(env(\hat{n}'))\}$ that is finite since (1) there are only a finite number of locations and lambda expressions in e, (2) program variables are unique, (3) and environments are finite partial functions.

A.6 Data Flow Refinement Type Semantics Proofs

As usual, we first prove that our definitions of refinement type meet and join are correct. A depth depth of a refinement type value is defined in the similar way as for the values of the collapsed semantics.

Proof (of Lemma 13).

Proof. Similar to the proof of Lemma 4, the proof is carried by case analysis on elements of U and induction on the minimum depth thereof when U consists of refinement function types only.

Proof (of Lemma 14 (Complete Meet Morphism of $\gamma_{\hat{N}}^{\mathsf{t}}$)).

Let $U \in \wp(\mathcal{V}_{\hat{N}}^{\mathsf{t}})$ and let \hat{N} be an abstract scope. We need to show that $\gamma^{\mathsf{t}}(\sqcap_{\hat{N}}^{\mathsf{t}}U) = \sqcap_{t \in U}^{\mathsf{cr}} \gamma^{\mathsf{t}}(t).$

Proof. We carry the proof by case analysis on elements of U and induction on the minimum depth thereof when U consists of tables only.

We first consider the trivial (base) cases where U is not a set of multiple refinement function types.

- $U = \emptyset$. Here, $\gamma^{\mathsf{t}}(\sqcap^{\mathsf{t}} \emptyset) = \gamma^{\mathsf{t}}(\top^{\mathsf{t}}) = \top^{\mathsf{cr}} = \sqcap^{\mathsf{cr}} \emptyset$.
- $U = \{t\}$. Trivial.
- $\{\nu: \cdot | \cdot\} \in U, T^{\mathsf{t}} \in U$. We have $\sqcap^{\mathsf{t}} \{\{\nu: \cdot | \cdot\}, T^{\mathsf{t}}\} = \bot^{\mathsf{t}}$ and since \bot^{t} is the bottom element, $\sqcap^{\mathsf{t}} U = \bot^{\mathsf{t}}$. Similarly, $\gamma^{\mathsf{t}}(\{\nu: \cdot | \cdot\}) \sqcap^{\mathsf{cr}} \gamma^{\mathsf{t}}(T^{\mathsf{t}}) = \bot^{\mathsf{cr}} = \gamma^{\mathsf{t}}(\bot^{\mathsf{t}})$ so by the definition of \sqcap^{cr} we have $\sqcap^{\mathsf{cr}}_{t \in U} \gamma^{\mathsf{t}}(t) = \bot^{\mathsf{r}}$.
- $\{\nu: \mathbf{b}_1 \mid \cdot\} \in U, \{\nu: \mathbf{b}_2 \mid \cdot\} \in U$, and $\mathbf{b}_1 \neq \mathbf{b}_2$. Similar to the previous case.

• $\exists \mathbf{b}. \forall t \in U. t = \{\nu : \mathbf{b} \mid \cdot\}.$ Here,

$$\begin{split} \gamma^{\mathsf{t}}(\sqcap^{\mathsf{t}} U) &= \gamma^{\mathsf{t}}(\sqcap^{\mathsf{t}}_{R^{\mathsf{t}} \in U} R^{\mathsf{t}}) \\ &= \{ D^{\mathsf{cr}} \mid D^{\mathsf{cr}} \in \gamma^{\mathsf{a}}(\sqcap^{\mathsf{a}} \{ a \mid \{ \nu \colon \mathsf{b} \mid a\} \in U \}) \land D^{\mathsf{cr}}(\nu) \in \gamma_{B}(\mathsf{b}) \} \end{split}$$

[by def. of γ^t and \sqcap^t]

$$= \bigcap_{\{\nu \mathsf{b} \mid a\} \in U} \{ D^{\mathsf{cr}} \mid D^{\mathsf{cr}} \in \gamma^{\mathsf{a}}(a) \land D^{\mathsf{cr}}(\nu) \in \gamma_B(\mathsf{b}) \}$$

[by the complete meet-morphism of $\gamma^{\rm a}]$

$$= \bigcap_{t \in U} \gamma^{\mathsf{t}}(t) = \sqcap_{t \in U}^{\mathsf{cr}} \gamma^{\mathsf{t}}(t) \qquad \qquad [\texttt{by def. of } \gamma^{\mathsf{t}} \texttt{ and } \sqcap^{\mathsf{cr}}]$$

• $\top^{\mathsf{t}} \in U$. Here, $\sqcap^{\mathsf{t}} U = \sqcap^{\mathsf{t}} (U/\{\top^{\mathsf{t}}\})$ and $\sqcap^{\mathsf{cr}}_{t \in U} \gamma^{\mathsf{t}}(t) = \sqcap^{\mathsf{cr}}_{t \in U, t \neq \top^{\mathsf{t}}} \gamma^{\mathsf{t}}(t)$ since $\gamma^{\mathsf{t}}(\top^{\mathsf{t}}) = \top^{\mathsf{cr}}$. The set $U/\{\top^{\mathsf{t}}\}$ either falls into one of the above cases or consists of multiple tables, which is the case we show next.

Let $U \subseteq \mathcal{T}_{\hat{N}}^{\mathsf{t}}$ and |U| > 1. Let d be the minimum depth of any table in U.

[by def. of γ^{t}]

$$= \langle z, \sqcap^{\mathsf{cr}}\{\gamma^{\mathsf{t}}(\pi_1(io(T^{\mathsf{t}}))) \mid T^{\mathsf{t}} \in U\}, \sqcap^{\mathsf{cr}}\{\gamma^{\mathsf{t}}(\pi_2(io(T^{\mathsf{t}}))) \mid T^{\mathsf{t}} \in U\}\rangle$$

[by i.h. on the min. depth of $\{\pi_1(io(T^t))) \mid T^t \in U\}$ and $\{\pi_2(io(T^t))) \mid T^t \in U\}$]

$$= \sqcap_{t \in U}^{\mathsf{cr}} \gamma^{\mathsf{t}}(t) \qquad \qquad [\mathbf{by \ def. \ of \ } \sqcap^{\mathsf{cr}}]$$

A.6.1 Domain Operations

Proof (of Lemma 15).
Proof. The proof goes by structural induction on the type being strengthened. The proof shows that structural strengthening on types overapproximates the structural strengthening on collapsed values. The argument then follows from Lemma 137.

It can be also shown that the structural strengthening indeed performs a strengthening given certain assumptions.

Lemma 138. Let the set of types $\mathcal{V}_{\hat{N}}^{\mathsf{t}}$ be instantiated with an abstract domain of type refinements $\mathcal{A}_{\hat{N}}$ with the strengthening property $\forall a \in \mathcal{A}_{\hat{N}}, a_1 \in \mathcal{A}_{\hat{N}}.a[z \leftarrow a_1] \sqsubseteq_{\hat{N}}^{\mathsf{a}} a$. Let $t \in \mathcal{V}_{\hat{N}}^{\mathsf{t}}$ and $t_1 \in \mathcal{V}_{\hat{N}}^{\mathsf{t}}$. Then, $t[z \leftarrow t_1] \sqsubseteq_{\hat{N}}^{\mathsf{t}} t$ where strengthening $\cdot [\cdot \leftarrow]$ is defined structurally as shown in Figure 3.10.

Proof. The proof goes by structural induction on t and case analysis of t_1 . For the cases where $t \in \{\perp^t, \top^t\}$ or $t_1 \in \{\perp^t, \top^t\} \cup \mathcal{T}^t$, the argument is trivial. For the case when both t and t_1 are basic refinement types, the argument follows from the strengthening property of abstract domain elements. The case when $t \in \mathcal{T}^t$ and $t_1 \in \mathcal{R}^t$ follows from the i.h.

We can also show that strengthening is monotone assuming that refinement domain strengthening and rescoping operations are monotone.

Lemma 139. Let $t \in \mathcal{V}_{\hat{N}}^{\mathsf{t}}$, $t' \in \mathcal{V}_{\hat{N}}^{\mathsf{t}}$, $t_1 \in \mathcal{V}_{\hat{N}}^{\mathsf{t}}$, $t_2 \in \mathcal{V}_{\hat{N}}^{\mathsf{t}}$, and $z \in DVar$. Assuming that strengthening operations on refinement domain $\mathcal{A}_{\hat{N}}$ are monotone, then $t_1 \sqsubseteq^{\mathsf{t}} t_2 \wedge t \sqsubseteq^{\mathsf{t}} t' \Longrightarrow t_1[z \leftarrow t] \sqsubseteq^{\mathsf{t}} t_2[z \leftarrow t']$.

Proof. When $z \notin \hat{N}$, the argument is straightforward. The proof goes by structural induction on t_1 and t_2 . But, let us first discuss the cases of t and t'. When $t = \perp^t$, the argument is trivial. The same applies when $t \in \mathcal{T}_{\hat{N}}^t$ or $t = \top^t$. Whenever $t' = \top^t$

or $t' \in \mathcal{T}_{\hat{N}}^{\mathsf{t}}$, the argument goes by Lemma 138. We can thus carry the induction argument by just considering the cases where $t, t' \in \mathcal{R}^{\mathsf{t}}$. Let $t = \{\nu : \mathsf{b}_0 \mid a_0\}$ and $t' = \{\nu : \mathsf{b}_0 \mid a'_0\}$. Let us cover the base cases first.

- $t_1 = \perp^t$. Here, $t_1[z \leftarrow t] = \perp^t$ and the result follows trivially.
- $t_2 = \perp^t$. Then it must be that $t_1 = \perp^t$ and the result follows from the previously analyzed case.
- $t_2 = \top^t$. We have $t_2[z \leftarrow t] = \top^t$ and the result holds trivially.
- $t_1 = \top^t$. Then, $t_2 = \top^t$ which is the case already covered.
- t₁ = {ν: b | a₁}, t₂ = {ν: b | a₂}, where a₁ ⊑^a a₂. We have t₁[z ← t] = {ν: b | a₁[z ← a₀]} and t₂[z ← t] = {ν: b | a₂[z ← a'₀]}. The result then follows from the monotonicity of the refinement strengthening.

For the induction case where $t_1 \in \mathcal{T}^t$ and $t_2 \in \mathcal{T}^t$, the argument follows straightforwardly from the i.h. and monotonicity of the refinement domain rescoping operation.

A.6.2 Abstract Propagation and Transformer

Proof (of Lemma 16).

Proof. The increasing property follows easily by the structural induction on both type arguments of $prop^t$. Regarding monotonicity, the argument goes on structural induction on the argument to $prop^t$ by closely following that of Lemma 28 and it relies on Lemma 139.

Proof (of Lemma 17).

Proof. The increasing property follows trivially. The argument for monotonicity follows that of Lemma 29 and it relies on the monotonicity of $prop^t$ established by Lemma 16.

