

An Efficient Type- and Control-Flow Analysis for System F

Connor Adsit

Rochester Institute of Technology
cda8519@rit.edu

Matthew Fluet

Rochester Institute of Technology
mtf@cs.rit.edu

Abstract

At IFL'12, we presented a novel monovariant flow analysis for System F (with recursion) that yields both *type-flow* and *control-flow* information. [4] The type-flow information approximates the type expressions that may instantiate type variables and the control-flow information approximates the λ - and Λ -expressions that may be bound to variables. Furthermore, the two flows are mutually beneficial: control flow determines which Λ -expressions may be applied to which type expressions (and, hence, which type expressions may instantiate which type variables), while type flow filters the λ - and Λ -expressions that may be bound to variables (by rejecting expressions with static types that are incompatible with the static type of the variable under the type flow).

Using a specification-based formulation of the type- and control-flow analysis, we proved the analysis to be sound, decidable, and computable. Unfortunately, naïvely implementing the analysis using a standard least fixed-point iteration yields an $O(n^{13})$ algorithm.

In this work, we give an alternative flow-graph-based formulation of the type- and control-flow analysis. We prove that the flow-graph-based formulation induces solutions satisfying the specification-based formulation and, hence, that the flow-graph-based formulation of the analysis is sound. We give a direct algorithm implementing the flow-graph-based formulation of the analysis and demonstrate that it is $O(n^4)$. By distinguishing the size l of expressions in the program from the size m of types in the program and performing an amortized complexity analysis, we further demonstrate that the algorithm is $O(l^3 + m^4)$.

1. Introduction

2. Language and Semantics

2.1 Syntax

In IFL'12, we introduced an ANF Intermediary Representation for System F. We give a specification for a modified ANF System F language in Figure 1.

We maintain the same types as before, but we restrict the form of syntactic types in order to remove the recursion found in the previous definition of types. Instead of having type constructors be composed of smaller types (ie. for function and forall types), they must

be constructed with type variables. We also introduce DeBruijn indices into the type system to describe any type parameterized by a Λ -abstraction.

In addition to adding a distinction between simple binds (values) and complex binds (applications), we expand upon the definition of expressions to accommodate the changes made to the type system. The syntax is extended to include `let`-bindings for type variables. Additionally, we mandate that all binding occurrences of expression variables be annotated with a type variable instead of a type. Similarly, the recursive function variables and λ -parameter variables appearing in simple binds must also have type variable annotations. We changed the specification of type applications so that only type variables may be passed as arguments.

Our motivation for using type variables is to promote type reuse, limiting the pool of possible types to as little a number as possible. As we will see later, having as few types as possible will reduce the overall complexity of the algorithm.

2.2 Scanning Input Program

Figure 2 introduces relations that finds all nested expressions, expression binds, type binds, type variables, and expression variables in a program.

We begin our recursive descent into a program by saying the entire program can be found inside itself. An expression can be found in the program if it is the body of a `let`-binding or if it belongs to a λ - or Λ -abstraction. An expression bind belongs to a program if the program contains a `let` x with that particular expression bind on the right-hand side. Similarly, a type bind belongs to a program if it participates in a `let` α expression that also belongs to the program. All type variables found in any nested `let` α or as parameters to Λ -abstractions belong to the program. The relation behaves similarly with expression variables and `let` x bindings and λ -abstractions. We must also consider that the recursive function variables, f , belong to a program if the encapsulating λ - or Λ -abstraction also belongs to the program.

2.3 Semantics

We assume the usual operational semantics for the ANF System using a CESK machine. More details can be found in [4].

The static semantics will need to be extended to include support for DeBruijn indices. In particular, when a value is used, we need to ensure that all indices present in the type are encapsulated by the proper amount of foralls. We also must mandate that type variables are bound before use, which has been implemented successfully in [12]. Whenever an expression variable is bound to a value, the annotated type variable must map to a type in the current context that is compatible with the type of the value.

It is worth noting that all forms of the analysis and algorithm will hold for a program even if it is untyped. It is only when the program is well typed that the flow-graph based analysis and the algorithm will yield a sound result.

Type variables	$TyVar$	\ni	α, β, \dots
Type indices	$TyIdx$	\ni	$n ::= 0 \mid 1 \mid \dots$
Type binds	$TyBnd$	\ni	$\tau ::= \alpha_a \rightarrow \alpha_b \mid \forall. \alpha_b \mid \#n$
Expression variables	$ExpVar$	\ni	x, y, z, f, g, \dots
Expression binds (simple)	$ExpBnd_s$	\ni	$b_s ::= \mu f : \alpha_f. \lambda z : \alpha_z. e_b \mid \mu f : \alpha_f. \Lambda \beta. e_b$
Expression binds (complex)	$ExpBnd_c$	\ni	$b_c ::= x_f x_a \mid x_f [\alpha_a]$
Expression binds	$ExpBnd$	\ni	$b ::= b_s \mid b_c$
Expressions	Exp	\ni	$e ::= \text{let } \alpha = \tau \text{ in } e \mid \text{let } x : \alpha_x = b \text{ in } e \mid x$
Programs	$Prog$	\ni	$P ::= e$
	$\text{ResOf}(\cdot)$	$::$	$Exp \rightarrow ExpVar$
	$\text{ResOf}(\text{let } \alpha = \tau \text{ in } e)$	$=$	$\text{ResOf}(e)$
	$\text{ResOf}(\text{let } x : \alpha_x = b \text{ in } e)$	$=$	$\text{ResOf}(e)$
	$\text{ResOf}(x)$	$=$	x
	$\text{TyOf}(\cdot)$	$::$	$ExpBnd_s \rightarrow TyVar$
	$\text{TyOf}(\mu f : \alpha_f. \lambda z : \alpha_z. e_b)$	$=$	α_f
	$\text{TyOf}(\mu f : \alpha_f. \Lambda \beta. e_b)$	$=$	α_f

Figure 1. Syntax of ANF System F

$e \preceq_{Exp} P$

$$\frac{}{P \preceq_{Exp} P} \quad \frac{\text{let } \alpha = \tau \text{ in } e \preceq_{Exp} P}{e \preceq_{Exp} P} \quad \frac{\text{let } x : \alpha_x = b \text{ in } e \preceq_{Exp} P}{e \preceq_{Exp} P}$$

$$\frac{\mu f : \alpha_f. \lambda z : \alpha_z. e_b \preceq_{ExpBnd} P}{e_b \preceq_{Exp} P} \quad \frac{\mu f : \alpha_f. \Lambda \beta. e_b \preceq_{ExpBnd} P}{e_b \preceq_{Exp} P}$$

$b \preceq_{ExpBnd} P$

$$\frac{\text{let } x : \alpha_x = b \text{ in } e \preceq_{Exp} P}{b \preceq_{ExpBnd} P}$$

$\tau \preceq_{TyBnd} P$

$$\frac{\text{let } \alpha = \tau \text{ in } e \preceq_{Exp} P}{\tau \preceq_{TyBnd} P}$$

$\alpha \preceq_{TyVar} P$

$$\frac{\text{let } \alpha = \tau \text{ in } e \preceq_{Exp} P}{\alpha \preceq_{TyVar} P} \quad \frac{\mu f : \alpha_f. \Lambda \beta. e_b \preceq_{ExpBnd} P}{\beta \preceq_{TyVar} P}$$

$x \preceq_{ExpVar} P$

$$\frac{\text{let } x : \alpha_x = b \text{ in } e \preceq_{Exp} P}{x \preceq_{ExpVar} P} \quad \frac{\mu f : \alpha_f. \lambda z : \alpha_z. e_b \preceq_{ExpBnd} P}{f \preceq_{ExpVar} P} \quad \frac{\mu f : \alpha_f. \lambda z : \alpha_z. e_b \preceq_{ExpBnd} P}{z \preceq_{ExpVar} P} \quad \frac{\mu f : \alpha_f. \Lambda \beta. e_b \preceq_{ExpBnd} P}{f \preceq_{ExpVar} P}$$

Figure 2. Sub-term Relations

3. Specification-Based Formulation of TCFA

The specification formulation is safe under an abstract type environment, which keeps track of type variables and the possibly many types that may be bound to the variables, and also an abstract value environment, which behaves the same way with expression variables and binders. In order to check for type compatibility, we make use of an abstract type environment to recursively replace all type variables with any type found in the environment’s entry for the variable. We say that a type is closed when there are no longer any type variables in the expanded type. Two type variables, α_1 and α_2 , are compatible under an abstract type environment if it is possible to derive the same closed type from both α_1 and α_2 .

