, , 1–48 ()
© Kluwer Academic Publishers, Boston. Manufactured in The Netherlands.

Optimal Representations of Polymorphic Types with Subtyping*

ALEXANDER AIKEN aiken@cs.berkeley.edu
EECS Department, University of California, Berkeley, Berkeley, CA 94720-1776

ED WIMMERS wimmers@almaden.ibm.com
IBM Almaden Research Center, 650 Harry Rd., San Jose, CA 95120-6099

JENS PALSBERG palsberg@cs.purdue.edu
Department of Computer Science, Purdue University, West lafayette, IN 47907

Editor: Carolyn Talcott

Abstract. Many type inference and program analysis systems include notions of subtyping and parametric polymorphism. When used together, these two features induce equivalences that allow types to be simplified by eliminating quantified variables. Eliminating variables both improves the readability of types and the performance of algorithms whose complexity depends on the number of type variables. We present an algorithm for simplifying quantified types in the presence of subtyping and prove it is sound and complete for non-recursive and recursive types. We also show that an extension of the algorithm is sound but not complete for a type language with intersection and union types, as well as for a language of constrained types.

Keywords: types, polymorphism, subtyping

1. Introduction

Contemporary type systems include a wide array of features, of which two of the most important are *subtyping* and *parametric polymorphism*. These two features are independently useful. Subtyping expresses relationships between types of the form "type τ_1 is less than type τ_2". Such relationships are useful, for example, in object-oriented type systems and in program analysis algorithms where a greatest (or least) element is required. Parametric polymorphism allows a parameterized type inferred for a program fragment to take on a different instance in every context where it is used. This feature has the advantage that the same program can be used at many different types.

A number of type systems have been proposed that combine subtyping and polymorphism, among other features. The intended purposes of these systems varies. A few examples are: studies of type systems themselves [5, 7, 2], proposals for type systems for object-oriented languages [8], and program analysis systems used in program optimization [3, 10]. In short, the combination of subtyping and polymorphism is useful, with a wide range of applications.

When taken together, subtyping and polymorphism induce equivalences on types that can be exploited to simplify the representation of types. Our main technical

* This work was supported by an NSF NYI award, CCR-9457812

result is that, in a simple type language with a least type $\bot$ and greatest type $\top$, for any type σ there is another type σ' that is equivalent to σ and σ' has the minimum number of quantified type variables. Thus, type simplification eliminates quantified variables wherever possible. Eliminating variables is desirable for three reasons. First, many type inference algorithms have computational complexity that is sensitive (both theoretically and practically) to the number of type variables. Second, eliminating variables makes types more readable. Third, simplification makes properties of types manifest that are otherwise implicit; in at least one case that we know of, these "hidden" properties are exactly the information needed to justify compiler optimizations based on type information [3].

The basic idea behind variable elimination is best illustrated with an example. A few definitions are needed first. Consider a simple type language defined by the following grammar:

$$\tau ::= \alpha \mid \top \mid \bot \mid \tau_1 \to \tau_2$$

In this grammar, α is a type variable. Following standard practice, we use $\alpha, \beta, \ldots$ for type variables and $\tau, \tau', \tau_1, \tau_2, \ldots$ for types. The subtyping relation is a partial order $\preceq$ on types, which is the least relation satisfying

$$\begin{aligned} \tau &\preceq \tau \\ \bot &\preceq \tau \\ \tau &\preceq \top \\ \tau_1 \preceq \tau_1' \wedge \tau_2 \preceq \tau_2' &\Leftrightarrow \tau_1' \to \tau_2 \preceq \tau_1 \to \tau_2' \end{aligned}$$

Quantified types are given by the grammar:

$$\sigma ::= \tau \mid \forall \alpha.\sigma$$

For the moment, we rely on the reader's intuition about the meaning of quantified types. A formal semantics of quantified types is presented in Section 2.

Consider the type $\forall\alpha.\forall\beta.\alpha \to \beta$. Any function with this type takes an input of an arbitrary type α and produces an output of any (possibly distinct) arbitrary type β. What functions have this type? The output β must be included in all possible types; there is only one such type $\bot$. The input α, however, must include all possible types; there is only one such type $\top$. Thus, one might suspect that this type is equivalent to $\top \to \bot$. The only function with this type is the one that diverges for all possible inputs.

It turns out that, in fact, $\forall\alpha.\forall\beta.\alpha \to \beta \equiv \top \to \bot$ in the standard *ideal* model of types [13]. As argued above, the type with fewer variables is better for human readability, the speed of type inference, and for the automatic exploitation of type information by a compiler. We briefly illustrate these three claims.

The reasoning required to discover that $\forall\alpha.\forall\beta.\alpha \to \beta$ represents an everywhere-divergent function is non-trivial. There is a published account illustrating how types inferred from ML programs (which have polymorphism but no subtyping) can be used to detect non-terminating functions exactly as above [12]. The previous example is the simplest one possible; the problem of understanding types only

$$\forall\alpha_1.\ldots.\forall\alpha_8.\,\alpha_6 \;/\; \left\{\begin{array}{rcl}
\alpha_4 \to \alpha_6 & \preceq & \alpha_1 \preceq \alpha_5 \to \alpha_6 \\
\alpha_1 & \preceq & \alpha_2 \preceq \alpha_3 \\
\alpha_1 & \preceq & \alpha_2 \preceq \alpha_5 \to \alpha_6 \\
\bot & \preceq & \alpha_3 \preceq \top \\
\alpha_2 & \preceq & \alpha_4 \preceq \alpha_5 \to \alpha_6 \\
\alpha_4 & \preceq & \alpha_5 \preceq \alpha_4 \\
\bot & \preceq & \alpha_6 \preceq \alpha_3 \\
\alpha_4 \to \alpha_6 & \preceq & \alpha_7 \preceq \alpha_1 \\
\alpha_3 & \preceq & \alpha_8 \preceq \top
\end{array}\right.$$

Figure 1. A quantified type of eight variables qualified by constraints.

increases with the size of the type and expressiveness of the type language. The following example is taken from the system of [2], a subtype inference system with polymorphism. In typing a term, the inference algorithm in this system generates a system of subtyping constraints that must be satisfied. The solution of the constraints gives the desired type. Constraints are generated as follows: If f has type $\alpha \to \beta$ and x has type γ, then for an application $f\,x$ to be well-typed it must be the case that $\gamma \preceq \alpha$. Figure 1 shows the type generated for the divergent lambda term $(\lambda x.x\,x)(\lambda x.x\,x)$. The type has the form

$$\forall\alpha_1,\ldots,\alpha_8.(\alpha_6/S)$$

Informally, the meaning of this type is α_6 for any assignment to the variables $\alpha_1,\ldots,\alpha_8$ that simultaneously satisfies all the constraints in S.