Lemma 140. For all programs e and $k \in \mathbb{N}$, either $\mathbf{C}^{\mathsf{t}}[\![e]\!](k)(\hat{n}) = M_{\mathsf{T}^{\mathsf{t}}}^{\mathsf{t}}$ or the set $\{\hat{n} \in \hat{\mathcal{N}} \mid \mathbf{C}^{\mathsf{t}}[\![e]\!](k)(\hat{n}) \neq \bot^{\mathsf{t}}\}$ is finite.

Proof. The proof is the same as the proof of Lemma 12.

A.6.3 Widenings

Lemma 141. Let $t_1 \in \mathcal{V}_{\hat{N}}^{\mathsf{t}}$, $t_2 \in \mathcal{V}_{\hat{N}}^{\mathsf{t}}$ such that $t_1 \sqsubseteq_{\hat{N}}^{\mathsf{t}} t_2$. Then, $\mathsf{sh}(t_1) \sqsubseteq_{\hat{N}}^{\mathsf{t}} \mathsf{sh}(t_2)$.

Proof. The proof easily follows by structural induction on both t_1 and t_2 . \Box

Lemma 142. Let $\nabla_{\hat{N}}^{a}$ be a widening operator for the domain of abstract refinements $\mathcal{A}_{\hat{N}}$ used to instantiate the set of types $\mathcal{V}_{\hat{N}}^{\mathsf{t}}$ and the widening operator $\nabla_{\hat{N}}^{\mathsf{ra}}$. Also, let $t_{1} \in \mathcal{V}_{\hat{N}}^{\mathsf{t}}$, $t_{2} \in \mathcal{V}_{\hat{N}}^{\mathsf{t}}$, and $t = t_{1} \nabla_{\hat{N}}^{\mathsf{ra}} t_{2}$. Then, depth $(t) \leq \max(\operatorname{depth}(t_{1}), \operatorname{depth}(t_{2}))$.

Proof. The proof goes by structural induction on both t_1 and t_2 and it relies on the properties of the sh function.

Proof (of Lemma 18).

Proof. The upper bound property of $\nabla_{\hat{N}}^{t}$ directly follows from the upper bound properties of $\nabla_{\hat{N}}^{ra}$ and $\nabla_{\hat{N}}^{sh}$. Let $t_0 \sqsubseteq_{\hat{N}}^{t} t_1 \sqsubseteq_{\hat{N}}^{t} \dots$ be an increasing sequence of types. Let y^i be the sequence of types defined by $y_0 = t_0$ and $y_i = y_{i-1} \nabla_{\hat{N}}^{t} t_{i-1}$ for all i > 0. By $\nabla_{\hat{N}}^{t}$ being an upper bound operator it follows that $y_0 \sqsubseteq_{\hat{N}}^{t} y_1 \bigsqcup_{\hat{N}}^{t} \dots$ By Lemma 141, it is then the case that $\mathsf{sh}(y_0) \bigsqcup_{\hat{N}}^{t} \mathsf{sh}(y_1) \bigsqcup_{\hat{N}}^{t} \dots$ Now, $\nabla_{\hat{N}}^{ra}$ increases the shape of the two argument types by being increasing and by Lemma 141, but only by introducing \top^{t} (Lemma 142). Hence, either $\nabla_{\hat{N}}^{\mathsf{ra}}$ eventually produces \top^{t} in the sequence y^i , thus stabilizing it, or it eventually stops affecting the shape of types in y^i . In that case, the shape of y^i eventually stabilizes by the widening properties of $\nabla_{\hat{N}}^{\mathsf{sh}}$. The whole y^i then stabilizes by the widening operator property of $\nabla_{\hat{N}}^{\mathsf{a}}$ and the definition of $\nabla_{\hat{N}}^{\mathsf{ra}}$.

A.7 Liquid Types Semantics Proofs

For the rest of this section, we assume the set of Liquid types $\mathcal{V}_{\hat{N}}^{t}$ is defined as described in § 3.6.

A.7.1 Domain Operations

Lemma 143. The strengthening operations on Liquid types are sound and strengthening.

Proof. The results follow from the soundness and strengthening properties of the refinement domain (follow from the semantics of logical conjunction), Lemma 138, and Lemma 15. $\hfill \Box$

We can also show that strengthening operations are idempotent. Here, we assume that the strengthening of the liquid types refinement domain is idempotent. This is true in practice as strengthening amounts to adding equality conjuncts.

$$a[z \leftarrow a_1] = a[z \leftarrow a_1][z \leftarrow a_1]$$
 idempotency

Lemma 144. Let $t \in \mathcal{V}_{\hat{N}}^{\mathsf{t}}$, $t_1 \in \mathcal{V}_{\hat{N}}^{\mathsf{t}}$, and $z \in DVar$. Then, $t[z \leftarrow t_1] = t[z \leftarrow t_1][z \leftarrow t_1]$.

Proof. The proof goes by the structural induction on t. First, observe that if $z \notin \hat{N}$, the argument is trivial. Otherwise, if $t_1 = \bot^t$ then $t[z \leftarrow \bot^t] = \bot^t = t[z \leftarrow \bot^t][z \leftarrow \bot^t]$. Next, if $t_1 \in \{\top^t\} \cup \mathcal{T}^t$, then $t[z \leftarrow t_1] = t = t[z \leftarrow t_1][z \leftarrow t_1]$. We can thus just focus on the case where $t \in \mathcal{R}^t$. Let $t = \{\nu : \mathbf{b}_1 \mid a_1\}$. We cover the base cases first.

• $t = \perp^{\mathsf{t}}$. Here, $\perp^{\mathsf{t}}[z \leftarrow t_1] = \perp^{\mathsf{t}} = \perp^{\mathsf{t}}[z \leftarrow t_1][z \leftarrow t_1]$.

- $t = \top^{\mathsf{t}}$. Similar to the above case, $\top^{\mathsf{t}}[z \leftarrow t_1] = \top^{\mathsf{t}} = \top^{\mathsf{t}}[z \leftarrow t_1][z \leftarrow t_1]$.
- $t = \{\nu : \mathbf{b} \mid a\}$. In this case, $\{\nu : \mathbf{b} \mid a\}[z \leftarrow t_1] = \{\nu : \mathbf{b} \mid a[z \leftarrow a_1]\}$ and $\{\nu : \mathbf{b} \mid a\}[z \leftarrow t_1][z \leftarrow t_1] = \{\nu : \mathbf{b} \mid a[z \leftarrow a_1]\}[z \leftarrow t_1] = \{\nu : \mathbf{b} \mid a[z \leftarrow a_1]][z \leftarrow a_1]\}$. The result then follows from the idempotency of the refinement strengthening.

For the induction case where $t \in \mathcal{T}^{t}$, the result holds by the i.h. \Box

Lemma 145. Strengthening operations on the liquid abstract domain of refinement are idempotent.

Proof. Follows the proof of Lemma 144.

A.7.2 Data Propagation

We can now prove that the subbyping between types implies fixpoint in data propagation established between those types.

Proof (of Lemma 20).

Proof. The proof goes by an induction on the maximum depth of types $d = max(depth(t_1), depth(t_2))$. We first cover the base cases where d = 1. We only need to consider few cases based on the definition of $<:^q$.

• $t_1 = \{\nu : \mathbf{b} \mid a_1\}, t_2 = \{\nu : \mathbf{b} \mid a_2\}, \text{ and } a_1 \sqsubseteq^{\mathbf{a}} a_2. \text{ We have } \mathsf{prop}^{\mathsf{t}}(\{\nu : \mathbf{b} \mid a_1\}, \{\nu : \mathbf{b} \mid a_2\}) = \langle \{\nu : \mathbf{b} \mid a_1\}, \{\nu : \mathbf{b} \mid a_1 \sqcup^{\mathbf{a}} a_2\} \rangle = \langle \{\nu : \mathbf{b} \mid a_1\}, \{\nu : \mathbf{b} \mid a_2\} \rangle.$

Let us now consider the induction case for d > 1 where $t_1 \in \mathcal{T}^t$ and $t_2 \in \mathcal{T}^t$. Let $\mathsf{prop}^t(t_1, t_2) = \langle t_3, t_4 \rangle$. Observe that the maximum depth of input and output types of t_1 and t_2 and their arbitrary strengthening is smaller than d by the definition of

depth and Lemma 143. Let $t_1 = z : t_{1i} \to t_{1o}$ and $t_2 = z : t_{2i} \to t_{2o}$. By definition of $\langle :^q$, we have that (1) $t_{2i} \langle :^q t_{1i}$ and (2) $t_{1o}[z \leftarrow t_{2i}] \langle :^q t_{2o}[z \leftarrow t_{2i}]$. By i.h., we thus have $\operatorname{prop}^{\mathsf{t}}(t_{2i}, t_{1i}) = \langle t_{2i}, t_{1i} \rangle$ and $\operatorname{prop}^{\mathsf{t}}(t_{1o}[z \leftarrow t_{2i}], t_{2o}[z \leftarrow t_{2i}]) = \langle t_{1o}[z \leftarrow t_{2i}], t_{2o}[z \leftarrow t_{2i}] \rangle$. By Lemma 143, we have $t_{2o}[z \leftarrow t_{2i}] \sqsubseteq^{\mathsf{t}} t_{2o}$ and $t_{1o}[z \leftarrow t_{2i}] \sqsubseteq^{\mathsf{t}} t_{1o}$. Therefore, by the definition of $\operatorname{prop}^{\mathsf{t}}$ it follows $t_3 = z : t_{1i} \to t_{1o} = t_1$ and $t_4 = z : t_{2i} \to t_{2o} = t_2$.