An abstract type environment, $\hat{\phi}$, and an abstract value environment, $\hat{\rho}$, safely approximate an expression by an inductive analysis of the expression.

If the expression is a `let- α` expression, we must make sure that the τ being bound in the syntax appears in the possible bindings for α as described by $\hat{\phi}$ and that the remaining expressions are also safe under the same $\hat{\phi}$ and $\hat{\rho}$.

If the expression is a `let- x` expression and the binder is a simple bind, we read the type off of x and the binder. If the two are compatible we need to make sure that the binder appears in the possible bindings of x in $\hat{\rho}$. Additionally, we need to push the analysis through the body of the binder and the remaining expressions.

Otherwise, if a complex bind appears on the right side of a `let- x` expression, we need to iterate through the possible λ -abstractions recorded in $\hat{\rho}$ if the binder is a value application. For every possible argument in the entry for x_a , if the type of the argument is compatible with the type of the formal parameter under $\hat{\phi}$, we need to assert that the argument also appears in the entry for the parameter. Likewise, for all values returned by the abstraction, if the type of the value is compatible with the type of the bound variable, it must appear in the entry for the bound variable in $\hat{\phi}$.

Similarly, if a type application is being performed, we iterate through all possible Λ -abstractions that could be bound to the function variable. We assert that all types potentially bound to the type argument must appear in the description of the type parameter in $\hat{\phi}$. As with value application, any type-compatible results returned from the abstraction must be present in $\hat{\phi}$ ’s entry for the bound variable.

Finally, in the case that the expression is a simple variable, we assume that the analysis is sound.

3.1 Soundness, Decidability, and Computability

Our previous work [4, 5] showed that the specification-based formulation of the type- and control-flow analysis is sound with respect to the operational semantics, that the acceptability of given (finite) abstract type and value environments with respect to a program is decidable, and that the minimum acceptable abstract type and value environments for a program are computable in polynomial time. We briefly recall the essence of these arguments.

Soundness of the specification-based formulation of the type- and control-flow analysis asserts that any acceptable pair of abstract environments for a well-typed program approximates the run-time behavior of the program. In particular, the abstract type and value environments approximate every concrete type and value environment that arises during execution of the program. Flow soundness relies crucially on the well-typedness of the program. Soundness of the type system guarantees that, at run time, an expression variable will only be bound to a well-typed closed value of a closed type and that the expression variable’s type annotation must be interpreted as that closed type. Hence, if there is no closed type at which both the static type of the expression variable and the static

type of the value might be instantiated, then that variable will never be bound to that value at run time. The critical component of the proof is that the type compatibility judgment $\hat{\phi}; \hat{\rho} \models_S \alpha_1 \approx \alpha_2$ is derivable whenever there is a common closed type at which both α_1 and α_2 are instantiated.

Although there are an infinite number of pairs of abstract type and value environments that are acceptable for a given program, we are primarily interested more precise pairs over less precise pairs. For a given program, we can limit our attention to the “finite” abstract type and value environments that map the type variables that occur in the program to sets of type binds that appear in the program (and map all type variables that do not occur in the program to the empty set) and map the expression variables that occur in the program to sets of simple expression binds that appear in the program (and map all expression variables that do not occur in the program to the empty set).

The decidability of the acceptability judgment $\hat{\phi}; \hat{\rho} \models_S e$ relies upon the decidability of the type compatibility judgment $\hat{\phi} \models_S \alpha_1 \approx \alpha_2$. Due to “recursion” in the abstract type environment, whereby a type variable may be mapped (directly or indirectly) to a set of type binds in which the type variable itself occurs free, we cannot simply enumerate the (potentially infinite sets of) closed types θ_1 and θ_2 such that $\hat{\phi} \models_S \alpha_1 \Rightarrow \theta_1$ and $\hat{\phi} \models_S \alpha_2 \Rightarrow \theta_2$ in order to decide whether or not the judgment $\hat{\phi} \models_S \alpha_1 \approx \alpha_2$ is derivable. To address this issue, we take inspiration from the theory and implementation of regular-tree grammars [1, 3, 6]. By interpreting an abstract type environment as (the productions for) a regular-tree grammar, a derivation of the judgment $\hat{\phi} \models_S \alpha \Rightarrow \theta$ is exactly a parse tree witnessing the derivation of the ground tree θ from the starting non-terminal α in the regular-tree grammar $\hat{\phi}$ and the judgment $\hat{\phi} \models_S \alpha_1 \approx \alpha_2$ is derivable iff languages generated by taking α_1 and α_2 , respectively, as the starting non-terminal in the regular-tree grammar $\hat{\phi}$ have a non-empty intersection. Since regular-tree grammars are closed under intersection and the emptiness of a regular-tree grammar is decidable [6, 9], the type compatibility judgment $\hat{\phi} \models_S \alpha_1 \approx \alpha_2$ is decidable.

Finally, the minimum acceptable pair of abstract type and value environments for a given program is computable via a standard least-fixed point iteration. We interpret the acceptability judgment $\hat{\phi}; \hat{\rho} \models_S e$ as defining a monotone function from pairs of abstract environments to pairs of abstract environments; the “output” abstract environments are formed from the “input” abstract environments joined with all discovered violations.

We conclude with a crude upper-bound on computing the minimum acceptable pair of abstract type and value environments for a given program, of size n , via a standard least-fixed point computation. We represent $\hat{\phi}$ and $\hat{\rho}$ as two-dimensional arrays (indexed by α/τ and x/b_s , respectively), requiring $O(n^2)$ space.¹ Thus, the two abstract environments are lattices of height $O(n^2)$. Each (naïve) iteration of the monotone function is syntax directed ($O(n)$) and dominated by the function-application bind, which loops over all of the elements of $\hat{\rho}(x_f)$ and $\hat{\rho}(\text{ResOf}(e_b))$ ($O(n)$), loops over all of the elements of $\hat{\rho}(x_a)$ ($O(n)$), and computes type compatibility via a regular-tree grammar intersection ($O((n^2)^2)$), because the output regular-tree grammar is, worst-case, quadratic space with respect to the size of the input regular-tree grammar) and emptiness test ($O(((n^2)^2)^2)$), because the emptiness query is quadratic time with respect to the input regular-tree grammar). Hence, our analysis is computable in polynomial time: $O(n^{13}) = (O(n^2) + O(n^2)) \times (O(n) \times O(n) \times O(n) \times (O(n^4) + O(n^8)))$.

¹ See Sections 5.2.1 and 5.2.2 for more discussion of assumptions about the representation of the input program and data structures and operations.

Types (closed) $TyClsd \ni \theta ::= \theta_a \rightarrow \theta_b \mid \forall. \theta_b \mid \#n$
Abstract type environments $ATyEnv = TyVar \rightarrow \mathcal{P}(TyBnd) \ni \hat{\phi} ::= \{\alpha \mapsto \{\tau, \dots\}, \dots\}$
Abstract value environments $AValEnv = ExpVar \rightarrow \mathcal{P}(ExpBnds) \ni \hat{\rho} ::= \{x \mapsto \{b_s, \dots\}, \dots\}$

$$\hat{\phi} \models_S \tau \Rightarrow \theta$$

$$\frac{\hat{\phi} \models_S \alpha_a \Rightarrow \theta_a \quad \hat{\phi} \models_S \alpha_b \Rightarrow \theta_b}{\hat{\phi} \models_S \alpha_a \rightarrow \alpha_b \Rightarrow \theta_a \rightarrow \theta_b} \quad \frac{\hat{\phi} \models_S \alpha_b \Rightarrow \theta_b}{\hat{\phi} \models_S \forall. \alpha_b \Rightarrow \forall. \theta_b} \quad \frac{}{\hat{\phi} \models_S \#n \Rightarrow \#n}$$

$$\hat{\phi} \models_S \alpha \Rightarrow \theta$$

$$\frac{\tau \in \hat{\phi}(\alpha) \quad \hat{\phi} \models_S \tau \Rightarrow \theta}{\hat{\phi} \models_S \alpha \Rightarrow \theta}$$

$$\hat{\phi} \models_S \alpha_1 \approx \alpha_2$$

$$\frac{\hat{\phi} \models_S \alpha_1 \Rightarrow \theta \quad \hat{\phi} \models_S \alpha_2 \Rightarrow \theta}{\hat{\phi} \models_S \alpha_1 \approx \alpha_2}$$

$$\hat{\phi}; \hat{\rho} \models_S e$$

$$\frac{\tau \in \hat{\phi}(\alpha) \quad \hat{\phi}; \hat{\rho} \models_S e}{\hat{\phi}; \hat{\rho} \models_S \text{let } \alpha = \tau \text{ in } e}$$