This type is equal to $\bot$, a fact proven by our algorithm extended to handle constraints. The type $\bot$ is sound, since the term is divergent. This example illustrates both improved readability and the possibility of more efficient inference. To use the polymorphic type $\forall\alpha_1,\ldots,\alpha_n.(\tau/S)$, the variables must be instantiated and the constraints duplicated for each usage context. Eliminating variables simplifies the representation, making this very expensive aspect of type inference less costly.

Finally, simplifying types can improve not only the speed but the quality of program analyses. For example, the *soft typing* system of [3] reduces the problem of identifying where runtime type checks are unneeded in a program to testing whether certain type variables can be replaced by $\bot$ in a quantified type. This is exactly the task performed by elimination of variables in quantified types.

Our main contribution is a variable elimination algorithm that is sound and complete (i.e., eliminates as many variables as possible) for the simple type language defined above, as well as for a type language with recursive types. We extend the algorithm to type languages with intersection and union types and to type languages with subsidiary constraints. For these latter two cases, the techniques we present are sound but not complete. Combining the completeness results for the simpler languages with examples illustrating the incompleteness of the algorithm in the more expressive settings, we shed some light on the sources of incompleteness.

The various algorithms are simple and efficient. Let n be the print size of the type and m be the number of variables. Then the time complexity is $\mathcal{O}(mn)$ for the cases of simple and recursive types and $\mathcal{O}(m^2n)$ for the cases of systems with intersection, union, or constrained types. The algorithm for simplifying quantified types with subsidiary constraints has been in use since 1993, but with the exception of code documentation little has been written previously on the subject. The algorithm has been implemented and used in Illyria[1], the systems reported in [2], a large scale program analysis system for the functional language FL [3], and a general-purpose constraint-based program analysis system [9]. These last two applications are by far the largest and best engineered. The quality of these systems depends on eliminating variables wherever possible.

Other recent systems based on constrained types have also pointed out the importance of variable elimination. In [8], Eifrig, Smith, and Trifonov describe a variable elimination method similar to, but not identical to, the one in Section 7. Pottier gives a method that can eliminate redundant variables from constraint sets [16]. Both of these methods are heuristic; i.e., they are sound but not complete. Constraint simplification is also a component of the systems described in [11, 17]. It is not claimed that either system performs complete constraint simplification.

Our focus in this paper is quite different. The question of variable elimination arises in any type system with polymorphism and subtyping, not just in systems with constrained types. Our purpose is to explore the basic structure of this problem in the simplest settings and to understand what makes the problem harder in the case of constrained types. To the best of our knowledge, we present the first sound and complete algorithms for variable elimination in any type system.

Rather than work in a specific semantic domain, we state axioms that a semantic domain must satisfy for our techniques to apply (Section 2). Section 3 gives the syntax for type expressions as well as their interpretation in the semantic domain.

Section 4 proves the results for the case of *simple type expressions*, which are non-recursive types. For quantified simple types, variable elimination produces an equivalent type with the minimum number of quantified variables. Furthermore, all equivalent types with the minimum number of quantified variables are α-*equivalent*—they are identical up to the names and order of quantified variables.

The intuition behind the variable elimination procedure is easy to convey. Type variables may be classed as *monotonic* (*positive*) or *anti-monotonic* (*negative*) based on their syntactic position in a type. Intuitively, the main lemma shows that quantified variables that are solely monotonic can be eliminated in favor of $\bot$; quantified variables that are solely anti-monotonic can be eliminated in favor of $\top$. Section 4.2 proves that the strategy of eliminating either monotonic or anti-monotonic variables is complete for the simple type language. Variables that are both monotonic and anti-monotonic cannot be eliminated.

Section 5 extends the basic variable elimination algorithm to a type language with recursive types. The extended algorithm is again both sound and complete, but it is no longer the case that all equivalent types with the minimum number of quantified variables are α-equivalent.

Section 6 extends the algorithm to intersection and union types. This language is the first extension for which the techniques are sound but not complete. Examples are given showing sources of incompleteness. Finally, Section 7 extends the algorithm to a type language with subsidiary constraints, as in Figure 1. This is the most general type language we consider. Section 8 concludes with a few remarks on related work.

2. Semantic Domains

Rather than work with a particular semantic domain, we axiomatize the properties needed to prove the corresponding theorems about eliminating quantified variables.

Definition 1. A semantic domain $\mathcal{D} = (\mathcal{D}_0, \mathcal{D}_1, \preceq, \sqcap)$ satisfies the following properties:

1. $\mathcal{D}_0 \subseteq \mathcal{D}_1$ or, more generally, there is monomorphism from $\mathcal{D}_0$ to $\mathcal{D}_1$.
2. a partial order on $\mathcal{D}_1$ denoted by $\preceq$.
3. a minimal element $\bot \in \mathcal{D}_0$ such that $\bot \preceq x$ for all $x \in \mathcal{D}_1$.
4. a maximal element $\top \in \mathcal{D}_0$ such that $x \preceq \top$ for all $x \in \mathcal{D}_1$.
5. a binary operation $\rightarrow$ on $\mathcal{D}_0$ such that if $y_1 \preceq x_1$ and $x_2 \preceq y_2$, then $x_1 \rightarrow x_2 \preceq y_1 \rightarrow y_2$.
   Furthermore, $\bot \rightarrow \top \neq \top$ and $\top \rightarrow \bot \neq \bot$.
6. a greatest lower bound operation $\sqcap$ on $\mathcal{D}_1$ such that if $D \subseteq \mathcal{D}_1$, then $\sqcap D$ is the greatest lower bound (or *glb*) of D.