A.7.3 Connection to Liquid Types

First, we define a typing relation $\vdash_{\mathsf{m}}^{\kappa}: \Xi^{\kappa} \times \lambda^{d} \times \mathcal{V}^{\mathsf{t}} \times \mathcal{M}^{\mathsf{t}}$ that mimics the earlier typing relation \vdash_{q}^{κ} , but also produces a type map that accumulates the results deduced in the typing derivation. To this end, given a κ and $\Gamma \in \Xi^{\kappa}$, we define $M_{\kappa,\Gamma}^{\mathsf{t}}$ to be a type map that holds $\Gamma(x)$ for all nodes $\hat{n} \in \mathsf{rng}(\kappa)$ where $\kappa(x) = \hat{n}$; otherwise, it returns \bot^{t} .

$$\begin{split} x \in \operatorname{dom}(\Gamma) & \Gamma(x) <:^{q} t \\ t = \Gamma(x)[\nu = \kappa(x)][\Gamma]^{\kappa} & t = c^{t}[\Gamma]^{\kappa} \\ \hline \Gamma \vdash_{\mathsf{m}}^{\kappa} x^{\ell} : t, M_{\Gamma,\kappa}^{t}[\kappa(\Gamma) \diamond \ell \mapsto t] & \Gamma \vdash_{\mathsf{m}}^{\kappa} c^{\ell} : t, M_{\Gamma,\kappa}^{t}[\kappa(\Gamma) \diamond \ell \mapsto t] \\ & \Gamma \vdash_{\mathsf{m}}^{\kappa} x^{\ell} : t, M_{\Gamma,\kappa}^{t}[\kappa(\Gamma) \diamond \ell \mapsto t] \\ & \frac{\Gamma \vdash_{\mathsf{m}}^{\kappa} e_{1}^{\ell_{1}} : t_{1}, M_{1}^{t} & \Gamma \vdash_{\mathsf{m}}^{\kappa} e_{2}^{\ell_{2}} : t_{2}, M_{2}^{t} \\ & \frac{t_{1} \in \mathcal{T}^{t} & t_{1} <:^{q} dx(t_{1}) : t_{2} \to t}{M^{t}, \Gamma \vdash_{\mathsf{m}}^{\kappa} (e_{1}^{\ell_{1}} e_{2}^{\ell_{2}})^{\ell} : t, (M_{1}^{t} \sqcup^{t} M_{2}^{t})[\kappa(\Gamma) \diamond \ell \mapsto t]} \\ & \frac{dx(t) : t_{x} \to t_{1} <:^{q} t & t <:^{q} t_{f}}{M^{t}, \Gamma \vdash_{\mathsf{m}}^{\kappa} (e_{1}^{\ell_{1}} e_{2}^{\ell_{2}})^{\ell} : t, (M_{1}^{t} \sqcup^{t} M_{2}^{t})[\kappa(\Gamma) \diamond \ell \mapsto t]} \\ & \frac{dx(t) : t_{x} \to t_{1} <:^{q} t & t <:^{q} t_{f}}{\Gamma_{1} = \Gamma.x : t_{x}.f : t_{f} & \kappa_{1} = \kappa.x : \kappa(\Gamma) \diamond x.f : \kappa(\Gamma) \diamond f & \Gamma_{1} \vdash_{\mathsf{m}}^{\kappa_{1}} e_{1}^{\ell_{1}} : t_{1}, M_{1}^{t}} \\ & \frac{\Gamma \vdash_{\mathsf{m}}^{\kappa} (\mu f. \lambda x. e_{1}^{\ell_{1}})^{\ell} : t, M_{1}^{t}[\kappa(\Gamma) \diamond \ell \mapsto t]}{\Gamma \vdash_{\mathsf{m}}^{\kappa} e_{1} : t_{1}, M_{1}^{t}} \\ & \frac{\Gamma \vdash_{\mathsf{m}}^{\kappa} x : t_{0}, M_{0}^{t} & t_{0} <:^{q} Bool^{t} & t_{1} <:^{q} t & t_{2} <:^{q} t}{\Gamma_{1} = \Lambda x \in \operatorname{dom}(\Gamma). \Gamma(x)[\kappa(x) \leftarrow t_{0} \sqcap^{t} true^{t}] & \Gamma_{1} \vdash_{\mathsf{m}}^{\kappa} e_{1} : t_{1}, M_{1}^{t}} \\ & \frac{\Gamma \vdash_{\mathsf{m}}^{\kappa} (x : e_{1}^{\ell_{1}} : e_{2}^{\ell_{2}})^{\ell} : t, (M_{0}^{t} \sqcup^{t} M_{1}^{t} \sqcup^{t} M_{2}^{t})[\kappa(\Gamma) \diamond \ell \mapsto t]} \\ \end{array}$$

Lemma 146. Let e be a program, κ an injective mapping from variables to abstract nodes, $t \in \mathcal{V}^{\mathsf{t}}$, and $\Gamma \in \Xi^{\kappa}$. Then, $\Gamma \vdash_{q}^{\kappa} e : t$ iff $\Gamma \vdash_{\mathsf{m}}^{\kappa} e : t$, M^{t} for some $M^{\mathsf{t}} \in \mathcal{M}^{\mathsf{t}}$.

Proof. The proof goes by structural induction on e. The result follows directly from the fact that $\vdash_{\mathsf{m}}^{\kappa}$ simply collects a type map M^{t} , without basing the rule premises on the contents of M^{t} .

We now show several interesting but rather obvious properties of $\vdash_{\mathsf{m}}^{\kappa}$.

Lemma 147. Let e be a program, κ an injective mapping from variables to abstract nodes, $t \in \mathcal{V}^{\mathsf{t}}$, and $\Gamma \in \Xi^{\kappa}$. If $\Gamma \vdash_{\mathsf{m}}^{\kappa} e : t, M^{\mathsf{t}}$ for some $M^{\mathsf{t}} \in \mathcal{M}^{\mathsf{t}}$, then $M^{\mathsf{t}}(\kappa(\Gamma) \diamond \ell) = t$.

Proof. The proof is carried trivially by structural induction on e.

Let domain validity domvalid(ℓ, p, Γ) of a type environment Γ subject to program p and a location ℓ in p holds iff dom(Γ) consists of exactly all variables bound above ℓ in p.

Lemma 148. Let e^{ℓ} be an expression of a program p, κ and injective mapping from variables to abstract nodes, $t \in \mathcal{V}^{\mathsf{t}}$, $\Gamma \in \Xi^{\kappa}$, and $\mathsf{domvalid}(\ell, p, \Gamma)$. If $\Gamma \vdash_{\mathsf{m}}^{\kappa} e : t, M^{\mathsf{t}}$ for some $M^{\mathsf{t}} \in \mathcal{M}^{\mathsf{t}}$, then $\forall x \in \mathsf{dom}(\Gamma)$. $M^{\mathsf{t}}(\kappa(x)) = \Gamma(x)$.

Proof. The proof is carried by structural induction on e. First, we note that programs $p \in \lambda^d$ have unique variables. Due to this, domain validity is preserved for updates to typing environments when typing the bodies of recursive functions. Further, every typing derivation has leaves at either variable or constant expressions where the environment assumptions are directly pushed to the maps. Finally, Γ_1 and Γ_2 are smaller than Γ when analyzing conditionals, due to strengthening, so the join of M_1^t and M_2^t with M_0^t ensures that the result holds since by i.h. $\forall x \in \mathsf{dom}(\Gamma). M_0^t(\kappa(x)) = \Gamma(x).$

Lemma 149. Let $e \in \lambda^d$, κ an injective mapping from variables to abstract nodes, $t \in \mathcal{V}^t$, $\Gamma \in \Xi^{\kappa}$, and $M^t \in \mathcal{M}^t$ such that $\Gamma \vdash_{\mathsf{m}}^{\kappa} e: t, M^t$. Then for every $\hat{n} \in \hat{\mathcal{N}}_e$, if $M^t(\hat{n}) \neq \bot^t$ then $loc(\hat{n}) \in Loc(e)$.

Proof. By structural induction on e.

Let $\operatorname{trange}(M^{\mathsf{t}}) \stackrel{\text{\tiny def}}{=} \{ \hat{n} \mid M^{\mathsf{t}}(\hat{n}) \neq \bot^{\mathsf{t}} \}$ be a set of non- \bot^{t} nodes of M^{t} .

Lemma 150. Let p be a program, ℓ_1 and ℓ_2 locations in p such that $Loc(p(\ell_1)) \cap Loc(p(\ell_2)) = \emptyset$, κ an injective mapping from variables to abstract nodes, $t_1 \in \mathcal{V}^t$, $t_2 \in \mathcal{V}^t$, $\Gamma \in \Xi^{\kappa}$, $M_1^t \in \mathcal{M}^t$, $M_2^t \in \mathcal{M}^t$, and domvalid (ℓ_1, p, Γ) as well as domvalid (ℓ_2, p, Γ) . Suppose $\Gamma \vdash_{\mathsf{m}}^{\kappa} p(\ell_1) : t_1, M_1^t$ and $\Gamma \vdash_{\mathsf{m}}^{\kappa} p(\ell_2) : t_2, M_2^t$. Then, trange $(M_1^t) \cap trange(M_2^t) = \mathsf{rng}(\kappa)$.

Proof. The argument follows from Lemma 149. Also, since p is a program, the program variables of lambda abstractions are unique. Therefore, the updates to Γ in the derivation for $p(\ell_1)$ and $p(\ell_2)$ are unique. Also, by domain validity of Γ , we have that κ and Γ are always extended (rather than being updated, i.e., some bindings being rewritten). Therefore, only the variable nodes associated with Γ , and hence in $\operatorname{rng}(\kappa)$, will and must appear in both M_1^t and M_2^t . \Box

Lemma 151. Let e^{ℓ} be an expression, κ an injective mapping from variables to abstract nodes, $t \in \mathcal{V}^{\mathsf{t}}$, $\Gamma \in \Xi^{\kappa}$ a valid typing environment, $M^{\mathsf{t}} \in \mathcal{M}^{\mathsf{t}}$, and $\mathsf{domvalid}(\ell, p, \Gamma)$. Suppose $\Gamma \vdash_{\mathsf{m}}^{\kappa} e^{\ell} : t, M^{\mathsf{t}}$. Then, M^{t} is safe and $t \neq \top^{\mathsf{t}}$.

Proof. The proof goes by structural induction on e. By Γ being valid and the definition of $\langle :^{q}$, further updates to the environment are also safe. Also, direct

updates to the maps are safe by the definition of $<:^q$. Finally, joins do not introduce errors by Lemma 150 and 148. The fact that $t \neq \top^t$ can be established then by Lemma 147.