$$\frac{\hat{\phi} \models_S \alpha_f \approx \alpha_x \Rightarrow \mu f: \alpha_f. \lambda z: \alpha_z. e_b \in \hat{\rho}(x) \quad \hat{\phi} \models_S \alpha_f \approx \alpha_f \Rightarrow \mu f: \alpha_f. \lambda z: \alpha_z. e_b \in \hat{\rho}(f) \quad \hat{\phi}; \hat{\rho} \models_S e_b \quad \hat{\phi}; \hat{\rho} \models_S e}{\hat{\phi}; \hat{\rho} \models_S \text{let } x: \alpha_x = \mu f: \alpha_f. \lambda z: \alpha_z. e_b \text{ in } e}$$

$$\frac{\hat{\phi} \models_S \alpha_f \approx \alpha_x \Rightarrow \mu f: \alpha_f. \Lambda \beta. e_b \in \hat{\rho}(x) \quad \hat{\phi} \models_S \alpha_f \approx \alpha_f \Rightarrow \mu f: \alpha_f. \Lambda \beta. e_b \in \hat{\rho}(f) \quad \hat{\phi}; \hat{\rho} \models_S e_b \quad \hat{\phi}; \hat{\rho} \models_S e}{\hat{\phi}; \hat{\rho} \models_S \text{let } x: \alpha_x = \mu f: \alpha_f. \Lambda \beta. e_b \text{ in } e}$$

$$\frac{\forall \mu f: \alpha_f. \lambda z: \alpha_z. e_b \in \hat{\rho}(x_f) \cdot \left(\begin{array}{l} \forall b_s \in \hat{\rho}(x_a) . \hat{\phi} \models_S \text{TyOf}(b_s) \approx \alpha_z \Rightarrow b_s \in \hat{\rho}(z) \wedge \\ \forall b_s \in \hat{\rho}(\text{ResOf}(e_b)) . \hat{\phi} \models_S \text{TyOf}(b_s) \approx \alpha_x \Rightarrow b_s \in \hat{\rho}(x) \end{array} \right) \quad \hat{\phi}; \hat{\rho} \models_S e}{\hat{\phi}; \hat{\rho} \models_S \text{let } x: \alpha_x = x_f \ x_a \text{ in } e}$$

$$\frac{\forall \mu f: \alpha_f. \Lambda \beta. e_b \in \hat{\rho}(x_f) \cdot \left(\begin{array}{l} \forall \tau \in \hat{\phi}(\alpha_a) . \tau \in \hat{\phi}(\beta) \wedge \\ \forall b_s \in \hat{\rho}(\text{ResOf}(e_b)) . \hat{\phi} \models_S \text{TyOf}(b_s) \approx \alpha_x \Rightarrow b_s \in \hat{\rho}(x) \end{array} \right) \quad \hat{\phi}; \hat{\rho} \models_S e}{\hat{\phi}; \hat{\rho} \models_S \text{let } x: \alpha_x = x_f [\alpha_a] \text{ in } e}$$

$$\frac{}{\hat{\phi}; \hat{\rho} \models_S x}$$

Figure 3. Specification-Based Formulation of TCFA

4. Flow-Graph-Based Formulation of TCFA

We give judgments pertaining to the behavior of the Flow-Graph analysis in Figure 4. These judgments are analogous to building a flow-graph of a program P where the edges are the definite flows between types, type variables, simple binders and expression variables. There is also a conditional edge between a binder and an expression variable that is guarded by the type compatibility of the two in P . Once it is learned that the types are indeed compatible, the edge is activated and belongs to the final result. After the graph is constructed, Typed- and Control-Flow Analysis is reduced to Graph Reachability across the flow-graph.

4.1 Flow-Graph Analysis

We define a series of judgments to perform a program parameterized flow-graph analysis, dependent upon the Sub-part relation in Figure 2.

The analysis uses a program based type compatibility, defined mutually with type variable compatibility. Any two DeBruijn indices are compatible if they are the same index number. Otherwise, for functions and forall types that have type variable components, the two types are compatible if the corresponding components are compatible under the program. Two type variables, α_1 and α_2 are compatible in a program if there is a τ_1 and τ_2 flowing to α_1 and α_2 , respectively, such that τ_1 is compatible with τ_2 in the program.

The flow-graph constructed by the analysis performed on a program, P , consists of both type-flow information and value-flow information.

We know that a type, τ , flows to a given type variable, α_1 , when one of two cases is true: there is an explicit $\text{let-}\alpha$ binding in the program involving τ and α_1 ; if not, there is some α_2 and we have learned that α_1 flows to α_2 . Such an edge between type variables is constructed if there exists a type application in P where α_2 is the formal type parameter and α_1 is the supplied type argument.

For a simple binder, b_s to flow to an expression variable, x , we need to know two pieces of information. We must have already seen that the binder could possibly flow to the variable and also that the type of the binder is compatible with the type of the variable. Initially, for all λ - and Λ -abstractions, we assert that the abstraction flows to its own recursive function variable. Other conditional flow edges are added whenever the program contains a $\text{let-}x$ expression using b_s and x and whenever we learn of a transitive variable flow. Whenever we see an expression application and we already know that a λ -abstraction flows to the function variable, we add a flow edge between the argument variable and the formal parameter variable dependent upon the type of the parameter and also between the return of the function and the variable being bound by the $\text{let-}x$ expression. The return of Λ -abstraction also flows to a $\text{let-}x$ bound variable dependent upon the bound variable's type if the binder is a type application and the Λ -abstraction flows the function variable.

4.2 Soundness

From our flow-based analysis, we prove that if we have a $\hat{\phi}$ and $\hat{\rho}$ that are "flow-induced" from a well-typed program, then they soundly model the program. Before we begin, we introduce the following lemma:

Lemma 1. *For all $P, \hat{\phi}$, if*

$$\forall \alpha, \tau. \tau \in \hat{\phi}(\alpha) \Leftrightarrow P \Vdash_G \tau \mapsto \alpha$$
then $P \Vdash_G \alpha_1 \approx \alpha_2 \Leftrightarrow \hat{\phi} \Vdash_S \alpha_1 \approx \alpha_2$

We assert that the lemma is true without an accompanying proof. By inspection, the forward case must be true because the derivation of $P \Vdash_G \alpha_1 \approx \alpha_2$ tells us that there exists a τ_1 flowing to α_1 and a τ_2 flowing to α_2 such that $P \Vdash_G \tau_1 \approx \tau_2$. We assume that this can be translated to $\hat{\phi} \Vdash_S \tau_1 \Rightarrow \theta$ and $\hat{\phi} \Vdash_S \tau_2 \Rightarrow \theta$,

thus allowing us to derive our conclusion. The backwards case is a reflection of the logic in the forward case.

Theorem 1. *For all $P, \hat{\phi}$, and $\hat{\rho}$, if*

- $\forall \alpha, \tau. \tau \in \hat{\phi}(\alpha) \Leftrightarrow P \Vdash_G \tau \mapsto \alpha$, and
- $\forall x, b_s. b_s \in \hat{\rho}(x) \Leftrightarrow P \Vdash_G b_s \mapsto x$

then $\hat{\phi}; \hat{\rho} \Vdash_S P$.

Proof. By induction on P . The interesting cases are when P is a $\text{let-}x$ bound to a λ -abstraction and also when P is a $\text{let-}x$ bound to a type application.

Case P of $\text{let } x:\alpha_x = \mu f:\alpha_f. \lambda z:\alpha_z. e_b \text{ in } e$: Knowing that P is unconditionally a sub-term of itself, we can deduce that $\text{let } x:\alpha_x = \mu f:\alpha_f. \lambda z:\alpha_z. e_b \text{ in } e$ is a subterm of P . Thus we can build a conditional edge in the flow graph from the abstraction to x . Since P is well-typed, we know that $P \Vdash_G \alpha_f \approx \alpha_x$ and can thus derive $P \Vdash_G \mu f:\alpha_f. \lambda z:\alpha_z. e_b \mapsto x$ and $\hat{\phi} \Vdash_S \alpha_f \approx \alpha_x$. Our assumption thus gives us $b_s \in \hat{\rho}(x)$ from $P \Vdash_G \mu f:\alpha_f. \lambda z:\alpha_z. e_b \mapsto x$. We have $\hat{\phi}; \hat{\rho} \Vdash_S e_b$ and $\hat{\phi}; \hat{\rho} \Vdash_S e$ by our inductive hypothesis, and thus we have our derivation for $\hat{\phi}; \hat{\rho} \Vdash_S \text{let } x:\alpha_x = \mu f:\alpha_f. \lambda z:\alpha_z. e_b \text{ in } e$.