In addition, the semantic domain $\mathcal{D}$ may satisfy some (or all) of the following properties:

standard function types
If $x_1 \rightarrow x_2 \preceq y_1 \rightarrow y_2$, then $y_1 \preceq x_1$ and $x_2 \preceq y_2$.

standard glb types
If $S_0 \subseteq \mathcal{D}_0$ and $x_1 \in \mathcal{D}_0$, then $\sqcap S_0 \preceq x_1$ iff $\exists x_0 \in S_0$ s.t. $x_0 \preceq x_1$.

Building the domain from two sets, as in Definition 1, permits more generality and is an example of a "predicative domain" ([14]). This structure parallels the two distinct operations provided in the type language: function space $t_1 \rightarrow t_2$ and universal quantification $\forall \ldots$ (see Section 3). These operations impose different requirements on the semantic domain. In particular, since the $\forall$ quantifier introduces a glb operation (and hence produces a value in $\mathcal{D}_1$) and the $\rightarrow$ operation can be performed only on elements of $\mathcal{D}_0$, the $\forall$ quantifier cannot appear inside of a $\rightarrow$ operation. If the semantic domain has the property that $\mathcal{D}_0 = \mathcal{D}_1$, then it supports $\forall$ quantifiers inside of the $\rightarrow$ operation. It is worth noting that separating $\mathcal{D}_0$ and $\mathcal{D}_1$

not only generalizes but also simplifies some of our results. Note that condition 5 requires that the functions be lifted. This assumption is frequently invalid (e.g., in the standard ideal model of [13]). Our conjecture is that minor technical variations such as not lifting the function spaces would require some minor variations in the proofs and algorithms but that completeness would still hold.

The following two examples illustrate the most important features of semantic domains and are used throughout the paper.

EXAMPLE: [Minimal Semantic Model] Let $\mathcal{D}_0 = \mathcal{D}_1$ be the three element set $\{\bot, \top \to \top, \top\}$ and let $\preceq$ be the partial order $\bot \preceq \top \to \top \preceq \top$. In this domain, all function types are the same and this type domain does little more than detect that something is a function. For all $x, y \in \mathcal{D}$, $x \to y = \top \to \top$. It is easy to check that $\mathcal{D}$ satisfies all properties required of a semantic domain as well as standard glb types. The only property missing is standard function types (e.g., because $\bot \to \bot \preceq \top \to \top$, but $\top \not\preceq \bot$). □

EXAMPLE: [Standard Model] Let $\mathcal{D}_0$ be the set consisting of $\bot$ and $\top$ and closed under the pairing operation (denoted using the $\to$ symbol). An obvious partial order is induced on $\mathcal{D}_0$. This partial order is constructed in such a way so as to ensure that the domain has standard function types. Let $\mathcal{D}_1$ consist of all the non-empty, upward-closed subsets of $\mathcal{D}_0$. Intuitively, each element of $\mathcal{D}_1$ represents the glb of its members. Define $d_0 \preceq d_1$ iff $d_0 \supseteq d_1$. Note that there is an obvious inclusion mapping from $\mathcal{D}_0$ to $\mathcal{D}_1$ by mapping each element of $\mathcal{D}_0$ to the upward-closure of the singleton set consisting of that element. It is easy to see that $\mathcal{D}_1$ has standard glb types. □

The construction of $\mathcal{D}_1$ from $\mathcal{D}_0$ used in Example 1 is a general procedure for building a $\mathcal{D}_1$. Given a domain $\mathcal{D}_0$, the domain $\mathcal{D}_1$ can be defined to be the non-empty, upward-closed subsets of $\mathcal{D}_0$. Each element of $\mathcal{D}_1$ represents the glb of its members.

3. Syntax

The first type language we consider has only type variables and function types. In this language, as in all extensions we consider, quantification is shallow (occurs only at the outermost level).

Definition 2. Unquantified simple type expressions are generated by the grammar:

$$\tau ::= \alpha \mid \top \mid \bot \mid \tau_1 \to \tau_2$$

where α ranges over a family of type variables.

A *quantified simple type expression* has the form

$$\forall\alpha_1 \ldots \forall\alpha_n.\tau$$

where α_i is a type variable for $i = 1, \ldots, n$ and τ is an unquantified simple type expression. The type τ is called the *body* of the type.

Since $n = 0$ is a possibility in Definition 2, every unquantified simple type expression is also a quantified simple type expression. In the sequel, we use σ for a quantified type expression (perhaps with no quantifiers), and τ for a type expression without quantifiers.

A type variable is *free* in a quantified type expression if it appears in the body but not in the list of quantified variables. To give meaning to a quantified type, it is necessary to specify the meaning of its free variables. An *assignment* $\theta : \text{Vars} \to \mathcal{D}_0$ is a map from variables to the semantic domain. The assignment $\theta[\alpha \leftarrow \tau]$ is the assignment θ modified at point α to return τ.

An assignment is extended from variables to (quantified) simple type expressions as follows:

Definition 3.

1. $\theta(\top) = \top$
2. $\theta(\bot) = \bot$
3. $\theta(\tau_1 \to \tau_2) = \theta(\tau_1) \to \theta(\tau_2)$
4. $\theta(\forall\alpha.\tau) = \sqcap\{\theta[\alpha \leftarrow x](\tau) | x \in \mathcal{D}_0\}$

Note that unquantified simple type expressions are assigned meanings in $\mathcal{D}_0$ whereas quantified simple type expressions typically have meanings in $\mathcal{D}_1$ but not in $\mathcal{D}_0$.

PROPOSITION 1 $\theta(\forall\alpha_1 \ldots \forall\alpha_n.\tau) = \sqcap\{\theta[\alpha_1 \leftarrow x_1 \ldots \alpha_n \leftarrow x_n](\tau) \mid x_1, \ldots, x_n \in \mathcal{D}_0\}$

Proof: Follows immediately from Definition 3. ■

Our results for eliminating variables in quantified types hinge on knowledge about when two type expressions have the same meaning in the semantic domain. However, because type expressions may have free variables, the notion of equality must also take into account possible assignments to those free variables. We say that two quantified type expressions σ_1 and σ_2 are *equivalent*, written $\sigma_1 \equiv \sigma_2$, if for all assignments θ, we have $\theta(\sigma_1) = \theta(\sigma_2)$.