We next define a typing relation $\vdash_{c}^{\kappa}: M^{t} \times \Gamma^{\kappa} \times \lambda^{d} \times t$ that checks if a given type map satisfies our Liquid typing relation. We say that two type maps M_{1}^{t} and M_{2}^{t}

$$\begin{split} x \in \operatorname{dom}(\Gamma) & \Gamma(x) <:^{q} t \\ \hline M^{\mathsf{t}}(\kappa(\Gamma) \diamond \ell) = t \quad t = \Gamma(x)[\nu = \kappa(x)][\Gamma]^{\kappa} & t = c^{\mathsf{t}}[\Gamma]^{\kappa} \quad M^{\mathsf{t}}(\kappa(\Gamma) \diamond \ell) = t \\ \hline M^{\mathsf{t}}, \Gamma \vdash_{\mathsf{c}}^{\kappa} x^{\ell} : t & \Pi^{\mathsf{t}}, \Gamma \vdash_{\mathsf{c}}^{\kappa} e^{\ell_{1}} : t_{1} & M^{\mathsf{t}}, \Gamma \vdash_{\mathsf{c}}^{\kappa} e^{\ell_{2}} : t_{2} & M^{\mathsf{t}}(\kappa(\Gamma) \diamond \ell) = t \\ \hline M^{\mathsf{t}}, \Gamma \vdash_{\mathsf{c}}^{\kappa} e^{\ell_{1}} : t_{1} & M^{\mathsf{t}}, \Gamma \vdash_{\mathsf{c}}^{\kappa} e^{\ell_{2}} : t_{2} & M^{\mathsf{t}}(\kappa(\Gamma) \diamond \ell) = t \\ \hline M^{\mathsf{t}}, \Gamma \vdash_{\mathsf{c}}^{\kappa} e^{\ell_{1}} : t_{1} & M^{\mathsf{t}}, \Gamma \vdash_{\mathsf{c}}^{\kappa} e^{\ell_{2}} : t_{2} & M^{\mathsf{t}}(\kappa(\Gamma) \diamond \ell) = t \\ \hline M^{\mathsf{t}}, \Gamma \vdash_{\mathsf{c}}^{\kappa} (e^{\ell_{1}} e^{\ell_{2}})^{\ell} : t & \\ \hline M^{\mathsf{t}}, \Gamma \vdash_{\mathsf{c}}^{\kappa} (e^{\ell_{1}} e^{\ell_{2}})^{\ell} : t & \\ \hline M^{\mathsf{t}}, \Gamma \vdash_{\mathsf{c}}^{\kappa} (e^{\ell_{1}} e^{\ell_{2}})^{\ell} : t & \\ \hline M^{\mathsf{t}}(\kappa(\Gamma) \diamond \ell) = t & M^{\mathsf{t}}(\kappa_{1}(x)) = t_{x} & M^{\mathsf{t}}(\kappa_{1}(f)) = t_{f} & M^{\mathsf{t}}(\kappa_{1}(\Gamma_{1}) \diamond \ell_{1}) = t_{1} \\ \hline M^{\mathsf{t}}, \Gamma \vdash_{\mathsf{c}}^{\kappa} (\mu f. \lambda x. e^{\ell_{1}})^{\ell} : t & \\ \hline M^{\mathsf{t}}(\kappa(\Gamma) \diamond \ell) = t & M^{\mathsf{t}}, \Gamma \vdash_{q}^{\kappa} x : t_{0} & t_{0} <:^{q} Bool^{\mathsf{t}} & t_{1} <:^{q} t & t_{2} <:^{q} t \\ \Gamma_{1} = \Lambda x \in \operatorname{dom}(\Gamma). \Gamma(x)[\kappa(x) \leftarrow t_{0} \sqcap^{\mathsf{t}} true^{\mathsf{t}}] & M^{\mathsf{t}}_{\Gamma_{1,\kappa}}, \Gamma_{2} \vdash_{q}^{\kappa} e_{1} : t_{1} \\ \Gamma_{2} = \Lambda x \in \operatorname{dom}(\Gamma). \Gamma(x)[\kappa(x) \leftarrow t_{0} \sqcap^{\mathsf{t}} false^{\mathsf{t}}] & M^{\mathsf{t}}_{\Gamma_{1,\kappa}}, \Gamma_{2} \vdash_{q}^{\kappa} e_{2} : t_{2} \\ \hline M^{\mathsf{t}}, \Gamma \vdash_{q}^{\kappa} (x^{\ell_{0}} ? e^{\ell_{1}} : e^{\ell_{2}})^{\ell} : t & \\ \hline \end{array}$$

agree, denoted tmagree (M_1^t, M_2^t) iff $\forall \hat{n}. M_1^t(\hat{n}) \neq \bot^t \land M_2^t(\hat{n}) \neq \bot^t \implies M_1^t(\hat{n}) = M_2^t(\hat{n}).$

Lemma 152. Let e^{ℓ} be an expression of p be a program, κ an injective mapping from variables to abstract nodes, $t \in \mathcal{V}^{\mathsf{t}}$, $\Gamma \in \Xi^{\kappa}$, $M^{\mathsf{t}} \in \mathcal{M}^{\mathsf{t}}$, $M_0^{\mathsf{t}} \in \mathcal{M}^{\mathsf{t}}$, and $\mathsf{domvalid}(\ell, p, \Gamma)$. If $\mathsf{tmagree}(M^{\mathsf{t}}, M_0^{\mathsf{t}})$ and $\Gamma \vdash_{\mathsf{m}}^{\kappa} e : t, M^{\mathsf{t}}$, then $M_0^{\mathsf{t}}, \Gamma \vdash_{\mathsf{c}}^{\kappa} e : t$.

Proof. By structural induction on e.

• $e^{\ell} = c^{\ell}$. Trivial.

- $e^{\ell} = x^{\ell}$. Also straightforward.
- e^ℓ = (e₁^{ℓ₁}e₂<sup>ℓ₂)^ℓ. By Lemma 150 and 148, we have that tmagree(M₁^t, M₀^t) and tmagree(M₂^t, M₀^t). The result then follows from the induction hypothesis.
 </sup>
- $e^{\ell} = (\mu f.\lambda x.e_1^{\ell_1})^{\ell}$. Similar to the previous case.
- $e^{\ell} = (x^{\ell_0}?e_1^{\ell_1}:e_2^{\ell_2})^{\ell}$. Similar, the argument again relies on Lemma 150 and 148.

We can now easily show an important connection between $\vdash_{\mathsf{m}}^{\kappa}$ and $\vdash_{\mathsf{c}}^{\kappa}$.

Lemma 153. Let e^{ℓ} be an expression of p be a program, κ an injective mapping from variables to abstract nodes, $t \in \mathcal{V}^{\mathsf{t}}$, $\Gamma \in \Xi^{\kappa}$, $M^{\mathsf{t}} \in \mathcal{M}^{\mathsf{t}}$, and $\mathsf{domvalid}(\ell, p, \Gamma)$. If $\Gamma \vdash_{\mathsf{m}}^{\kappa} e : t, M^{\mathsf{t}}$, then $M^{\mathsf{t}}, \Gamma \vdash_{\mathsf{c}}^{\kappa} e : t$.

Proof. Direct corollary of Lemma 152.

We say a triple (Γ, M^{t}, κ) matches if $\forall x \in \mathsf{dom}(\Gamma)$. $\Gamma(x) = M^{t}(\kappa(x))$.

Theorem 15. Let e^{ℓ} be an expression, κ an injective encoding of abstract variable nodes, $\Gamma \in \Xi^{\kappa}$, M^{t} a safe type map, and $t \in \mathcal{V}_{\hat{N}_{\kappa}}^{t}$. If $M^{t}, \Gamma \vdash_{\mathsf{c}}^{\kappa} e^{\ell} : t$ and (Γ, M^{t}, κ) match, then $\mathsf{Step}^{\mathsf{t}}[\![e^{\ell}]\!](\kappa(\Gamma) \diamond \ell)(M^{\mathsf{t}}) = M^{\mathsf{t}}$ and $M^{\mathsf{t}}(\kappa(\Gamma) \diamond \ell) = t$.

Proof. The proof goes by structural induction on e.

• $e = x^{\ell}$. By matching (Γ, M^{t}, κ) , we have that $\kappa(\Gamma)(x) = \kappa(x)$ and $M^{t}(\kappa(\Gamma)(x)) = \Gamma(x)$. Hence, $M^{t}(\kappa(\Gamma)(x)) <:^{q} M^{t}(\kappa(\Gamma) \diamond \ell)$. Then by the definition of $\vdash_{\mathsf{c}}^{\kappa}$, we have that $M^{t}(\kappa(\Gamma) \diamond \ell) = M^{t}(\kappa(\Gamma)(x))[\nu = \kappa(x)][\Gamma]^{\kappa}$ and hence by matching (Γ, M^{t}, κ) we have $M^{t}(\kappa(\Gamma) \diamond \ell) = M^{t}(\kappa(\Gamma)(x))[\nu = \kappa(\Gamma)(x)][M^{t}]$. Therefore,

 $M^{\mathsf{t}}(\kappa(\Gamma)(x))[\nu = \kappa(\Gamma)(x)] <:^{q} M^{\mathsf{t}}(\kappa(\Gamma) \diamond \ell)$ by the structural definition of strengthening. The result then follows from Lemma 20 and Lemma 145.

- $e = c^{\ell}$. The result follows from the definition of $c^{t}[\Gamma]^{\kappa}$. That is, $c^{t}[\Gamma]^{\kappa} = c^{t}[M^{t}]$ by matching (Γ, M^{t}, κ) .
- e = (e₁^{ℓ₁}e₂<sup>ℓ₂)^ℓ. By i.h., we have that Step^t [[e₁^{ℓ₁}]](κ(Γ))(M^t) = M^t and M^t(κ(Γ) ◊ℓ₁) = t₁ where t₁ ∈ T^t. Similarly, Step^t [[e₂^{ℓ₂}]](κ(Γ))(M^t) = M^t and M^t(κ(Γ) ◊ℓ₂) = t₂ where t₂ ∉ {⊥^t, ^{¬t}} by the definition of <:^q. Then by the definition of Step^t [[·]], we thus have that M₁^t = M₂^t = M^t and the result then follows from Lemma 20.
 </sup>
- $e = (\mu f.\lambda x. e^{\ell_1})^{\ell}$. First, by the definition of $<:^q$, it must be that $\pi_1(io(T^t)) \notin \{\perp^t, \top^t\}$. Further, by the definition of $\mathsf{Step}^t[\![\cdot]\!]$ and Lemma 20, we have that $T^t = T_1^t = T_2^t$, $t_x = t'_x$, $t_f = t'_f$, and $t_1 = t'_1[\hat{n}_x = dx(T^t)]$ since $dx(T^t)$ is not initially in the scope of t_1 (\hat{n}_x and $dx(T^t)$ are simply used as synonym variables). Next, by $\mathsf{domvalid}(\ell, p, \Gamma)$ it follows that $(\Gamma_1, \kappa_1, M^t)$ is matching and that domain validity is preserved $\mathsf{domvalid}(\ell_1, p, \Gamma_1)$. Also, by $\mathsf{Step}^t[\![\cdot]\!]$ we have that $\kappa(\Gamma_1) = \hat{E}_1$. Hence, $M_1^t = M^t$ by the induction hypothesis.
- $e = (x^{\ell_0} ? e_1^{\ell_1} : e_2^{\ell_2})^{\ell}$. First, observe that by definition of $<:^q$, if $t_0 <:^q Bool^t$, then $t_0 \neq \perp^t$ and $t_0 \equiv^t Bool^t$. Next, by i.h., we have that $M_0^t = M^t$ and t_0 in \vdash_c^{κ} and $\mathsf{Step}^t[\![\cdot]\!]$ are the same. Now, $\mathsf{Step}^t[\![\cdot]\!]$ strengthens the environment nodes of $\kappa(\Gamma)$ in M_0^t , using the node $\hat{E} \diamond \ell_0$, while the typing relation \vdash_c^{κ} uses $\kappa(x)$. However, by the typing rule of \vdash_c^{κ} for variable nodes, we know that $\Gamma(x) <:^q t_0$ where $t_0 = \Gamma(x)[\nu = \kappa(x)][M_0^t]$. By matching (Γ, κ, M_0^t) (that is implied by matching (Γ, κ, M^t) and $M_0^t = M^t$), we have $t_0 = M_0^t(\kappa(\Gamma)(x))[\nu = \kappa(\Gamma)(x)][M_0^t]$. By the definition of the liquid abstract