Case P of $\text{let } x:\alpha_x = x_f [\alpha_a] \text{ in } e$: We again start by asserting that $\alpha_x = x_f [\alpha_a] \text{ in } e$ is a subterm of P . If there is a $\mu f:\alpha_f. \Lambda \beta. e_b$ that flows to x_f , then we can create a conditional edge from the result of the Λ -abstraction to x and an unconditional edge from α_a and β . The conditional edge only allows binders that are type compatible in the flow specification to flow transitively from the result variable to x . By our Lemma we can show that the binders that flow to x are type compatible with α_x and our assumption tells us that those binders are in the entry for x in $\hat{\rho}$. The graph also tells us that any τ flowing to α_a flows to β , which we can use to show that τ is in the entry for β in $\hat{\phi}$ from our assumption. Our inductive hypothesis allows us to show that $\hat{\phi}$ and $\hat{\rho}$ soundly model e , and we thus derive that $\hat{\phi}; \hat{\rho} \Vdash_G \text{let } x:\alpha_x = x_f [\alpha_a] \text{ in } e$. \square

$$P \vDash_G \tau_1 \approx \tau_2$$

$$\frac{\text{TYCOMPATARROW} \quad P \vDash_G \alpha_{a1} \approx \alpha_{a2} \quad P \vDash_G \alpha_{b1} \approx \alpha_{b2}}{P \vDash_G \alpha_{a1} \rightarrow \alpha_{b1} \approx \alpha_{a2} \rightarrow \alpha_{b2}}$$

$$\frac{\text{TYCOMPATFORALL} \quad P \vDash_G \alpha_{b1} \approx \alpha_{b2}}{P \vDash_G \forall. \alpha_{b1} \approx \forall. \alpha_{b2}}$$

$$\frac{\text{TYCOMPATTYIDX}}{P \vDash_G \#n \approx \#n}$$

$$P \vDash_G \alpha_1 \approx \alpha_2$$

$$\frac{\text{TYVARCOMPAT} \quad P \vDash_G \tau_1 \rightarrow \alpha_1 \quad P \vDash_G \tau_2 \rightarrow \alpha_2 \quad P \vDash_G \tau_1 \approx \tau_2}{P \vDash_G \alpha_1 \approx \alpha_2}$$

$$P \vDash_G \tau \rightarrow \alpha$$

$$\frac{\text{LETTYBND} \quad \text{let } \alpha = \tau \text{ in } e \preceq_{Exp} P}{P \vDash_G \tau \rightarrow \alpha}$$

$$\frac{\text{TRANSTYBND} \quad P \vDash_G \tau \rightarrow \alpha \quad P \vDash_G \alpha \rightarrow \beta}{P \vDash_G \tau \rightarrow \beta}$$

$$P \vDash_G \alpha \rightarrow \beta$$

$$\frac{\text{TYAPPARG} \quad P \vDash_G \mu f : \alpha_f . \Lambda \beta . e_b \rightarrow x \quad \text{let } x_r : \alpha_r = x [\alpha_a] \text{ in } e \preceq_{Exp} P}{P \vDash_G \alpha_a \rightarrow \beta}$$

$$P \vDash_G b_s \rightarrow x$$

$$\frac{\text{TYVARCOMPATEXPBND}_s \quad P \vDash_G b_s \rightarrow^? x : \alpha_x \quad P \vdash \text{TyOf}(b_s) \approx \alpha_x}{P \vDash_G b_s \rightarrow x}$$

$$P \vDash_G b_s \rightarrow^? x : \alpha_x$$

$$\frac{\text{LETEXPBND}_s \quad \text{let } x : \alpha_x = b_s \text{ in } e \preceq_{Exp} P}{P \vDash_G b_s \rightarrow^? x : \alpha_x}$$

$$\frac{\text{TRANSEXPBND}_s \quad P \vDash_G b_s \rightarrow x \quad P \vDash_G x \rightarrow y : \alpha_y}{P \vDash_G b_s \rightarrow^? y : \alpha_y}$$

$$\frac{\mu \lambda \text{EXPBND}_s \quad \text{let } x : \alpha_x = \mu f : \alpha_f . \lambda z : \alpha_z . e_b \text{ in } e \preceq_{Exp} P}{P \vDash_G \mu f : \alpha_f . \lambda z : \alpha_z . e_b \rightarrow^? f : \alpha_f}$$

$$\frac{\mu \Lambda \text{EXPBND}_s \quad \text{let } x : \alpha_x = \mu f : \alpha_f . \Lambda \beta . e_b \text{ in } e \preceq_{Exp} P}{P \vDash_G \mu f : \alpha_f . \Lambda \beta . e_b \rightarrow^? f : \alpha_f}$$

$$P \vDash_G x \rightarrow y : \alpha_y$$

$$\frac{\text{EXPAPPARG} \quad P \vDash_G \mu f : \alpha_f . \lambda z : \alpha_z . e_b \rightarrow x \quad \text{let } x_r : \alpha_r = x x_a \text{ in } e \preceq_{Exp} P}{P \vDash_G x_a \rightarrow z : \alpha_z}$$

$$\frac{\text{EXPAPPRES} \quad P \vDash_G \mu f : \alpha_f . \lambda z : \alpha_z . e_b \rightarrow x \quad \text{let } x_r : \alpha_r = x f x_a \text{ in } e \preceq_{Exp} P}{P \vDash_G \text{ResOf}(e_b) \rightarrow x_r : \alpha_r}$$

$$\frac{\text{TYAPPRES} \quad P \vDash_G \mu f : \alpha_f . \Lambda \beta . e_b \rightarrow x \quad \text{let } x_r : \alpha_r = x [\alpha_a] \text{ in } e \preceq_{Exp} P}{P \vDash_G \text{ResOf}(e_b) \rightarrow x_r : \alpha_r}$$

Figure 4. Flow-Graph-Based Formulation of TCFA

5. Algorithm

In Figures 5, 6, and 7, we give a direct algorithm implementing the flow-graph-based formulation of the type- and control-flow analysis. The algorithm returns a result set R whose elements correspond to judgements from Figure 4 that are proven to be derivable with respect to the input program P . After an initialization phase, the algorithm uses a work-queue W to process each element that is added to R ; when a newly added element is processed, all of the inference rules for which the newly added element could be an antecedent are inspected to determine if the corresponding conclusion can now be added to R . In order to achieve our desired time complexity, there is a map T from elements of the form $\alpha_1 \approx \alpha_2$ to a queue of conclusions that may be added to R when $\alpha_1 \approx \alpha_2$ is proved to be derivable; the queues in T will serve as “banks” holding credit for the amortized complexity analysis.

5.1 Commentary

5.1.1 Initialization Phase

The first initialization phase (lines 5–18) adds to R and W all elements of the form $b_s \mapsto^? x : \alpha_x$ that are derivable using the rules LETEXPBND_s , μEXPBND_s , and $\mu\Delta\text{EXPBND}_s$, rules whose conclusion follows directly from the input program. Similarly, the second initialization phase (lines 19–22) adds to R and W all elements of the form $\tau \mapsto \alpha$ that are derivable using the rule LETTYBND_s .

The third initialization phase (lines 23–27) prepares the map T , creating an empty queue for each pair of type variables that appear in the input program.

The fourth initialization phase (lines 28–48) handles the rules TYCOMPATARROW , TYCOMPATFORALL , and TYCOMPATTYIDX for all type binds that appear in the input program. When τ_1 and τ_2 are arrow types, then $\tau_1 \approx \tau_2$ is derivable using the rule TYCOMPATARROW when the argument type variables are known to be compatible and the result type variables are known to be compatible. Therefore, we create a counter c initialized with the value 2 and add the element $\langle c, \tau_1 \approx \tau_2 \rangle$ to the queues in map T for the elements $\alpha_{a1} \approx \alpha_{a2}$ and $\alpha_{b1} \approx \alpha_{b2}$. The element $\langle c, \tau_1 \approx \tau_2 \rangle$ indicates that τ_1 and τ_2 will be known to be compatible when two pairs of type variables are known to be compatible; when $\alpha_{a1} \approx \alpha_{a2}$ and $\alpha_{b1} \approx \alpha_{b2}$ are known to be compatible, the counter will be decremented and when the counter is zero, $\tau_1 \approx \tau_2$ will be added to R and W (see lines 148–156). Similarly, when τ_1 and τ_2 are forall types, then $\tau_1 \approx \tau_2$ is derivable using the rule TYCOMPATFORALL when the result type variables are known to be compatible and we create a counter c initialized with the value 1 and add the element $\langle c, \tau_1 \approx \tau_2 \rangle$ to the queue in map T for the element $\alpha_{b1} \approx \alpha_{b2}$. Finally, when τ_1 and τ_2 are the same type index, then $\tau_1 \approx \tau_2$ is immediately derivable using the rule TYCOMPATTYIDX and $\tau_1 \approx \tau_2$ is added to R and W .