4. Simple Type Expressions

This section presents an algorithm for eliminating quantified type variables in simple type expressions and proves that the algorithm is sound and complete. The following definition formalizes what it means to correctly eliminate as many variables from a type as possible:

Definition 4. A type expression σ is *irredundant* if for all σ' such that $\sigma' \equiv \sigma$, it is the case that σ has no more quantified variables than σ'.

In general, irredundant types are not unique. It is easy to show that renaming quantified variables does not change the meaning of a type, provided we observe the usual rules of capture. Thus, $\forall\alpha.\sigma \equiv \forall\beta.\sigma[\alpha \leftarrow \beta]$ provided that β does not occur in σ. It is also true that types distinguished only by the order of quantified variables are equivalent. That is, $\forall\alpha.\forall\beta.\sigma \equiv \forall\beta.\forall\alpha.\sigma$. Our main result is that for every type there is a unique (up to renaming and reordering of bound variables) irredundant type that is equivalent.

Since equivalence ($\equiv$) is a semantic notion, irredundancy is also semantic in nature and cannot be determined by a trivial examination of syntax. The key question is: Under what circumstances can a type $\forall\alpha.\tau$ be replaced by some type $\tau[\alpha \leftarrow \tau']$ (for some type expression τ' not containing α)? In one direction we have

$$\theta(\forall\alpha.\tau) \preceq \theta[\alpha \leftarrow \tau'](\tau) = \theta(\tau[\alpha \leftarrow \tau'])$$

Then, using Definition 3, it follows that

$$\forall\alpha.\tau \equiv \tau[\alpha \leftarrow \tau']$$

if and only if for all assignments θ

$$\forall d \in \mathcal{D}_0.\ \theta(\tau[\alpha \leftarrow \tau']) \preceq \theta[\alpha \leftarrow d](\tau)$$

In other words, a type $\sigma = \forall\alpha.\tau$ is equivalent to $\tau[\alpha \leftarrow \tau']$ whenever for all assignments θ, we have $\theta(\tau[\alpha \leftarrow \tau'])$ is the minimal element of the set $\{\theta[\alpha \leftarrow x](\tau) | x \in \mathcal{D}_0\}$ to which the glb operation is applied in computing σ's meaning under θ.

The difficulty in computing irredundant types is that the function-space constructor $\rightarrow$ is *anti-monotonic* in its first position. That is, $\tau_1 \preceq \tau_2$ implies that $\tau_1 \rightarrow \tau \succeq \tau_2 \rightarrow \tau$. Thus, determining the minimal element of a greatest lower bound computation may require maximizing or minimizing a variable, depending on whether the type is monotonic or anti-monotonic in that variable. Intuitively, to eliminate as many variables as possible, variables in anti-monotonic positions should be set to $\top$, while others in monotonic positions should be set to $\bot$. We define functions *Pos* and *Neg* that compute a type's set of monotonic and anti-monotonic variables, respectively.

Definition 5. *Pos* and *Neg* are defined as follows:

$$\begin{aligned}
Pos(\alpha) &= \{\alpha\} \\
Pos(\tau_1 \rightarrow \tau_2) &= Neg(\tau_1) \cup Pos(\tau_2) \\
Pos(\top) &= \emptyset \\
Pos(\bot) &= \emptyset \\
\\
Neg(\alpha) &= \emptyset
\end{aligned}$$

$$\begin{aligned} Neg(\tau_1 \to \tau_2) &= Pos(\tau_1) \cup Neg(\tau_2) \\ Neg(\top) &= \emptyset \\ Neg(\bot) &= \emptyset \end{aligned}$$

As an example, for the type $\alpha \to \beta$ we have

$$\begin{aligned} Pos(\alpha \to \beta) &= \{\beta\} \\ Neg(\alpha \to \beta) &= \{\alpha\} \end{aligned}$$

The following lemma shows that *Pos* and *Neg* correctly characterize variables in monotonic and anti-monotonic positions respectively.

LEMMA 1 Let $d_1, d_2 \in \mathcal{D}_0$ where $d_1 \preceq d_2$. Let θ be any assignment.

1. If $\alpha \notin Pos(\tau)$, then $\theta[\alpha \leftarrow d_2](\tau) \preceq \theta[\alpha \leftarrow d_1](\tau)$.
2. If $\alpha \notin Neg(\tau)$, then $\theta[\alpha \leftarrow d_1](\tau) \preceq \theta[\alpha \leftarrow d_2](\tau)$.

Proof: This proof is an easy induction on the structure of τ.

- If $\tau = \bot$ or $\tau = \top$, then $\theta[\alpha \leftarrow d_1](\tau) = \theta(\tau) = \theta[\alpha \leftarrow d_2](\tau)$, so both (1) and (2) hold.
- If $\tau = \alpha$, then $\alpha \in Pos(\alpha)$, so (1) holds vacuously. For (2), we have

  $$\theta[\alpha \leftarrow d_1](\alpha) = d_1 \preceq d_2 = \theta[\alpha \leftarrow d_2](\alpha)$$

- Let $\tau = \tau_1 \to \tau_2$. We prove only (1), as the proof for (2) is symmetric. So assume that $\alpha \notin Pos(\tau)$. By the definition of *Pos*, we know

  $$\alpha \notin Neg(\tau_1) \cup Pos(\tau_2)$$

  Applying the lemma inductively to τ_1 and τ_2, we have

  $$\begin{aligned} \theta[\alpha \leftarrow d_1](\tau_1) &\preceq \theta[\alpha \leftarrow d_2](\tau_1) \\ \theta[\alpha \leftarrow d_2](\tau_2) &\preceq \theta[\alpha \leftarrow d_1](\tau_2) \end{aligned}$$

  Combining these two lines using axiom 5 of a semantic domain (Definition 1) it follows that

  $$\theta[\alpha \leftarrow d_2](\tau_1 \to \tau_2) \preceq \theta[\alpha \leftarrow d_1](\tau_1 \to \tau_2)$$

  which proves the result.