domain, t_0 implies that $\kappa(\Gamma)(x) = \hat{E} \diamond \ell_0$ and thus $\kappa(x) = \hat{E} \diamond \ell_0$. Since $\hat{E} \diamond \ell_0$ is not in the scope of any node in $\hat{N}_{\hat{E} \diamond \ell_0}$, we have $M_i^t = M_0^t [\hat{E} \diamond \ell_0 \leftarrow true^t \sqcap^t t_0] =$ $M_0^t [\kappa(x) \leftarrow true^t \sqcap^t t_0] = M_{\Gamma_1,\kappa}^t$ and $M_j^t = M_0^t [\hat{E} \diamond \ell_0 \leftarrow false^t \sqcap^t t_0] =$ $M_0^t [\kappa(x) \leftarrow false^t \sqcap^t t_0] = M_{\Gamma_2,\kappa}^t$. Therefore, $(\Gamma_1, \kappa_1, M_i^t)$ is matching and so is $(\Gamma_2, \kappa_2, M_j^t)$. By i.h. then, we have $M_1^t = M_i^t$ and $M_2^t = M_j^t$. By Lemma 20, we thus have that t = t' = t'', $t_1 = t'_1$, and $t_2 = t'_2$. The result then follows from the fact that both M_i^t and M_j^t , and consequently M_1^t and M_2^t , are smaller than M_0^t that is equal to M^t .

Proof (of Theorem 4).

Proof. Our implicit assumption is that e^{ℓ} is an expression of a program p such that domvalid (ℓ, p, Γ) . Then by Lemma 146, there exists a type map M^{t} such that $\Gamma \vdash_{q}^{\kappa} e : t$ iff $\Gamma \vdash_{m}^{\kappa} e : t, M^{t}$. Let M^{t} be such map. We know this map is safe by Lemma 151 and that it satisfies $M^{t}, \Gamma \vdash_{c}^{\kappa} e : t$. We also know by Lemma 148 that (Γ, M^{t}, κ) is matching. The result then follows from Theorem 15.

Appendix B

Type Error Localization

B.1 Proof of Lemma 25

The main idea behind our proof is centered around how the set of error sources for a particular program changes if we expand some let usage. In particular, we show that by expanding a well-typed let usage the number of error sources decreases. At the same time, we show that the number of proper error sources actually increases in the same scenario. In the case of a full expansion, when for instance $\mathfrak{L} = \mathfrak{L}_p$, it follows that all error sources are in fact proper. Using this, we show that a minimum error source that is proper for some \mathfrak{L} will be a minimum error source when we extend \mathfrak{L} to \mathfrak{L}_p .

We first state and prove few lemmas that the main proofs rely on. The next lemma states that typing derivations for the same expression using the same typing environments are equivalent modulo type variables used.

Lemma 154. Let $\Gamma, \Pi \vdash_{\mathfrak{L}} e^{\ell} \mid \mathcal{A}_1 \text{ and } \Gamma, \Pi \vdash_{\mathfrak{L}} e^{\ell} \mid \mathcal{A}_2$. Then, \mathcal{A}_1 and \mathcal{A}_2 are the same modulo consistent renaming of type variables not bound in Γ and Π .

Proof. This follows from the fact that the only source of non-determinism in our typing rules is the choice of names for type variables. \Box

The purpose of the next lemma is to show that every generated set of constraints can be trivially satisfied.

Lemma 155. Let $\Gamma, \Pi \vdash_{\mathfrak{L}} e^{\ell} \mid \mathcal{A}$. Then, \mathcal{A} is of the form $\{T_{\ell} \implies \mathcal{A}'\}$.

Proof. The proof goes by induction on typing derivations. By focusing on the last rule in the derivation, we can see that every rule generates a constraint with the above form. By setting T_{ℓ} to false, the generated set is then satisfied.

We now introduce notation and definitions used in our main arguments. Given a program p, we say that typing environment Γ_1 pointwise implies (for p) Γ_2 , written $\Gamma_1 \twoheadrightarrow_p \Gamma_2$, iff for all typing bindings $x : \forall \vec{\alpha_1} . (\mathcal{A}_1 \Longrightarrow \beta_1) \in \Gamma_1$ and x : $\forall \vec{\alpha_2} . (\mathcal{A}_2 \Longrightarrow \beta_2) \in \Gamma_2$ for the same variable x we have that $\forall \beta_1, \beta_2.\beta_1 = \beta_2 \implies$ $(\exists \vec{\alpha_1'}.\mathcal{A}_1 \land \exists \vec{\alpha_2}.\mathcal{A}_2 \land PDefs(p) \implies \exists \vec{\alpha_2'}.\mathcal{A}_2)$ is valid. Here, $\vec{\alpha_1'}$ is $\vec{\alpha_1}$ without β_1 (similarly for $\vec{\alpha_2'}$). When $\Gamma_1 \twoheadrightarrow_p \Gamma_2$ and $\Gamma_2 \twoheadrightarrow_p \Gamma_1$, we simply write $\Gamma_1 \ll p_2$.

We now state two facts that are important in the proofs that follow and which can be deduced from the proof of correctness of W algorithm.

- If e is a subexpression of p and $\Gamma_1 \xleftarrow{}{}_p \Gamma_2$, then $\rho(\Gamma_1, e) = \rho(\Gamma_2, e)$.
- Let e be a subexpression of p and $\Gamma, \Pi \vdash_{\mathfrak{L}} e : \beta_1 \mid \mathcal{A}_1$ where all T and P variables in \mathcal{A}_1 are assumed to be true. Given $\rho(\Pi, e) = \forall \vec{\delta}.\tau_{prin}$ and $\mathcal{A}_2 = \{\beta_2 = \tau_{prin}\}, \text{ then } \forall \beta_1, \beta_2.\beta_1 = \beta_2 \implies (\exists \vec{\alpha'}.\mathcal{A}_1 \land \exists \vec{\delta}, \beta_2.\mathcal{A}_2) \iff$ $(\exists \vec{\delta}.\mathcal{A}_2 \land \exists \vec{\alpha}.\mathcal{A}_1) \text{ holds. Here, } \vec{\alpha} \text{ is a vector of all free variables in } \mathcal{A}_1 \text{ except}$ those bound in Γ and Π . Also, $\vec{\alpha'}$ is $\vec{\alpha}$ without β_1 .

Given $\Gamma, \Pi \vdash_{\mathfrak{L}} e^{\ell} : \beta \mid \mathcal{A}$ where e is a subexpression of a program p, we define $C_{\Gamma,\Pi,p,\mathcal{A},\beta}$ to be the formula $\exists \vec{\alpha}.\mathcal{A} \land PDefs(p)$. Here, $\vec{\alpha}$ is a vector of the set $fv(\mathcal{A})/(fv(\Gamma) \cup fv(\Pi) \cup \{\beta\})$. When Γ, Π, p , and β are clear from the context, we simply write $C_{\mathcal{A}}$. Also for convenience, let $pwi_p(\Gamma_1, \Pi_2, \Gamma_2, \Pi_2)$ for a program phold iff $\Gamma_2 \twoheadrightarrow_p \Gamma_1, \Gamma_2 \twoheadrightarrow_p \Pi_2$, and $\Pi_1 \twoheadleftarrow \twoheadrightarrow_p \Pi_2$.

We are now ready to prove the first main lemma. We show that the number of error sources decreases for successive expansions of well typed let usages.

Lemma 156. Let e be an expression in a program p, \mathfrak{L} and \mathfrak{L}_u set of locations where $\mathfrak{L}_0 \subseteq \mathfrak{L} \subseteq \mathfrak{L}_u \subseteq \mathfrak{L}_p$. Also, assume $\mathfrak{L}_u = \mathfrak{L} \cup \{\ell_u\}$ such that ℓ_u is a location of a well-typed let variable usage in e. If $\Gamma, \Pi \vdash_{\mathfrak{L}} e^{\ell} : \alpha \mid \psi$ and $pwi_p(\Gamma, \Pi, \Gamma_u, \Pi_u)$, then from the typing derivation using \mathfrak{L} we can create a typing derivation using \mathfrak{L}_u such that the following holds:

- $\Gamma_u, \Pi_u \vdash_{\mathfrak{L}_u} e^{\ell} : \alpha \mid \psi_u$
- $M \models C_{\psi_u} \implies M \models C_{\psi}$

Proof. We will refer to typing derivations using \mathfrak{L} and \mathfrak{L}_u with d and d_u , respectively, using in general subscript u for parts of d_u . We carry the proof by induction on d, focusing on the last rule in the derivation. We start by observing that if $M \models C_{\psi_u}$ where $M \not\models T_\ell$, then by Lemma 155 $M \models C_{\psi}$ also holds. For simplicity, we will skip this trivial analysis of such models below; we only consider case where $M \models T_\ell$ and assume that C_{ψ} and C_{ψ_u} are simplified accordingly.

[A-Abs]. Suppose $pwi_p(\Gamma, \Pi, \Gamma_u, \Pi_u)$. Since $x : \alpha$ binding is trivial $(x : \alpha \text{ is})$ the same as $x : \forall \alpha.(\emptyset \Rightarrow \alpha))$, we have $pwi_p(\Gamma.x : \alpha, \Pi.x : \alpha, \Gamma_u.x : \alpha, \Pi_u.x; \alpha)$. We then create d_u by induction hypothesis, applying the same rule, and choosing the same name for α and γ as in d by appropriate type variable renaming

(Lemma 154). Now, suppose M is a model of C_{ψ_u} . Although α and β are quantified out in C_{ψ_u} , by semantics of existential quantification there exist a valuation to α and β that satisfies the constraint when those variables are free; let those values be α' and β' , respectively. Let M' then be $M[\alpha \to \alpha', \beta \to \beta']$. Then, M'is a model of $C_{\mathcal{A}_u}$ too. By induction hypothesis, it is also a model of $C_{\mathcal{A}}$, thus also being a model of C_{ψ} when α and β are made free. Then, M' is also a model of C_{ψ} and so must be M.

Similar argument can be given for the rules [A-HOLE], [A-INT], and [A-BOOL].