5.1.2 Work-queue Phase

The work-queue phase repeatedly pops an element from the work-queue W and processes the element (possibly adding new elements to R and W) until W is empty. To process an element, all of the inference rules for which the element could be an antecedent are inspected to determine if the corresponding conclusion can now be added to R and W .

When the work-queue element is of the form $x \mapsto y : \alpha_y$ (lines 51–58), only the rule TRANSEXPBND_s need be inspected. For each $b_s \mapsto x$ that is already known to be derivable, then TRANSEXPBND_s may derive $b_s \mapsto^? y : \alpha_y$ and it is added to R and W .

When the work-queue element is of the form $b_s \mapsto x : \alpha_x$ (lines 59–68), only the rule $\text{TYVARCOMPATEXPBND}_s$ need be inspected. If $\text{TyOf}(b_s)$ and α_x are already known to be compatible, then $\text{TYVARCOMPATEXPBND}_s$ may derive $b_s \mapsto x$ and it is added

to R and W . If $\text{TyOf}(b_s)$ and α_x are not yet known to be compatible, then the element $b_s \mapsto x$ is added to the queue given by $\text{Map.get}(T, \text{TyOf}(b_s) \approx \alpha_x)$, indicating that when $\text{TyOf}(b_s)$ and α_x are known to be compatible, $b_s \mapsto x$ will be added to R and W (see lines 142–147).

When the work-queue element is of the form $b_s \mapsto x$ (lines 69–104), the rules TRANSEXPBND_s , EXPAPPARG , EXPAPPRES , TYAPPARG , and TYAPPRES need to be inspected. For each $x \mapsto y : \alpha_y$ that is already known to be derivable, then TRANSEXPBND_s may derive $b_s \mapsto^? y : \alpha_y$ and it is added to R and W . When b_s is of the form $\mu f : \alpha_f . \lambda z : \alpha_z . e_b$ where $x_b = \text{ResOf}(e_b)$ (lines 77–89), all expression applications $\text{let } x_r : \alpha_r = x \alpha_a \text{ in } e$ in the input program are examined to determine if EXPAPPARG may derive $x_a \mapsto z : \alpha_z$ and if EXPAPPRES may derive $x_b \mapsto x_r : \alpha_r$. Similarly, when b_s is of the form $\mu f : \alpha_f . \Lambda \beta . e_b$ where $x_b = \text{ResOf}(e_b)$ (lines 90–102), all expression applications $\text{let } x_r : \alpha_r = x [\alpha_a] \text{ in } e$ in the input program are examined to determine if TYAPPARG may derive $\alpha_a \mapsto \beta$ and if TYAPPRES may derive $x_b \mapsto x_r : \alpha_r$.

When the work-queue element is of the form $\tau \mapsto \alpha$ (lines 105–120), the rules TRANSTYBND and TYVARCOMPAT need to be inspected. For each $\alpha \mapsto \beta$ that is already known to be derivable, then TRANSTYBND may derive $\tau \mapsto \beta$ and it is added to R and W . For each $\pi \mapsto \beta$ that is already known to be derivable, if τ and π are already known to be compatible, then TYVARCOMPAT may derive $\alpha \approx \beta$ and it is added to R and W .

When the work-queue element is of the form $\alpha \mapsto \beta$ (lines 121–128), only the rule TRANSTYBND need be inspected. For each $\tau \mapsto \alpha$ that is already known to be derivable, then TRANSTYBND may derive $\tau \mapsto \beta$ and it is added to R and W .

When the work-queue element is of the form $\tau_1 \approx \tau_2$ (lines 129–138), only the rule TYVARCOMPAT need be inspected. For each $\tau_1 \mapsto \alpha_1$ and $\tau_2 \mapsto \alpha_2$ that are known to be derivable, then TYVARCOMPAT may derive $\alpha_1 \approx \alpha_2$ and it is added to R and W .

Finally, when the work-queue element is of the form $\alpha_1 \approx \alpha_2$ (lines 139–159), the rules $\text{TYVARCOMPATEXPBND}_s$ and TYVARCOMPAT need to be inspected. Each time that $b_s \mapsto x : \alpha_x$ was known to be derivable but $\text{TyOf}(b_s)$ and α_x were not yet known to be compatible (preventing $\text{TYVARCOMPATEXPBND}_s$ from deriving $b_s \mapsto x$), an element of the form $b_s \mapsto x$ was added to the queue given by $\text{Map.get}(T, \text{TyOf}(b_s) \approx \alpha_x)$ (see line 66); hence, processing these elements of the queue will add each $b_s \mapsto x$ that may be derived by $\text{TYVARCOMPATEXPBND}_s$ to R and W . For each pair of type binds τ_1 and τ_2 whose compatibility depends upon the compatibility of α_1 and α_2 (and possibly upon the compatibility of other pairs of type variables), an element of the form $\langle c, \tau_1 \approx \tau_2 \rangle$, where c indicates the total number of pairs of type variables whose compatibility will establish the compatibility of τ_1 and τ_2 , was added to the queue given by $\text{Map.get}(T, \alpha_1 \approx \alpha_2)$ (see lines 33, 34, and 38); hence, processing these elements of the queue will add each $\tau_1 \approx \tau_2$ that may be derived by TYVARCOMPAT to R and W .

5.1.3 Termination

Note that throughout the algorithm, whenever an element is added to the result set R , it is simultaneously added to the work-queue W . Furthermore, an element is added to R and W only after checking that the element is not already in R , except during the initialization phase when all elements added to R and W are necessarily not already in R . Hence, elements are only added to W once and the work-queue phase of the algorithm terminates because, for a given input program, there are only a finite number of elements that may be added to R and W .

Ensure: $\forall \tau_1 \preceq_{TyBnd} P. \forall \tau_2 \preceq_{TyBnd} P. \tau_1 \approx \tau_2 \in R \Leftrightarrow P \models_G \tau_1 \approx \tau_2$
Ensure: $\alpha_1 \approx \alpha_2 \in R \Leftrightarrow P \models_G \alpha_1 \approx \alpha_2$
Ensure: $\tau \mapsto \alpha \in R \Leftrightarrow P \models_G \tau \mapsto \alpha$
Ensure: $\alpha \mapsto \beta \in R \Leftrightarrow P \models_G \alpha \mapsto \beta$
Ensure: $b_s \mapsto x \in R \Leftrightarrow P \models_G b_s \mapsto x$
Ensure: $b_s \mapsto^? x : \alpha_x \in R \Leftrightarrow P \models_G b_s \mapsto^? x : \alpha_x$
Ensure: $x \mapsto y : \alpha_y \in R \Leftrightarrow P \models_G x \mapsto y : \alpha_y$

```

1: procedure TCFA( $P$ )
2:    $R \leftarrow \text{Set.newEmpty}()$ 
3:    $W \leftarrow \text{Queue.newEmpty}()$ 
4:    $T \leftarrow \text{Map.newEmpty}()$ 

5:   for all  $\text{let } x : \alpha_x = b_s \text{ in } e \preceq_{Exp} P$  do
6:      $\text{Set.insert}(R, b_s \mapsto^? x : \alpha_x)$ 
7:      $\text{Queue.push}(W, b_s \mapsto^? x : \alpha_x)$ 
8:     match  $b_s$  with
9:       case  $\mu f : \alpha_f. \lambda z : \alpha_z. e_b$  do
10:         $\text{Set.insert}(R, b_s \mapsto^? f : \alpha_f)$ 
11:         $\text{Queue.push}(W, b_s \mapsto^? f : \alpha_f)$ 
12:      end case
13:      case  $\mu f : \alpha_f. \Lambda \beta. e_b$  do
14:         $\text{Set.insert}(R, b_s \mapsto^? f : \alpha_f)$ 
15:         $\text{Queue.push}(W, b_s \mapsto^? f : \alpha_f)$ 
16:      end case
17:    end match
18:  end for

19:  for all  $\text{let } \alpha = \tau \text{ in } e \preceq_{Exp} P$  do
20:     $\text{Set.insert}(R, \tau \mapsto \alpha)$ 
21:     $\text{Queue.push}(W, \tau \mapsto \alpha)$ 
22:  end for

23:  for all  $\alpha_1 \preceq_{TyVar} P$  do
24:    for all  $\alpha_2 \preceq_{TyVar} P$  do
25:       $\text{Map.set}(T, \alpha_1 \approx \alpha_2, \text{Queue.newEmpty}())$ 
26:    end for
27:  end for

28:  for all  $\tau_1 \preceq_{TyBnd} P$  do
29:    for all  $\tau_2 \preceq_{TyBnd} P$  do
30:      match  $\langle \tau_1, \tau_2 \rangle$  with
31:        case  $\langle \alpha_{a1} \rightarrow \alpha_{b1}, \alpha_{a2} \rightarrow \alpha_{b2} \rangle$  do
32:           $c \leftarrow \text{Counter.new}(2)$ 
33:           $\text{Queue.push}(\text{Map.get}(T, \alpha_{a1} \approx \alpha_{a2}), \langle c, \tau_1 \approx \tau_2 \rangle)$ 
34:           $\text{Queue.push}(\text{Map.get}(T, \alpha_{b1} \approx \alpha_{b2}), \langle c, \tau_1 \approx \tau_2 \rangle)$ 
35:        end case
36:        case  $\langle \forall. \alpha_{b1}, \forall. \alpha_{b2} \rangle$  do
37:           $c \leftarrow \text{Counter.new}(1)$ 
38:           $\text{Queue.push}(\text{Map.get}(T, \alpha_{b1} \approx \alpha_{b2}), \langle c, \tau_1 \approx \tau_2 \rangle)$ 
39:        end case
40:        case  $\langle \#n, \#m \rangle$  do
41:          if  $n = m$  then
42:             $\text{Set.insert}(R, \tau_1 \approx \tau_2)$ 
43:             $\text{Queue.push}(W, \tau_1 \approx \tau_2)$ 
44:          end if
45:        end case
46:      end match
47:    end for
48:  end for