■

COROLLARY 1

1. If $\alpha \notin Pos(\tau)$, then $\theta(\tau[\alpha \leftarrow \top]) \preceq \theta(\tau) \preceq \theta(\tau[\alpha \leftarrow \bot])$ holds for all assignments θ.
2. If $\alpha \notin Neg(\tau)$, then $\theta(\tau[\alpha \leftarrow \bot]) \preceq \theta(\tau) \preceq \theta(\tau[\alpha \leftarrow \top])$ holds for all assignments θ.

4.1. Variable Elimination

Our algorithm for eliminating variables from quantified types is based on the computation of *Pos* and *Neg*. Before presenting the variable elimination procedure, we extend *Pos* and *Neg* to quantified types:

$$\begin{aligned} Pos(\forall\alpha.\sigma) &= Pos(\sigma) - \{\alpha\} \\ Neg(\forall\alpha.\sigma) &= Neg(\sigma) - \{\alpha\} \end{aligned}$$

The following lemma gives sufficient conditions for a variable to be eliminated.

LEMMA 2 If σ is a quantified simple type expression, then

$$\begin{aligned} \alpha \notin Neg(\sigma) &\Rightarrow \forall\alpha.\sigma \equiv \sigma[\alpha \leftarrow \bot] \\ \alpha \notin Pos(\sigma) &\Rightarrow \forall\alpha.\sigma \equiv \sigma[\alpha \leftarrow \top] \end{aligned}$$

Proof: Assume first that $\forall\alpha.\sigma = \forall\alpha.\tau$ where τ is an unquantified simple type expression and that $\alpha \notin Neg(\tau)$. Note that

$$\begin{aligned} & \theta(\forall\alpha.\tau) & \\ = \; & \sqcap\{\theta[\alpha \leftarrow x](\tau) | x \in \mathcal{D}_0\} & \\ \preceq \; & \theta[\alpha \leftarrow \bot](\tau) & \text{since } \bot \text{ is a possible choice for } x \\ = \; & \theta(\tau[\alpha \leftarrow \bot]) & \\ = \; & \sqcap\{\theta[\alpha \leftarrow x](\tau[\alpha \leftarrow \bot]) | x \in \mathcal{D}_0\} & \text{since } \alpha \text{ does not occur in } \tau[\alpha \leftarrow \bot] \\ \preceq \; & \sqcap\{\theta[\alpha \leftarrow x](\tau) | x \in \mathcal{D}_0\} & \text{by part 2 of Corollary 1} \\ = \; & \theta(\forall\alpha.\tau) & \end{aligned}$$

Therefore, $\theta(\forall\alpha.\tau) = \theta(\tau[\alpha \leftarrow \bot])$ for all assignments θ. For the general (quantified) case $\forall\alpha_1, \ldots, \alpha_n.\tau$, observe that any variable α_i for $1 \leq i \leq n$ can be moved to the innermost position of the type by a sequence of bound variable interchanges and renamings, at which point the reasoning for the base case above can be applied. The proof for the second statement ($\alpha \notin Pos(\sigma)$) is symmetric. ■

We are interested in quantified types for which as many variables have been eliminated using the conditions of Lemma 2 as possible. Returning to our canonical example,

$$
\begin{array}{ll}
& \forall\alpha.\forall\beta.\alpha \to \beta \\
\equiv & \forall\beta.\top \to \beta \qquad \text{since } \alpha \notin \mathit{Pos}(\forall\beta.\alpha \to \beta) \\
\equiv & \top \to \bot \qquad \text{since } \alpha \notin \mathit{Neg}(\top \to \beta)
\end{array}
$$

Definition 6. A quantified simple type expression σ is *reduced* if

- σ is unquantified; or
- $\sigma = \forall\alpha.\sigma'$ and furthermore $\alpha \in \mathit{Pos}(\sigma') \wedge \alpha \in \mathit{Neg}(\sigma')$ and σ' is reduced.

Note that the property of being reduced is distinct from the property of being irredundant. "Reduced" is a syntactic notion and does not depend on the semantic domain. Irredundancy is a semantic notion, because it involves testing the expression's meaning against the meaning of other type expressions.

Procedure 1 (Variable Elimination Procedure (VEP)) Given a quantified type expression
$\forall\alpha_1 \ldots \forall\alpha_n.\tau$, compute the sets $\mathit{Pos}(\tau)$ and $\mathit{Neg}(\tau)$. Let $VEP(\forall\alpha_1 \ldots \forall\alpha_n.\tau)$ be the type obtained by:

1. dropping any quantified variable not used in τ,
2. setting any quantified variable α where $\alpha \notin \mathit{Pos}(\forall\alpha_1 \ldots \forall\alpha_n.\tau)$ to $\bot$,
3. setting any quantified variable α where $\alpha \notin \mathit{Neg}(\forall\alpha_1 \ldots \forall\alpha_n.\tau)$ to $\top$,
4. and retaining any other quantified variable.

THEOREM 2 Let σ be any quantified simple type expression. Then $\sigma \equiv VEP(\sigma)$ and $VEP(\sigma)$ is reduced.

Proof: Equivalence follows easily from Lemma 2. To see that $VEP(\sigma)$ is reduced, observe that any quantified variable not satisfying conditions (1)–(3) of the Variable Elimination Procedure must occur both positively and negatively in the body of σ. ■

A few remarks on the Variable Elimination Procedure are in order. The algorithm can be implemented very efficiently. Two passes over the structure of the type are needed: one to compute the *Pos* and *Neg* sets (which can be done using a using a hash-table or bit-vector implementation of sets) and another to perform any substitutions. In addition, the algorithm need only be applied once, as $VEP(VEP(\sigma)) = VEP(\sigma)$.

THEOREM 3 Every irredundant simple type expression is reduced.