[A-App]. Suppose $pwi_p(\Gamma, \Pi, \Gamma_u, \Pi_u)$. Since the environments are the same for the rule assumptions, we can use the induction hypothesis. We create d_u similarly as in the previous rule. Suppose now M is a model of C_{ψ_u} . As in the previous rule, we can easily modify M into M' such that $M' \models C_{\mathcal{A}_{1_u}}, C_{\mathcal{A}_{2_u}}$; then by induction hypothesis it also must be that $M' \models C_{\mathcal{A}_1}, C_{\mathcal{A}_2}$. The result then follows from induction hypothesis on α and β as earlier.

Similar argument can be given for the rule [A-COND].

[A-Let-Prin]. Suppose $pwi_p(\Gamma, \Pi, \Gamma_u, \Pi_u)$. We first apply induction hypothesis for derivation for e_1 using \mathfrak{L} . We create derivation using \mathfrak{L}_u as before using induction hypothesis and renaming type variables where necessary. By induction hypothesis we then have $C_{\mathcal{A}_{1u}} \implies C_{\mathcal{A}_1}$. By this and semantics of existential quantification we also have that $\Gamma_u.x: \tau_{exp_u} \twoheadrightarrow_p \Gamma.x: \tau_{exp}$.

By the above first property of W algorithm, we have $\Pi.x : \tau_{prin} \ll \to_p \Pi_u.x : \tau_{prin_u}$. Finally, we show that $\Gamma_u.x : \tau_{exp_u} \to_p \Pi_u.x : \tau_{prin_u}$. If we assume that all T and P variables of \mathcal{A}_1 are set to **true**, then by the *PDefs* definition P_{ℓ_1} holds, and the argument follows from the second property of W algorithm we mention above. Otherwise, P_{ℓ_1} is set to false and the argument is trivial.

The proof then continues by applying the induction hypothesis on typing derivation for e_2 . The final argument is made as in the previous rules using induction hypothesis on \mathcal{A}_{2u} and FOL substitution lemma together with induction hypothesis on $\mathcal{A}_{1u}[\vec{\beta}/\vec{\alpha}]$.

Similar argument can be given for the rule [A-LET-EXP].

[A-Var-Prin]. We only consider the case when $\ell \neq \ell_u$. The argument is similar for the other case. Suppose $pwi_p(\Gamma, \Pi, \Gamma_u, \Pi_u)$. We then create d_u in the usual way also making sure it ends in rule [A-VAR-EXP], since $\ell \in \mathfrak{L}_u$. Now, let M be a model of C_{ψ_u} . The result then directly follows from the fact that $\Gamma_u \twoheadrightarrow_p \Pi_u$ and $\Pi_u \twoheadrightarrow_p \Pi$: by definition of pointwise implication it is also the case that $\Gamma_u \twoheadrightarrow_p \Pi$. Similar argument can be given for the rule [A-VAR-EXP].

To avoid clutter, we now change a bit the definition of pointwise implication, needed for the remainder of the proof, instead of redefining it from scratch. The definition again states the same, except the case where \mathcal{A}_1 is of the form $\{P_\ell \implies \beta_1 = \tau\}$. Then we require that $\forall \beta_1, \beta_2, \beta_1 = \beta_2 \implies (\exists \vec{\alpha'_1}.\mathcal{A}_1 \land \exists \vec{\alpha_2}.\mathcal{A}_2 \land PDefs(p) \implies$ $(P_\ell \implies \exists \vec{\alpha'_2}.\mathcal{A}_2))$. Also, we redefine $pwi_p(\Gamma_1, \Pi_1, \Gamma_2, \Pi_2)$ to hold for a program piff $\Gamma_2 \twoheadrightarrow_p \Gamma_1, \Pi_1 \twoheadrightarrow_p \Gamma_1$, and $\Pi_1 \ll \twoheadrightarrow_p \Pi_2$.

Let M be an FOL model and Γ a typing environment. We write $M \models \Gamma$ if $M \models \exists \vec{\alpha}.\mathcal{A}$ for every $x : \forall \vec{\alpha}.(\mathcal{A} \Rightarrow \beta) \in \Gamma$. We say $ms(M, \Gamma, \Pi, \Gamma', \Pi')$ iff $M \models \Gamma$, $M \models \Pi, M \models \Gamma'$, and $M \models \Pi'$.

We now show that every proper error source is still an error source after we

perform an expansion of a well-typed let variable. Argued later, such an error source is also proper after the expansion.

Lemma 157. Let e be an expression in a program p and \mathfrak{L} and \mathfrak{L}_u set of locations where $\mathfrak{L}_0 \subseteq \mathfrak{L} \subseteq \mathfrak{L}_u \subseteq \mathfrak{L}_p$. Also, let $\mathfrak{L}_u = \mathfrak{L} \cup \{\ell_u\}$ where ℓ_u is a location of a well-typed let variable usage. If $\Gamma, \Pi \vdash_{\mathfrak{L}} e : \alpha \mid \psi, M$ is an FOL model such that $M \models P_{dloc(\ell_u)}, ms(M, \Gamma, \Pi, \Gamma', \Pi')$, and $pwi_p(\Gamma', \Pi', \Gamma, \Pi)$, then from the typing derivation using \mathfrak{L} we can create a typing derivation using \mathfrak{L}_u such that:

- $\Gamma_u, \Pi_u \vdash_{\mathfrak{L}_u} e : \alpha \mid \psi_u$
- $M \models C_{\psi} \implies M \models C_{\psi_u}$

Proof. We borrow the notation from the proof of the previous lemma. We also carry the proof using the same technique. Likewise, we skip the trivial analysis when $M \not\models T_{\ell}$.

[A-Abs]. Suppose $ms(M, \Gamma, \Pi, \Gamma_u, \Pi_u)$ and $pwi_p(\Gamma_u, \Pi_u, \Gamma, \Pi)$. Assume $M \models C_{\psi}$ and $M \models P_{dloc(\ell_u)}$. Since $x : \alpha$ binding is trivial, we have $ms(M, \Gamma.x : \alpha, \Pi.x : \alpha, \Gamma_u.x : \alpha, \Pi_u.x : \alpha)$ and $pwi_p(\Gamma_u.x : \alpha, \Pi_u.x : \alpha, \Gamma.x : \alpha, \Pi.x : \alpha)$. The argument is then same as for the same rule in Lemma 156.

Similar argument can be given for the rules [A-HOLE], [A-INT], and [A-BOOL].

[A-App]. Suppose $ms(M, \Gamma, \Pi, \Gamma_u, \Pi_u)$ and $pwi_p(\Gamma_u, \Pi_u, \Gamma, \Pi)$. Assume $M \models C_{\psi}$ and $M \models P_{dloc(\ell_u)}$. Since the environments are the same for the rule assumptions, we can use the induction hypothesis. We proceed by using the same argument as for the same rule in Lemma 156.

Similar argument can be given for the rule [A-COND].

[A-Let-Prin]. Suppose $ms(M, \Gamma, \Pi, \Gamma_u, \Pi_u)$, $pwi_p(\Gamma_u, \Pi_u, \Gamma, \Pi)$, and M is a model of C_{ψ} where $M \models P_{dloc(\ell_u)}$. We first apply induction hypothesis for derivation for e_1 using \mathfrak{L} . We create derivation using \mathfrak{L}_u as before using induction hypothesis renaming type variables when necessary. By induction hypothesis we then have $C_{\mathcal{A}_1} \implies C_{\mathcal{A}_{1u}}$. By this and semantics of existential quantification we also have that $\Gamma.x : \tau_{exp} \twoheadrightarrow_p \Gamma_u.x : \tau_{exp_u}$.

By the above first property of W algorithm, we have $\Pi.x : \tau_{prin} \leftarrow \twoheadrightarrow_p \Pi_u.x : \tau_{prin_u}$. Finally, we show that $\Pi.x : \tau_{prin} \twoheadrightarrow_p \Gamma.x : \tau_{exp}$. We first note that if P_{ℓ_1} is set to **false**, the statement trivially follows. Otherwise, the statement follows from the properties of correctness of W algorithm.

The fact that $ms(M, \Gamma.x : \tau_{exp}, \Pi.x : \tau_{prin}, \Gamma_u.x : \tau_{exp_u}, \Pi_u.x : \tau_{prin_u})$ follows from the fact that $M \models \mathcal{A}_1[\vec{\beta}/\vec{\alpha}] \ (\mathcal{A}_{1_u}[\vec{\beta}/\vec{\alpha}]), \ C_{\mathcal{A}_1} \implies C_{\mathcal{A}_{1_u}}$, and FOL substitution lemma.

Again, the argument then similarly proceeds as in Lemma 156. Similar proof can be given for the rule [A-LET-EXP].

[A-Var-Prin]. We only consider the case when $\ell \neq \ell_u$. The argument is similar for the other case. We create d_u in the usual way also making sure it ends in rule [A-VAR-EXP], since $\ell \in \mathfrak{L}_u$. Suppose $ms(M, \Gamma, \Pi, \Gamma_u, \Pi_u)$, $pwi_p(\Gamma_u, \Pi_u, \Gamma, \Pi)$, and M is a model of C_{ψ} where $M \models P_{dloc(\ell_u)}$. First, P_{ℓ_1} in ψ , coming from $\mathcal{A}[\vec{\beta}/\vec{\alpha}]$, is in fact $P_{dloc(\ell_u)}$ which is set to **true** by assumption. Using this, the argument follows from $\Pi \twoheadrightarrow_p \Gamma_u$, which transitively holds from $\Pi \twoheadrightarrow_p \Gamma$ and $\Gamma \twoheadrightarrow_p \Gamma_u$.

Lemma 158. Let p be a program, R a cost function, \mathfrak{L} and \mathfrak{L}_u set of locations where $\mathfrak{L}_0 \subseteq \mathfrak{L} \subseteq \mathfrak{L}_u \subseteq \mathfrak{L}_p$ such that $\mathfrak{L}_u = \mathfrak{L} \cup \{\ell_u\}$ and $\ell_u \in Scope(p, \mathfrak{L})$, where ℓ_u is a location of a well-typed **let** variable usage. If $M = SOLVE(p, R, \mathfrak{L})$ where L_M is minimum and proper, then $M = SOLVE(p, R, \mathfrak{L}_u)$ and L_M is minimum and proper again.