```

$$\begin{aligned}
\triangleright O(l^3 + m^4) &= O(l^2) + O(m^2) + O(1) + O(m^2) \\
&+ O(l) + O(m) + O(m^2) + O(m^2) \\
&+ O(l^3) + O(l^2) + O(l^3) \\
&+ O(m^4) + O(m^3) + O(m^4) + O(m^2)
\end{aligned}$$

$$\triangleright O(l^2) + O(m^2)$$

$$\triangleright O(1)$$

$$\triangleright O(m^2)$$

$$\triangleright O(l) = O(l) \times O(1)$$

$$\triangleright O(m) = O(m) \times O(1)$$

$$\begin{aligned}
\triangleright O(m^2) &= O(m) \times O(m) \\
\triangleright O(m) &= O(m) \times O(1)
\end{aligned}$$

$$\begin{aligned}
\triangleright O(m^2) &= O(m) \times O(m) \\
\triangleright O(m) &= O(m) \times O(1)
\end{aligned}$$

$$\triangleright O(1) \text{ credit into } T[\alpha_{a1} \approx \alpha_{a2}] \text{ queue}$$

$$\triangleright O(1) \text{ credit into } T[\alpha_{b1} \approx \alpha_{b2}] \text{ queue}$$

$$\triangleright O(1) \text{ credit into } T[\alpha_{b1} \approx \alpha_{b2}] \text{ queue}$$

Figure 5. TCFA Algorithm


```

49:   while ¬Queue.empty?(W) do
50:     match Queue.pop(W) with

51:       case  $x \mapsto y : \alpha_y$  do
52:         for all  $b_s \mapsto x \in R$  do
53:           if  $b_s \mapsto^? y : \alpha_y \notin R$  then
54:             Set.insert( $R, b_s \mapsto^? y : \alpha_y$ )
55:             Queue.push( $W, b_s \mapsto^? y : \alpha_y$ )
56:           end if
57:         end for
58:       end case

59:       case  $b_s \mapsto^? x : \alpha_x$  do
60:         if  $\text{TyOf}(b_s) \approx \alpha_x \in R$  then
61:           if  $b_s \mapsto x \notin R$  then
62:             Set.insert( $R, b_s \mapsto x$ )
63:             Queue.push( $W, b_s \mapsto x$ )
64:           end if
65:         else
66:           Queue.push( $\text{Map.get}(T, \text{TyOf}(b_s) \approx \alpha_x), b_s \mapsto x$ )
67:         end if
68:       end case

69:       case  $b_s \mapsto x$  do
70:         for all  $x \mapsto y : \alpha_y \in R$  do
71:           if  $b_s \mapsto^? y : \alpha_y \notin R$  then
72:             Set.insert( $R, b_s \mapsto^? y : \alpha_y$ )
73:             Queue.push( $W, b_s \mapsto^? y : \alpha_y$ )
74:           end if
75:         end for
76:         match  $b_s$  with
77:           case  $\mu f : \alpha_f . \lambda z : \alpha_z . e_b$  do
78:              $x_b \leftarrow \text{ResOf}(e_b)$ 
79:             for all let  $x_r : \alpha_r = x x_a$  in  $e \preceq_{Exp} P$  do
80:               if  $x_a \mapsto z : \alpha_z \notin R$  then
81:                 Set.insert( $R, x_a \mapsto z : \alpha_z$ )
82:                 Queue.push( $W, x_a \mapsto z : \alpha_z$ )
83:               end if
84:               if  $x_b \mapsto x_r : \alpha_r \notin R$  then
85:                 Set.insert( $R, x_b \mapsto x_r : \alpha_r$ )
86:                 Queue.push( $W, x_b \mapsto x_r : \alpha_r$ )
87:               end if
88:             end for
89:           end case
90:           case  $\mu f : \alpha_f . \Lambda \beta . e_b$  do
91:              $x_b \leftarrow \text{ResOf}(e_b)$ 
92:             for all let  $x_r : \alpha_r = x [\alpha_a]$  in  $e \preceq_{Exp} P$  do
93:               if  $\alpha_a \mapsto \beta \notin R$  then
94:                 Set.insert( $R, \alpha_a \mapsto \beta$ )
95:                 Queue.push( $W, \alpha_a \mapsto \beta$ )
96:               end if
97:               if  $x_b \mapsto x_r : \alpha_r \notin R$  then
98:                 Set.insert( $R, x_b \mapsto x_r : \alpha_r$ )
99:                 Queue.push( $W, x_b \mapsto x_r : \alpha_r$ )
100:             end if
101:           end for
102:         end match
103:       end case
104:     end case

```

$\triangleright O(l^3) = O(l^2) \times O(l)$
 $\triangleright O(l) = O(l) \times O(1)$

$\triangleright O(l^2) = O(l^2) \times O(1)$

$\triangleright O(1)$ credit into $T[\text{TyOf}(b_s) \approx \alpha_x]$ queue

$\triangleright O(l^3) = O(l^2) \times O(l)$
 $\triangleright O(l) = O(l) \times O(1)$

$\triangleright O(l) = O(l) \times O(1)$

Figure 6. TCFA Algorithm (continued)

```

105:     case  $\tau \mapsto \alpha$  do
106:         for all  $\alpha \mapsto \beta \in R$  do
107:             if  $\tau \mapsto \beta \notin R$  then
108:                 Set.insert( $R, \tau \mapsto \beta$ )
109:                 Queue.push( $W, \tau \mapsto \beta$ )
110:             end if
111:         end for
112:         for all  $\pi \mapsto \beta \in R$  do
113:             if  $\tau \approx \pi \in R$  then
114:                 if  $\alpha \approx \beta \notin R$  then
115:                     Set.insert( $R, \alpha \approx \beta$ )
116:                     Queue.push( $W, \alpha \approx \beta$ )
117:                 end if
118:             end if
119:         end for
120:     end case

121:     case  $\alpha \mapsto \beta$  do
122:         for all  $\tau \mapsto \alpha \in R$  do
123:             if  $\tau \mapsto \beta \notin R$  then
124:                 Set.insert( $R, \tau \mapsto \beta$ )
125:                 Queue.push( $R, \tau \mapsto \beta$ )
126:             end if
127:         end for
128:     end case

129:     case  $\tau_1 \approx \tau_2$  do
130:         for all  $\tau_1 \mapsto \alpha_1 \in R$  do
131:             for all  $\tau_2 \mapsto \alpha_2 \in R$  do
132:                 if  $\alpha_1 \approx \alpha_2 \notin R$  then
133:                     Set.insert( $R, \alpha_1 \approx \alpha_2$ )
134:                     Queue.push( $W, \alpha_1 \approx \alpha_2$ )
135:                 end if
136:             end for
137:         end for
138:     end case

139:     case  $\alpha_1 \approx \alpha_2$  do
140:         while  $\neg$ Queue.empty?(Map.get( $T, \alpha_1 \approx \alpha_2$ )) do
141:             match Queue.pop(Map.get( $T, \alpha_1 \approx \alpha_2$ )) with
142:                 case  $b_s \mapsto x$  do
143:                     if  $b_s \mapsto x \notin R$  then
144:                         Set.insert( $R, b_s \mapsto x$ )
145:                         Queue.push( $W, b_s \mapsto x$ )
146:                     end if
147:                 end case
148:                 case  $\langle c, \tau_1 \approx \tau_2 \rangle$  do
149:                     Counter.dec( $c$ )
150:                     if Counter.get( $c$ ) = 0 then
151:                         if  $\tau_1 \approx \tau_2 \notin R$  then
152:                             Set.insert( $R, \tau_1 \approx \tau_2$ )
153:                             Queue.push( $W, \tau_1 \approx \tau_2$ )
154:                         end if
155:                     end if
156:                 end case
157:             end match
158:         end while
159:     end case

160: end match
161: end while

162: return  $R$ 
163: end procedure