Proof: Let σ be an irredundant simple type expression. Since σ is irredundant, $VEP(\sigma)$ has at least as many quantified variables as σ. Therefore $VEP(\sigma) = \sigma$; i.e., the Variable Elimination Procedure does not remove any variables from σ. Since $VEP(\sigma)$ is reduced, σ is a reduced simple type expression. ■

4.2. Completeness

If σ is a quantified simple type expression, then $VEP(\sigma)$ is an equivalent reduced simple type expression, possibly with fewer quantified variables. In this section, we address whether additional quantified variables can be eliminated from a reduced type. In other words, is a reduced simple type expression irredundant? We show that if the semantic domain $\mathcal{D}$ has standard function types (Definition 1) then every reduced simple type expression is irredundant (Theorem 5).

For semantic domains with standard function types, the Variable Elimination Procedure is complete in the sense that no other algorithm can eliminate more quantified variables and preserve equivalence. The completeness proof shows that whenever two reduced types are equivalent, then they are syntactically identical, up to renamings and reorderings of quantified variables.

To simplify the presentation that follows, we introduce some new notation and terminology. By analogy with the α-reduction of the lambda calculus, two quantified simple type expressions are *α-equivalent* iff either can be obtained from the other by a series of reorderings or capture-avoiding renamings of quantified variables. We sometimes use the notation $\forall\{\alpha_1, \ldots, \alpha_n\}.\tau$ to denote $\forall\alpha_1 \ldots \forall\alpha_n.\tau$. Using a set instead of an ordered list involves no loss of generality since duplicates never occur in reduced expressions and variable order can be permuted freely. We generally use the letters s and t to range over type expressions.

4.2.1. Constraint Systems Proving completeness requires a detailed comparison of the syntactic structure of equivalent reduced types. This comparison is more intricate than might be expected; in addition, in the sequel we perform a similar analysis to prove that variable elimination is complete for recursive types. This section develops the technical machinery at the heart of both completeness proofs.

Definition 7. A *system of constraints* is a set of inclusion relations between unquantified simple type expressions $\{\ldots s \preceq t \ldots\}$. A *solution* of the constraints is any assignment θ such that $\theta(s) \preceq \theta(t)$ holds for all constraints $s \preceq t$ in the set.

Definition 8 gives an algorithm B that compares two unquantified simple type expressions t_1 and t_2. The comparison is expressed in terms of constraints; the function B transforms a constraint $t_1 \preceq t_2$ into a system of constraints such that at least one side of each inequality in the system of constraints is a variable of t_1 or t_2. Intuitively, $B(\{t_1 \preceq t_2\})$ summarizes what must be true about the variables of the two types whenever the relationship $t_1 \preceq t_2$ holds.

Definition 8. Let S be a set of constraints. $B(S)$ is a set of constraints defined by the following rules. These clauses are to be applied in order with the earliest one that applies taking precedence.

1. $B(\emptyset) = \emptyset$
2. $B(\{t \preceq t\} \cup S) = B(S)$.

3. $B(\{s_1 \to s_2 \preceq t_1 \to t_2\} \cup S) = B(\{t_1 \preceq s_1, s_2 \preceq t_2\} \cup S)$.

4. Otherwise, $B(\{s \preceq t\} \cup S) = \{s \preceq t\} \cup B(S)$.

LEMMA 3 Let S be a system of constraints. If $\mathcal{D}$ is a semantic domain with standard function types, then every solution of S is a solution of $B(S)$.

Proof: Let the *complexity* of S be the pair *(number of* $\to$ *symbols in* S*, number of constraints in* S*)*. Complexity is ordered lexicographically, so $(i, j) < (i', j')$ if $i < i'$ or $i = i'$ and $j < j'$. The result is proven by induction on the complexity of S, with one case for each clause in the definition of B:

1. $B(\emptyset) = \emptyset$. The result clearly holds.

2. Since any assignment is a solution of $t \preceq t$, any solution θ of $\{t \preceq t\} \cup S$ is also a solution of S. By induction, θ is a solution of $B(S)$.

3. Let θ be a solution of $\{s_1 \to s_2 \preceq t_1 \to t_2\} \cup S$. Since the domain has standard function types, it follows that θ is also a solution of $\{t_1 \preceq s_1, s_2 \preceq t_2\} \cup S$. By induction, θ is a solution of $B(\{t_1 \preceq s_1, s_2 \preceq t_2\} \cup S)$.

4. In the final case, by induction every solution of S is a solution of $B(S)$. Therefore all solutions of $\{s \preceq t\} \cup S$ are solutions of $\{s \preceq t\} \cup B(S)$.

■

The completeness proof uses an analysis of the constraints $B(\{t_1 \preceq t_2\})$ where t_1 and t_2 are the bodies of reduced equivalent types. Observe that if t_1 and t_2 differ only in the names of variables, then $B(\{t_1 \preceq t_2\})$ is a system of constraints between variables. Furthermore, it turns out that if t_1 and t_2 are actually renamings of each other (and if t_1 and t_2 are the bodies of reduced equivalent types) then the constraints $B(\{t_1 \preceq t_2\})$ define this renaming in both directions. Proving this claim is a key step in the proof. This discussion motivates the following definition:

Definition 9. A system S of constraints is (V_1, V_2)-*convertible* iff V_1, V_2 are disjoint sets of variables and there is a bijection f from V_1 to V_2 such that $S = \{\alpha \preceq f(\alpha) | \alpha \in V_1\} \cup \{f(\alpha) \preceq \alpha | \alpha \in V_1\}$

EXAMPLE: For example, let S be the system of constraints

$$\alpha \preceq \gamma \quad \gamma \preceq \alpha$$
$$\beta \preceq \delta \quad \delta \preceq \beta$$

Let $V_1 = \{\alpha, \beta\}, V_2 = \{\gamma, \delta\}$, and define $f : V_1 \to V_2$ such that $f(\alpha) = \gamma, f(\beta) = \delta$. It is easy to check that S is (V_1, V_2)-convertible. □

The idea behind Definition 9 is that if two reduced types $\forall V_1.\tau_1$ and $\forall V_2.\tau_2$ are α-convertible, then $B(\{\tau_1 \preceq \tau_2\})$ is a (V_1, V_2)-convertible system of constraints (provided V_1 and V_2 are disjoint). It is easiest to prove this fact by first introducing an alternative characterization of convertible constraint systems, which is given in the following technical definition and lemma.