Proof. We first note that SOLVE returns solutions for \mathcal{A} where $\emptyset, \emptyset \vdash_{\mathfrak{L}} e : \alpha \mid \mathcal{A}$. Then, by Lemma 156, Lemma 154, and semantics of existential quantification, we have that the number of error sources reduces or stays the same when instead of using \mathfrak{L} we use \mathfrak{L}_u . Therefore, since M is minimum among all models before expansion, it will also be a minimum model after the expansion as R is fixed. Now, since $\ell_u \notin \mathfrak{L}$ and M is proper for \mathfrak{L} , we have that $\ell_u \notin Usages(p, \mathfrak{L}, M)$. Given our assumption that $\ell_u \in Scope(p, \mathfrak{L})$, it must be the that $M \models P_{dloc(\ell_u)}$. Then we have by Lemma 157 that proper error sources are still error sources when instead of \mathfrak{L} we use \mathfrak{L}_u . Also by the definition of *PDefs* and *Scope*, we have that $Usages(p, \mathfrak{L}, M) = Usages(p, \mathfrak{L}_u, M)$. We hence have that L_M is also proper. \Box

We can finally prove Lemma 25.

Proof. From Lemma 24, we know it suffices to use \mathfrak{L}_p in order for the generated constraints to have the minimum error source as the correct solution. We also know $\mathfrak{D} = \mathfrak{L}_p \setminus \mathfrak{L}$ consists of locations corresponding to usages of well-typed let definitions. By the definition of *Scope* and *PDefs*, this set can be ordered in a list so that each consecutive usage location is in the scope if we expand one by one from the beginning all usage locations appearing earlier in the list. For each such consecutive expansion inductively, Lemma 158 guarantees that the proper minimum error source before the expansion is still a proper minimum error source

after the expansion. After expanding all usages, our expansion location set \mathfrak{L} is exactly \mathfrak{L}_p , guaranteeing the correct solution since in that case all error sources are in fact proper.

Bibliography

- A. Aiken and E. L. Wimmers. Type inclusion constraints and type inference. In *FPCA*, pages 31–41. ACM, 1993.
- [2] A. W. Appel and D. A. McAllester. An indexed model of recursive types for foundational proof-carrying code. ACM Trans. Program. Lang. Syst., 23(5):657–683, 2001.
- [3] V. Arceri, M. D. Preda, R. Giacobazzi, and I. Mastroeni. SEA: string executability analysis by abstract interpretation. *CoRR*, abs/1702.02406, 2017.
- [4] T. Ball, R. Majumdar, T. D. Millstein, and S. K. Rajamani. Automatic predicate abstraction of C programs. In *Proceedings of the 2001 ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI), Snowbird, Utah, USA, June 20-22, 2001*, pages 203–213, 2001.
- [5] C. Barrett, C. L. Conway, M. Deters, L. Hadarean, D. Jovanović, T. King, A. Reynolds, and C. Tinelli. CVC4. In CAV, pages 171–177. Springer-Verlag, 2011.
- [6] C. Barrett, R. Sebastiani, S. A. Seshia, and C. Tinelli. Satisfiability Modulo Theories, chapter 26, pages 825–885. Volume 185 of Biere et al. [9], February 2009.

- [7] C. Barrett, I. Shikanian, and C. Tinelli. An abstract decision procedure for a theory of inductive data types. *Journal on Satisfiability, Boolean Modeling* and Computation, 3:21–46, 2007.
- [8] C. Barrett, A. Stump, and C. Tinelli. The SMT-LIB standard version 2.0. In SMT, 2010.
- [9] A. Biere, M. J. H. Heule, H. van Maaren, and T. Walsh, editors. Handbook of Satisfiability, volume 185 of Frontiers in Artificial Intelligence and Applications. IOS Press, February 2009.
- [10] N. Bjørner and A. Phan. ν Z: Maximal Satisfaction with Z3. In SCSS, 2014.
- [11] N. Bjørner, A. Phan, and L. Fleckenstein. ν Z: An Optimizing SMT Solver. In *TACAS*, 2015.
- [12] N. E. Boustani and J. Hage. Improving type error messages for generic java. *Higher-Order and Symbolic Computation*, 24(1-2):3–39, 2011.
- [13] S. Chen and M. Erwig. Counter-factual typing for debugging type errors. In POPL, pages 583–594. ACM, 2014.
- [14] O. Chitil. Compositional explanation of types and algorithmic debugging of type errors. In *ICFP*, ICFP '01, pages 193–204. ACM, 2001.
- [15] P. Cousot. Types as abstract interpretations. In Proceedings of the 24th ACM SIGPLAN-SIGACT symposium on Principles of programming languages, pages 316–331. ACM, 1997.
- [16] P. Cousot and R. Cousot. Abstract interpretation: a unified lattice model for static analysis of programs by construction or approximation of fixpoints. In

Proceedings of the 4th ACM SIGACT-SIGPLAN symposium on Principles of programming languages, pages 238–252. ACM, 1977.

- [17] P. Cousot and R. Cousot. Systematic design of program analysis frameworks. In Proceedings of the 6th ACM SIGACT-SIGPLAN symposium on Principles of programming languages, pages 269–282. ACM, 1979.
- [18] P. Cousot and R. Cousot. Invited talk: Higher order abstract interpretation (and application to comportment analysis generalizing strictness, termination, projection, and PER analysis. In *Proceedings of the IEEE Computer Society* 1994 International Conference on Computer Languages, May 16-19, 1994, Toulouse, France, pages 95–112, 1994.
- [19] P. Cousot and N. Halbwachs. Automatic discovery of linear restraints among variables of a program. In Conference Record of the Fifth Annual ACM Symposium on Principles of Programming Languages, Tucson, Arizona, USA, January 1978, pages 84–96, 1978.
- [20] L. Damas and R. Milner. Principal type-schemes for functional programs. In POPL, pages 207–212. ACM, 1982.
- [21] L. De Moura and N. Bjørner. Z3: An Efficient SMT Solver. In TACAS, pages 337–340. Springer-Verlag, 2008.
- [22] D. Duggan and F. Bent. Explaining type inference. In Science of Computer Programming, pages 37–83, 1995.
- [23] B. Dutertre and L. de Moura. The Yices SMT solver. Technical report, SRI International, 2006.

- [24] EasyOCaml. http://easyocaml.forge.ocamlcore.org. [Online; accessed 10-March-2014].
- [25] J. Eremondi, W. Swierstra, and J. Hage. A framework for improving error messages in dependently-typed languages. Open Computer Science, 9(1):1–32, 2019.
- [26] C. Flanagan and K. R. M. Leino. Houdini, an annotation assistant for esc/java. In FME 2001: Formal Methods for Increasing Software Productivity, International Symposium of Formal Methods Europe, Berlin, Germany, March 12-16, 2001, Proceedings, volume 2021 of Lecture Notes in Computer Science, pages 500–517. Springer, 2001.
- [27] C. Flanagan and S. Qadeer. Predicate abstraction for software verification. In Conference Record of POPL 2002: The 29th SIGPLAN-SIGACT Symposium on Principles of Programming Languages, Portland, OR, USA, January 16-18, 2002, pages 191–202, 2002.
- [28] T. S. Freeman and F. Pfenning. Refinement types for ML. In Proceedings of the ACM SIGPLAN'91 Conference on Programming Language Design and Implementation (PLDI), Toronto, Ontario, Canada, June 26-28, 1991, pages 268–277, 1991.
- [29] R. Garcia, A. M. Clark, and É. Tanter. Abstracting gradual typing. In Proceedings of the 43rd Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL 2016, St. Petersburg, FL, USA, January 20 - 22, 2016, pages 429–442, 2016.

- [30] H. Gast. Explaining ML type errors by data flows. In Implementation and Application of Functional Languages, pages 72–89. Springer, 2005.
- [31] R. Gori and G. Levi. An experiment in type inference and verification by abstract interpretation. In Verification, Model Checking, and Abstract Interpretation, Third International Workshop, VMCAI 2002, Venice, Italy, January 21-22, 2002, Revised Papers, pages 225–239, 2002.
- [32] R. Gori and G. Levi. Properties of a type abstract interpreter. In Verification, Model Checking, and Abstract Interpretation, 4th International Conference, VMCAI 2003, New York, NY, USA, January 9-11, 2002, Proceedings, pages 132–145, 2003.
- [33] C. Haack and J. B. Wells. Type error slicing in implicitly typed higher-order languages. Sci. Comput. Program., pages 189–224, 2004.
- [34] J. Hage. Language Implementation Patterns: Create your own Domain-Specific and General Programming Languages, by terence parr, pragmatic bookshelf, http://www.pragprog.com, ISBN 9781934356456. J. Funct. Program., 21(2):215-217, 2011.
- [35] J. Hage and B. Heeren. Heuristics for type error discovery and recovery. In Implementation and Application of Functional Languages, 18th International Symp osium, IFL 2006, Budapest, Hungary, September 4-6, 2006, Revised Selected Papers, pages 199–216, 2006.
- [36] J. Hage and B. Heeren. Strategies for solving constraints in type and effect systems. *Electr. Notes Theor. Comput. Sci.*, 236:163–183, 2009.

- [37] R. Harper. Constructing type systems over an operational semantics. J. Symb. Comput., 14(1):71–84, 1992.
- [38] The Haskell Programming Language. http://www.haskell.org/. [Online; accessed 15-March-2014].
- [39] M. Hassan, C. Urban, M. Eilers, and P. Müller. Maxsmt-based type inference for python 3. In Computer Aided Verification - 30th International Conference, CAV 2018, Held as Part of the Federated Logic Conference, FloC 2018, Oxford, UK, July 14-17, 2018, Proceedings, Part II, pages 12–19, 2018.
- [40] J. R. Hindley. The principal type-scheme of an object in combinatory logic. Transactions of the American Mathematical Society, 146:2960, 1969.
- [41] S. Jagannathan and S. Weeks. A unified treatment of flow analysis in higherorder languages. In Proceedings of the 22Nd ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, pages 393–407. ACM, 1995.
- [42] S. H. Jensen, A. Møller, and P. Thiemann. Type analysis for javascript. In Static Analysis, 16th International Symposium, SAS 2009, Los Angeles, CA, USA, August 9-11, 2009. Proceedings, pages 238–255, 2009.
- [43] R. Jhala, R. Majumdar, and A. Rybalchenko. HMC: verifying functional programs using abstract interpreters. In Computer Aided Verification - 23rd International Conference, CAV 2011, Snowbird, UT, USA, July 14-20, 2011. Proceedings, pages 470–485, 2011.
- [44] N. D. Jones and N. Andersen. Flow analysis of lazy higher-order functional programs. *Theor. Comput. Sci.*, 375(1-3):120–136, Apr. 2007.