```

$\triangleright O(m^4) = O(m^2) \times O(m^2)$
 $\triangleright O(m) = O(m) \times O(1)$

$\triangleright O(m^2) = O(m^2) \times O(1)$

$\triangleright O(m^3) = O(m^2) \times O(m)$
 $\triangleright O(m) = O(m) \times O(1)$

$\triangleright O(m^4) = O(m^2) \times O(m^2)$
 $\triangleright O(m^2) = O(m) \times O(m)$
 $\triangleright O(m) = O(m) \times O(1)$

$\triangleright O(m^2) = O(m^2) \times O(1)$
 $\triangleright O(1)$ credit from $T[\alpha_{a1} \approx \alpha_{a2}]$ queue

Figure 7. TCFA Algorithm (continued)

5.2 Complexity

5.2.1 Preliminaries

Before analyzing the time complexity of the algorithm, we first make some (standard) assumptions about the representation of the input program.

We assume that all `let`-, μ -, and λ -bound expression variables and all `let`-, and Λ -bound type variables in the program are distinct. We further assume that expression variables and type variables can be mapped (in $O(1)$ time) to unique integers (for $O(1)$ time indexing into an array) and that integers can be mapped (in $O(1)$ time) to corresponding expression variables and type variables.² Given the assumption that all `let`-, μ -, and λ -bound expression variables in the program are unique, each expression variable in the program is annotated with a single type variable at its unique binding occurrence. We therefore assume that expression variables can be mapped (in $O(1)$ time) to its annotating type variable. Given the assumption that all μ -bound expression variables in the program are unique, each simple expression bind in the program is unique and can be mapped (in $O(1)$ time) to and from unique integers.³ We do not assume that each type bind in the program is unique, but we do assume that each type bind in the program can be mapped (in $O(1)$ time) to and from unique integers. Finally, we assume that `ResOf`(\cdot) can be computed in $O(1)$ time.⁴

5.2.2 Data Structures and Operations

We next consider the data structures used to implement the result set R and the map T and the cost of various operations.

The result set R is implemented as seven multi-dimensional arrays, each corresponding to one of the seven judgements from Figure 4. Given the assumptions above, it is easy to see that the arrays corresponding to $\tau_1 \approx \tau_2$, $\alpha_1 \approx \alpha_2$, $\tau \mapsto \alpha$, $\alpha \mapsto \beta$, and $b_s \mapsto x$ are simple two-dimensional arrays with $O(1)$ time indexing by mapping components to unique integers. Furthermore, the arrays corresponding to $b_s \mapsto x : \alpha_x$ and $x \mapsto y : \alpha_y$ can also be implemented with simple two-dimensional arrays (indexed by b_s/x and x/y , respectively), because the type variable is always the single type variable at the unique binding occurrence of the expression variable and can be left implicit. Thus, queries like $b_s \mapsto x \notin R$ and operations like `Set.insert`($R, b_s \mapsto x$) can be performed in constant time. Loops like “**for all** $b_s \mapsto x \in R$ **do**” for fixed b_s instantiating x or for fixed x instantiating b_s can be implemented as a linear scan of an array column or array row.

The map T is implemented with a simple two-dimensional array, indexed by pairs of type variables. Operations like `Map.set`($T, \alpha_1 \approx \alpha_2, q$) and `Map.get`($T, \alpha_1 \approx \alpha_2$) can be performed in constant time.

The work-queue W and the queues in map T are implemented with a simple linked-list queue. Queries like `Queue.empty?`(W) and operations like `Queue.push`($W, b_s \mapsto x$) and `Queue.pop`(W) can be performed in constant time.

5.2.3 Coarse Analysis

We first argue that the algorithm is $O(n^4)$ time, where n is the size of the input program P . First, note that there are $O(n)$ type

² These mappings can be established with a linear-time preprocessing step.

³ In a richer language with simple-expression-bind forms that do not include a bound expression variable (e.g., $\langle x_1, x_2 \rangle$ pairs), we can assume a numbering of all simple expression binds in the program, similar to the labeling found in textbook presentations of CFA [10].

⁴ This can be established either by a linear-time preprocessing step (associating each result variable with its corresponding abstraction) or by changing the representation of expressions to a list of $\alpha = \tau$ and $x : \alpha_x = b$ bindings paired with the result variable.

variables, $O(n)$ type binds, $O(n)$ expression variables, and $O(n)$ simple expression binds in the program. Thus, the result set R requires $O(n^2)$ space for (and is $O(n^2)$ time to create) each of the seven two-dimensional arrays and the map T requires $O(n^2)$ space for (and is $O(n^2)$ time to create) the two-dimensional array.

The first initialization phase is $O(n)$ time to traverse the program and process each simple expression bind. Similarly, the second initialization phase is $O(n)$ time to traverse the program and process each type bind. The third initialization phase is $O(n^2)$ time to process each pair of type variables. The fourth initialization phase is $O(n^2)$ time to process each pair of type binds. Altogether, the initialization phase is $O(n^2) = O(n) + O(n) + O(n^2) + O(n^2)$ time.

As noted above, elements are only added to W once. Therefore, the time complexity of the “**while** `!Queue.empty?`(W) **do**”-loop is the sum of the time required to process an element of each kind times the number of elements of that kind. There are $O(n^2)$ elements of the form $x \mapsto y : \alpha_y$ (recall that the α_y is implicitly determined by the y) and processing an $x \mapsto y : \alpha_y$ element is $O(n)$ time to scan for all $b_s \mapsto x \in R$. There are $O(n^2)$ elements of the form $b_s \mapsto x : \alpha_x$ and processing a $b_s \mapsto x : \alpha_x$ element is $O(1)$ time. There are $O(n^2)$ elements of the form $b_s \mapsto x$ and processing a $b_s \mapsto x$ element is $O(n)$ time to scan all $x \mapsto y : \alpha_y \in R$ and $O(n)$ time to find all `let` $x_r : \alpha_r = x$ x_a **in** $e \preceq_{Exp} P$ and to find all `let` $x_r : \alpha_r = x$ $[\alpha_a]$ **in** $e \preceq_{Exp} P$. There are $O(n^2)$ elements of the form $\tau \mapsto \alpha$ and processing a $\tau \mapsto \alpha$ element is $O(n)$ time to scan for all $\alpha \mapsto \beta \in R$ and $O(n^2)$ to process all $\pi \mapsto \beta \in R$. There are $O(n^2)$ elements of the form $\alpha \mapsto \beta$ and processing an $\alpha \mapsto \beta$ element is $O(n)$ time to scan for all $\tau \mapsto \alpha \in R$. There are $O(n^2)$ elements of the form $\tau_1 \approx \tau_2$ and processing a $\tau_1 \approx \tau_2$ element is $O(n^2)$ time to scan for all $\tau_1 \mapsto \alpha_1 \in R$ and $\tau_2 \mapsto \alpha_2 \in R$. There are $O(n^2)$ elements of the form $\alpha_1 \approx \alpha_2$, processing an $\alpha_1 \approx \alpha_2$ element must process each element in the queue `Map.get`($T, \alpha_1 \approx \alpha_2$), and, therefore, the time complexity to process an $\alpha_1 \approx \alpha_2$ element is the sum of the time required to process the elements in the queue of each kind times the number of elements of that kind; there are $O(n^2)$ elements of the form $b_s \mapsto x$ in the queue (since an element of the form $b_s \mapsto x$ are added at most once to at most one queue (see line 66)) and processing an $b_s \mapsto x$ element is $O(1)$ time and there are $O(n^2)$ elements of the form $\langle c, \tau_1 \approx \tau_2 \rangle$ (since an element of the form $\langle c, \tau_1 \approx \tau_2 \rangle$ is added at most twice to at most one queue (see lines 33–34 and line 38)) and processing a $\langle c, \tau_1 \approx \tau_2 \rangle$ element is $O(1)$ time. Altogether, the work-queue phase is $O(n^4) = O(n^2) \times O(n) + O(n^2) \times O(1) + O(n^2) \times (O(n) + O(n) + O(n)) + O(n^2) \times (O(n) + O(n^2)) + O(n^2) \times O(n) + O(n^2) \times O(n^2) + O(n^2) \times (O(n^2) \times O(1) + O(n^2) \times O(1))$.