Definition 10. A system of constraints $\{s_1 \preceq t_1, \ldots, s_n \preceq t_n\}$ is (V_1, V_2)-*miniscule* iff the following all hold:

1. V_1 and V_2 are disjoint sets of variables.
2. for all $i \leq n$, at most one of s_i and t_i is a $\rightarrow$ expression.
3. for all $i \leq n$, s_i and t_i are different expressions.
4. for each $v \in V_1 \cup V_2$, there exists $i \leq n$ such that $v \in \mathit{Pos}(s_i) \cup \mathit{Neg}(t_i)$
5. for each $v \in V_1 \cup V_2$, there exists $i \leq n$ such that $v \in \mathit{Neg}(s_i) \cup \mathit{Pos}(t_i)$
6. for every assignment θ there is a assignment θ' such that $\theta(v) = \theta'(v)$ for all $v \notin V_1$ and
   $\theta'(s_i) \preceq \theta'(t_i)$ holds for all $i \leq n$.
7. for every assignment θ there is a assignment θ' such that $\theta(v) = \theta'(v)$ for all $v \notin V_2$ and
   $\theta'(t_i) \preceq \theta'(s_i)$ holds for all $i \leq n$. (Note the reverse order of t_i and s_i.)

LEMMA 4 A system of constraints is (V_1, V_2)-miniscule iff it is (V_1, V_2)-convertible.

Proof: It is easy to check that any (V_1, V_2)-convertible system of constraints is (V_1, V_2)-miniscule.

To prove the converse, let θ_0 be the assignment that assigns $\bot$ to every variable, let θ_1 be the assignment that assigns $\top$ to every variable, and let S be a (V_1, V_2)-miniscule system of constraints. The first step is to show that no $\rightarrow$ expressions can occur in S. It is easy to check that if we reverse all inequalities we get a (V_2, V_1)-miniscule system of constraints. Thus, by symmetry, to show that $\rightarrow$ cannot occur in S it suffices to show that $\rightarrow$ cannot occur in any upper bound in S.

For the sake of obtaining a contradiction, assume that $s_i \preceq t'_i \rightarrow t''_i \in S$. We show that each of the four possible forms for s_i is impossible.

1. $s'_i \rightarrow s''_i \preceq t'_i \rightarrow t''_i$ is ruled out by Property 2 of Definition 10.
2. $\bot \preceq t'_i \rightarrow t''_i$ is ruled out by Property 7 of Definition 10, since, by requirement 5 in the definition of domains (Definition 1), no assignment satisfies $t'_i \rightarrow t''_i \preceq \bot$.
3. $\top \preceq t'_i \rightarrow t''_i$ is ruled out by Property 6 of Definition 10, since, by requirement 5 in the definition of domains (Definition 1), no assignment satisfies $\top \preceq t'_i \rightarrow t''_i$.

4. Suppose s_i is the variable v.

   If v is a variable not in V_1, let $\theta = \theta_1$. Then Property 6 of Definition 10 is violated because for all θ' that agree with θ off of V_1, we have $\theta'(v) = \theta(v) = \theta_1(v) = \top \not\preceq \theta'(t_i' \to t_i'')$.

   If $v \in V_1$, let $\theta = \theta_0$. Note that $v \notin V_2$ since V_1 and V_2 are disjoint. Then Property 7 of Definition 10 is violated because for all θ' that agree with θ off of V_2, we have $\theta'(t_i' \to t_i'') \not\preceq \bot = \theta_0(v) = \theta(v) = \theta'(v)$.

This completes the proof that $\to$ cannot occur in S.

The next step is to show that $\bot$ cannot occur in S. By symmetry it suffices to show that $\bot$ cannot occur as an upper bound in S. There are three cases to consider.

1. $\bot \preceq \bot$ is ruled out by Property 3 in Definition 10.

2. $\top \preceq \bot$ is ruled out by Property 6 in Definition 10 since no assignment satisfies $\top \preceq \bot$.

3. Consider $v \leq \bot$ where v is a variable. If $v \notin V_1$, let $\theta = \theta_1$. Then Property 6 in Definition 10 is violated since for all θ' that agree with θ off of V_1, we have that $\theta'(v) = \theta(v) = \theta_1(v) = \top \not\preceq \bot = \theta'(\bot)$.

   If $v \in V_1$, a complex case argument is needed because Property 6 is not directly violated. By Property 5 of Definition 10, there is a constraint $s' \preceq v$ in S. There are four possible cases for s':

   (A) $s' = \bot$. In this case, $\bot \preceq v \preceq \bot$ is in S and hence Property 7 is violated by taking $\theta = \theta_1$.

   (B) $s' = \top$. In this case, $\top \preceq v \preceq \bot$ violates Property 6 since it is never satisfied by any assignment.

   (C) $s' = v' \in V_1$. In this case, $v' = v$ is ruled out by Property 3. So we may assume that v' and v are different variables. Property 7 is violated by taking $\theta = \theta_0[v \leftarrow \top]$ since if θ' agrees with θ off of V_2 the constraint $v \preceq v'$ is violated since $\theta'(v) = \theta(v) = \top \not\preceq \bot = \theta(v') = \theta'(v')$.

   (D) $s' = v' \notin V_1$. In this case, $v' \preceq v \preceq \bot$ violates Property 6 by taking $\theta = \theta_1$ since $\theta(v') = \top$.

This proves that $\bot$ cannot occur as an upper bound in S. By symmetry, $\bot$ cannot occur as a lower bound in S, and hence $\bot$ cannot occur anywhere in S. An analogous argument shows that $\top$ can not occur anywhere in S either.