- [45] N. D. Jones and A. Mycroft. Data flow analysis of applicative programs using minimal function graphs. In Conference Record of the Thirteenth Annual ACM Symposium on Principles of Programming Languages, St. Petersburg Beach, Florida, USA, January 1986, pages 296–306, 1986.
- [46] N. D. Jones and M. Rosendahl. Higher-order minimal function graphs. Journal of Functional and Logic Programming, 1997(2), 1997.
- [47] M. Jose and R. Majumdar. Bug-Assist: Assisting Fault Localization in ANSI-C Programs. In CAV, pages 504–509. Springer-Verlag, 2011.
- [48] M. Kazerounian, N. Vazou, A. Bourgerie, J. S. Foster, and E. Torlak. Refinement types for ruby. In Verification, Model Checking, and Abstract Interpretation - 19th International Conference, VMCAI 2018, Los Angeles, CA, USA, January 7-9, 2018, Proceedings, pages 269–290, 2018.
- [49] A. J. Kfoury, J. Tiuryn, and P. Urzyczyn. ML Typability is DEXTIME-Complete. In CAAP, pages 206–220, 1990.
- [50] S. Kim and K. Choe. String analysis as an abstract interpretation. In Verification, Model Checking, and Abstract Interpretation - 12th International Conference, VMCAI 2011, Austin, TX, USA, January 23-25, 2011. Proceedings, pages 294–308, 2011.
- [51] K. Knowles and C. Flanagan. Type reconstruction for general refinement types. In Programming Languages and Systems, 16th European Symposium on Programming, ESOP 2007, Held as Part of the Joint European Conferences on Theory and Practics of Software, ETAPS 2007, Braga, Portugal, March 24 - April 1, 2007, Proceedings, pages 505–519, 2007.

- [52] S. K. Lahiri and S. Qadeer. Complexity and algorithms for monomial and clausal predicate abstraction. In Automated Deduction - CADE-22, 22nd International Conference on Automated Deduction, Montreal, Canada, August 2-7, 2009. Proceedings, volume 5663 of Lecture Notes in Computer Science, pages 214–229. Springer, 2009.
- [53] N. Lehmann and É. Tanter. Gradual refinement types. In Proceedings of the 44th ACM SIGPLAN Symposium on Principles of Programming Languages, POPL 2017, Paris, France, January 18-20, 2017, pages 775–788, 2017.
- [54] B. Lerner, M. Flower, D. Grossman, and C. Chambers. Searching for typeerror messages. In *PLDI*. ACM Press, 2007.
- [55] X. Leroy, D. Doligez, A. Frisch, J. Garrigue, D. Rémy, and J. Vouillon. OCaml Manual: Module Pervasives. http://caml.inria.fr/pub/docs/ manual-ocaml/libref/Pervasives.html. [Online; accessed 14-March-2014].
- [56] C. M. Li and F. Manyà. MaxSAT, Hard and Soft Constraints, chapter 19, pages 613–631. Volume 185 of Biere et al. [9], February 2009.
- [57] H. G. Mairson. Deciding ML Typability is Complete for Deterministic Exponential Time. In POPL, pages 382–401, 1990.
- [58] J. Midtgaard. Control-flow analysis of functional programs. ACM Comput. Surv., 44(3):10:1–10:33, 2012.
- [59] R. Milner. A theory of type polymorphism in programming. J. Comput. Syst. Sci., 17(3):348–375, 1978.

- [60] B. Monsuez. Polymorphic typing by abstract interpretation. In International Conference on Foundations of Software Technology and Theoretical Computer Science, pages 217–228. Springer, 1992.
- [61] B. Monsuez. Polymorphic types and widening operators. In *Static Analysis*, pages 267–281. Springer, 1993.
- [62] B. Monsuez. Polymorphic typing for call-by-name semantics. In Formal Methods in Programming and Their Applications, pages 156–169. Springer, 1993.
- [63] B. Monsuez. System f and abstract interpretation. In International Static Analysis Symposium, pages 279–295. Springer, 1995.
- [64] B. Monsuez. Using abstract interpretation to define a strictness type inference system. In Proceedings of the 1995 ACM SIGPLAN symposium on Partial evaluation and semantics-based program manipulation, pages 122–133. ACM, 1995.
- [65] C. Mossin. Higher-order value flow graphs. Nord. J. Comput., 5(3):214–234, 1998.
- [66] N. Narodytska and F. Bacchus. Maximum satisfiability using core-guided maxsat resolution. In Twenty-Eighth AAAI Conference on Artificial Intelligence, 2014.
- [67] M. Neubauer and P. Thiemann. Discriminative sum types locate the source of type errors. In *ICFP*, pages 15–26. ACM Press, 2003.
- [68] F. Nielson, H. R. Nielson, and C. Hankin. Principles of program analysis. Springer, 1999.

- [69] H. R. Nielson and F. Nielson. Infinitary control flow analysis: a collecting semantics for closure analysis. In Conference Record of POPL'97: The 24th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, Papers Presented at the Symposium, Paris, France, 15-17 January 1997, pages 332–345, 1997.
- [70] OCaml. http://ocaml.org. [Online; accessed 2-February-2014].
- [71] M. Odersky, M. Sulzmann, and M. Wehr. Type inference with constrained types. TAPOS, 5(1):35–55, 1999.
- [72] Z. Pavlinovic, T. King, and T. Wies. Finding minimum type error sources. In Proceedings of the 2014 ACM International Conference on Object Oriented Programming Systems Languages & Applications, OOPSLA 2014, part of SPLASH 2014, Portland, OR, USA, October 20-24, 2014, pages 525–542, 2014.
- [73] Z. Pavlinovic, T. King, and T. Wies. Practical smt-based type error localization. In Proceedings of the 20th ACM SIGPLAN International Conference on Functional Programming, ICFP 2015, Vancouver, BC, Canada, September 1-3, 2015, pages 412–423, 2015.
- [74] B. C. Pierce. Types and programming languages. MIT press, 2002.
- [75] R. Piskac, T. Wies, and D. Zufferey. Grasshopper. In Tools and Algorithms for the Construction and Analysis of Systems, pages 124–139. Springer, 2014.
- [76] J. Plevyak and A. A. Chien. Iterative flow analysis, 1995.

- [77] V. Rahli, J. B. Wells, J. Pirie, and F. Kamareddine. Skalpel: A type error slicer for standard ML. *Electr. Notes Theor. Comput. Sci.*, 312:197–213, 2015.
- [78] P. M. Rondon, M. Kawaguchi, and R. Jhala. Liquid types. In Proceedings of the ACM SIGPLAN 2008 Conference on Programming Language Design and Implementation, Tucson, AZ, USA, June 7-13, 2008, pages 159–169, 2008.
- [79] P. M. Rondon, M. Kawaguchi, and R. Jhala. Low-level liquid types. In Proceedings of the 37th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL 2010, Madrid, Spain, January 17-23, 2010, pages 131–144, 2010.
- [80] D. S. Scott. Data types as lattices. SIAM J. Comput., 5(3):522–587, 1976.
- [81] E. L. Seidel, H. Sibghat, K. Chaudhuri, W. Weimer, and R. Jhala. Learning to blame: localizing novice type errors with data-driven diagnosis. *PACMPL*, 1(OOPSLA):60:1–60:27, 2017.
- [82] SHErrLoc. http://www.cs.cornell.edu/projects/sherrloc/. [Online; accessed 22-April-2015].
- [83] G. Singh, M. Püschel, and M. T. Vechev. Fast polyhedra abstract domain. In Proceedings of the 44th ACM SIGPLAN Symposium on Principles of Programming Languages, POPL 2017, Paris, France, January 18-20, 2017, pages 46–59, 2017.
- [84] P. J. Stuckey, M. Sulzmann, and J. Wazny. Interactive type debugging in haskell. In *Haskell*, pages 72–83. ACM, 2003.

- [85] P. J. Stuckey, M. Sulzmann, and J. Wazny. Improving type error diagnosis. In ACM SIGPLAN Workshop on Haskell, pages 80–91. ACM, 2004.
- [86] M. Sulzmann. An overview of the chameleon system. In The Third Asian Workshop on Programming Languages and Systems, APLAS'02, Shanghai Jiao Tong University, Shanghai, China, November 29 - December 1, 2002, Proceedings, pages 16–30, 2002.
- [87] M. Sulzmann, M. Müller, and C. Zenger. Hindley/Milner style type systems in constraint form. Res. Rep. ACRC-99-009, University of South Australia, School of Computer and Information Science. July, 1999.
- [88] A. Tarski. A lattice-theoretical fixpoint theorem and its applications. Pacific Journal of Mathematics, 5, 06 1955.
- [89] F. Tip and T. B. Dinesh. A slicing-based approach for locating type errors. ACM Trans. Softw. Eng. Methodol., pages 5–55, 2001.
- [90] N. Vazou, A. Bakst, and R. Jhala. Bounded refinement types. In Proceedings of the 20th ACM SIGPLAN International Conference on Functional Programming, ICFP 2015, Vancouver, BC, Canada, September 1-3, 2015, pages 48–61, 2015.
- [91] N. Vazou, P. M. Rondon, and R. Jhala. Abstract refinement types. In Programming Languages and Systems - 22nd European Symposium on Programming, ESOP 2013, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2013, Rome, Italy, March 16-24, 2013. Proceedings, pages 209–228, 2013.

- [92] N. Vazou, E. L. Seidel, R. Jhala, D. Vytiniotis, and S. L. P. Jones. Refinement types for haskell. In Proceedings of the 19th ACM SIGPLAN international conference on Functional programming, Gothenburg, Sweden, September 1-3, 2014, pages 269–282, 2014.
- [93] P. Vekris, B. Cosman, and R. Jhala. Refinement types for typescript. In Proceedings of the 37th ACM SIGPLAN Conference on Programming Language Design and Implementation, PLDI 2016, Santa Barbara, CA, USA, June 13-17, 2016, pages 310–325, 2016.
- [94] M. Wand. Finding the source of type errors. In POPL, pages 38–43. ACM, 1986.
- [95] H. Xi and F. Pfenning. Dependent types in practical programming. In POPL '99, Proceedings of the 26th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, San Antonio, TX, USA, January 20-22, 1999, pages 214–227, 1999.
- [96] H. Xi and F. Pfenning. Dependent types in practical programming. In Proceedings of the 26th ACM SIGPLAN-SIGACT symposium on Principles of programming languages, pages 214–227. ACM, 1999.
- [97] D. Zhang and A. C. Myers. Toward general diagnosis of static errors. In POPL, pages 569–581. ACM, 2014.
- [98] D. Zhang, Y. Wang, G. E. Suh, and A. C. Myers. A hardware design language for timing-sensitive information-flow security. In *Proceedings of the Twentieth International Conference on Architectural Support for Programming Lan-*

guages and Operating Systems, ASPLOS '15, Istanbul, Turkey, March 14-18, 2015, pages 503–516, 2015.

[99] H. Zhu and S. Jagannathan. Compositional and lightweight dependent type inference for ML. In Verification, Model Checking, and Abstract Interpretation, 14th International Conference, VMCAI 2013, Rome, Italy, January 20-22, 2013. Proceedings, pages 295–314, 2013.