Thus, the entire algorithm is $O(n^4)$. Recall that algorithms for classic (untyped) control-flow analysis have been shown to be $O(n^3)$ [2, 7, 10, 11], though recently improved to $O(n^2 \log n)$ [8].

5.2.4 Refined Analysis

In order to clarify the relationship between the time complexity of algorithms for classic (untyped) control-flow analysis and our algorithm for type- and control-flow analysis, we perform a refined analysis of our algorithm.

First, note that the quartic components of the algorithm are due to the processing of elements of the form $\tau \mapsto \alpha$, $\tau_1 \approx \tau_2$, and $\alpha_1 \approx \alpha_2$. Intuitively, the increased time complexity of the algorithm for type- and control-flow analysis compared to algorithms for classic (untyped) control-flow analysis is due to the computation of the type-compatibility relations.

Second, in typical programs of interest, we expect that the total size of the program to be dominated by the contribution of (bound) expression variables and expression binds, with the contribution of (bound) type variables and type binds significantly (asymptotically?) less. For example, a program may have many definitions of and uses of $\text{int} \rightarrow \text{int}$ functions, all of which can share the same (top-level) $\text{let } \alpha_i = \text{int} \text{ in } \text{let } \alpha_{i \rightarrow i} = \alpha_i \rightarrow \alpha_i \text{ in } \dots$ type bindings. Indeed, our ANF representation of types encourages type-level optimizations such as let -floating, common subexpression elimination (CSE), and copy propagation, which would further reduce the contribution of types to the total program size. Therefore, we consider it useful to distinguish l , the size of (bound) expression variables and expression binds, and m , the size of (bound) type variables and type binds, where we have $O(l) + O(m)$ is $O(n)$ and we expect $O(l) \gg O(m)$, though, in the worst-case, both $O(l)$ and $O(m)$ are $O(n)$. We further assume an $O(n)$ preprocessing step that provides an enumeration of all $\text{let } x : \alpha_x = b \text{ in } e \preceq_{Exp} P$ in $O(l)$ time and an enumeration of all $\text{let } \alpha = \tau \text{ in } e \preceq_{Exp} P$ in $O(m)$ time.

We now argue that the algorithm is $O(l^3 + m^4)$ time. First, note that there are $O(m)$ type variables, $O(m)$ type binds, $O(l)$ expression variables, and $O(l)$ simple expression binds in the program. Thus, the result set R requires $O(l^2 + m^2)$ space for (and is $O(l^2 + m^2)$ time to create) the seven two-dimensional arrays and the map T requires $O(m^2)$ space for (and is $O(m^2)$ time to create) the two-dimensional array.

The first initialization phase is $O(l)$ time to process each simple expression bind. Similarly, the second initialization phase is $O(m)$ time to process each type bind. The third initialization phase is $O(m^2)$ time to process each pair of type variables. The fourth initialization phase is $O(m^2)$ time to process each pair of type binds; included in this processing time is an $O(1)$ credit “deposited” into the queues in T when pushing elements, which “pre-pays” for the processing of the elements when popped. Altogether, the initialization phase is $O(l + m^2) = O(l) + O(m) + O(m^2) + O(m^2)$ time.

The analysis of the work-queue phase is similar to that performed above: the time complexity of the “**while** $\neg \text{Queue.empty?}(W)$ **do**”-loop is the sum of the time required to process an element of each kind times the number of elements of that kind; we simply refine n to l or m as appropriate. We further perform an amortized analysis of the time complexity to process an $b_s \mapsto^? x : \alpha_x$ element and to process an $\alpha_1 \approx \alpha_2$ element. Included in the time to process an $b_s \mapsto^? x : \alpha_x$ element is an $O(1)$ credit “deposited” into the queue given by $\text{Map.get}(T, \text{TyOf}(b_s) \approx \alpha_x)$ when pushing elements, which “pre-pays” for the processing of the elements when popped. As before, processing an $\alpha_1 \approx \alpha_2$ element must process each element in the queue $\text{Map.get}(T, \alpha_1 \approx \alpha_2)$; however, an $O(1)$ credit may be “withdrawn” from the queue $\text{Map.get}(T, \alpha_1 \approx \alpha_2)$ when popping elements and this $O(1)$ credit may be used to “pay” for the popping and processing of the element. Thus, processing an $\alpha_1 \approx \alpha_2$ element is (amortized) $O(1)$ time.⁵ Altogether, the work-queue phase is $O(l^3 + m^4) = O(l^2) \times O(l) + O(l^2) \times O(1) + O(l^2) \times (O(l) + O(l) + O(l)) + O(m^2) \times (O(m) + O(m^2)) + O(m^2) \times O(m) + O(m^2) \times O(m^2) + O(m^2) \times O(1)$.

Thus, the entire algorithm is $O(l^3 + m^4)$.

⁵Note that without the amortized analysis, processing an $\alpha_1 \approx \alpha_2$ element would be $O(l^2) + O(m^2)$ time and the entire algorithm would be $O(l^3 + l^2m^2 + m^4)$.

6. Conclusion

References

- [1] A. Aiken and B. R. Murphy. Implementing regular tree expressions. In J. Hughes, editor, *FPCA'91: Proceedings of the Fifth ACM Conference on Functional Programming Languages and Computer Architecture*, volume 523 of *Lecture Notes in Computer Science*, pages 427–447, Cambridge, Massachusetts, Aug. 1991. Springer-Verlag.
- [2] A. E. Ayers. Efficient closure analysis with reachability. In M. Bilaud, P. Castéran, M.-M. Corsini, K. Musumbu, and A. Rauzy, editors, *Actes WSA'92 Workshop on Static Analysis*, Bigre, pages 126–134, Bordeaux, France, Sept. 1992. Atelier Irisa, IRISA, Campus de Beaulieu.
- [3] P. Cousot and R. Cousot. Formal language, grammar and set-constraint-based program analysis by abstract interpretation. In S. Peyton Jones, editor, *FPCA'95: Proceedings of the Seventh International Conference on Functional Programming Languages and Computer Architecture*, pages 170–181, La Jolla, California, June 1995.
- [4] M. Fluet. A type- and control-flow analysis for System F. In R. Hinze, editor, *IFL'12: Post-Proceedings of the 24th International Symposium on Implementation and Application of Functional Languages*, Lecture Notes in Computer Science, Oxford, England, 2013. Springer-Verlag. To appear.
- [5] M. Fluet. A type- and control-flow analysis for System F. Technical report, Rochester Institute of Technology, February 2013. <https://ritdml.rit.edu/handle/1850/15920>.
- [6] F. Geceseg and M. Steinby. *Tree Automata*. Akademiai Kiado, Budapest, Hungary, 1984.
- [7] N. Heintze. Set-based program analysis of ML programs. In C. L. Talcott, editor, *LFP'94: Proceedings of the 1994 ACM Conference on Lisp and Functional Programming*, pages 306–317, Orlando, Florida, June 1994.
- [8] J. Midtgaard and D. V. Horn. Subcubic control flow analysis algorithms. Computer Science Research Report 125, Roskilde University, Roskilde, Denmark, May 2009. Revised version to appear in *Higher-Order and Symbolic Computation*.
- [9] P. Mishra and U. S. Reddy. Declaration-free type checking. In M. S. Van Deusen, Z. Galil, and B. K. Reid, editors, *POPL'85: Proceedings of the Twelfth Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, pages 7–21, New Orleans, Louisiana, Jan. 1985. ACM, ACM.
- [10] F. Nielson, H. R. Nielson, and C. Hankin. *Principles of Program Analysis*. Springer-Verlag, 1999.
- [11] J. Palsberg and M. I. Schwartzbach. Safety analysis versus type inference. *Information and Computation*, 118(1):128–141, 1995.
- [12] C. A. Stone. Type definitions. In B. C. Pierce, editor, *Advanced Topics in Types and Programming Languages*. The MIT Press, 2005.