Thus, every element of S is of the form $v' \preceq v''$ for variables v', v''. We now show that $v', v'' \in V_1 \cup V_2$. Suppose that $v' \notin V_1 \cup V_2$. If $v'' \in V_1$, then Property 7 is violated by taking $\theta = \theta_0[v'' \leftarrow \top]$ since for all θ' that agree with θ off of V_2, we have that $\theta'(v'') = \theta(v'') = \top \not\preceq \bot = \theta(v') = \theta'(v')$. If $v'' \notin V_1$, then Property 6 is violated by taking $\theta = \theta_0[v' \leftarrow \top]$ since for all θ' that agree with θ off of V_1, we have that $\theta'(v') = \theta(v') = \top \not\preceq \bot = \theta(v'') = \theta'(v'')$. Therefore, the supposition

that $v' \notin V_1 \cup V_2$ is false and it follows that $v' \in V_1 \cup V_2$. A similar argument shows that $v'' \in V_1 \cup V_2$.

If both v' and v'' are in V_1, then Property 7 is violated by taking $\theta = \theta_0[v'' \leftarrow \top]$ since for all θ' that agree with θ off of V_2, we have that $\theta'(v'') = \theta(v'') = \top \not\preceq \bot = \theta(v') = \theta'(v')$. This shows that not both v' and v'' are in V_1. A symmetric argument shows that not both v' and v'' are in V_2. Thus, it follows that for every constraint $s_i \preceq t_i$ in S, either $s_i \in V_1$ and $t_i \in V_2$ or $s_i \in V_2$ and $t_i \in V_1$.

Next we show that if $v_0 \preceq v_1 \preceq v_2$, then $v_0 = v_2$. First assume that $v_1 \in V_1$. If v_0 and v_2 are different variables, then Property 6 is violated by taking $\theta = \theta_0[v_0 \leftarrow \top]$ since for all θ' that agree with θ off of V_1, we have that $\theta'(v_0) = \theta(v_0) = \top \not\preceq \bot = \theta(v_2) = \theta'(v_2)$. Hence, in the case that $v_1 \in V_1$, it follows that $v_0 = v_2$. A similar argument shows that if $v_1 \in V_2$, then $v_0 = v_2$.

The next goal is to show that for every $v_1 \in V_1$, there exists a unique $v_2 \in V_2$ such that $v_1 \preceq v_2$ is in S. By Property 4, there is at least one such v_2. Let v_2' be any variable such that $v_1 \preceq v_2'$ is in S. By Property 5, there is a v_0 such that $v_0 \preceq v_1$ is in S. It follows that $v_2 = v_0 = v_2'$ which proves that v_2 is unique.

Define a function f mapping V_1 to V_2 so that $v_1 \preceq f(v_1)$ is in S. By Property 5, for any $v_1 \in V_1$, there is a v_0 such that $v_0 \preceq v_1$ is in S. It follows that $v_0 = f(v_1)$. This proves that $S \subseteq \{\alpha \preceq f(\alpha) | \alpha \in V_1\} \cup \{f(\alpha) \preceq \alpha | \alpha \in V_1\}$. Since every constraint in S has the form $v' \preceq v''$ where either v' or v'' is in V_1 and since the upper and lower bounds are unique (because $v_0 \preceq v_1 \preceq v_2 \in S$ implies that $v_0 = v_2$), it follows that there are no extra elements of S. Therefore, $S = \{\alpha \preceq f(\alpha) | \alpha \in V_1\} \cup \{f(\alpha) \preceq \alpha | \alpha \in V_1\}$. Thus S is (V_1, V_2)-convertible as desired. ■

4.2.2. From Constraints to Completeness The definitions and lemmas of Section 4.2.1 are the building blocks of the completeness proof. Before finally presenting the proof, we need one last definition:

Definition 11. Two simple type expressions $\forall V_1.\tau_1$ and $\forall V_2.\tau_2$ are *compatible* iff $\forall V_1.\tau_1$ and $\forall V_2.\tau_2$ are equivalent reduced simple type expressions such that V_1 and V_2 are disjoint and no variable in V_1 occurs in τ_2 and no variable in V_2 occurs in τ_1.

The important part of the definition of compatibility is that the type expressions are reduced and equivalent. The conditions regarding quantified variables are there merely to simplify proofs. There is no loss of generality because α-conversion can be applied to convert any two equivalent reduced type expressions into compatible expressions.

LEMMA 5 Let $\forall V_1.\tau_1$ and $\forall V_2.\tau_2$ be compatible type expressions. If the semantic domain has standard function types and standard glb types, then $B(\{\tau_1 \preceq \tau_2\})$ is a (V_1, V_2)-miniscule system of constraints.

Proof: Let $B(\{\tau_1 \preceq \tau_2\}) = \{s_1 \preceq t_1, \ldots, s_n \preceq t_n\}$. We prove that the conditions in Definition 10 all hold:

1. By compatibility V_1 and V_2 are disjoint sets of variables.

2. By Part 3 of Definition 8 at most one of s_i and t_i is a $\rightarrow$ expression.

3. For all $i \leq n$, s_i and t_i are different expressions by Part 2 of Definition 8.

4. After a number of applications of B, the intermediate result for the calculation of $B(\{\tau_1 \preceq \tau_2\})$ is of the form:

$$\{s'_1 \preceq t'_1, \ldots, s'_k \preceq t'_k\} \cup B(\{s'_{k+1} \preceq t'_{k+1}, \ldots, s'_m \preceq t'_m\}) \ .$$

   It is sufficient to show that

   (A) for all $v \in V_1 \cup V_2$, there is an $i \leq m$ such that $v \in \mathit{Pos}(s'_i) \cup \mathit{Neg}(t'_i)$, and
   (B) if, for some $i \leq m$, we have $s'_i = t'_i$, then s'_i, t'_i contain no variables.

   We proceed by induction on the number of steps needed to compute $B(\{\tau_1 \preceq \tau_2\})$. In the base case, consider $B(\{\tau_1 \preceq \tau_2\})$. The result follows from the observations that τ_1, τ_2 are reduced, and that V_1, V_2 are disjoint. In the induction step, each of the four cases from the definition of B follow immediately from the induction hypothesis.

5. Proof similar to the previous step.

6. Let θ be any assignment. We must show that there is an assignment θ' such that $\theta(v) = \theta'(v)$ for all $v \notin V_1$ and $\theta'(s_i) \preceq \theta'(t_i)$ holds for all $i \leq n$. Since $\theta(\forall V_1.\tau_1) \preceq \theta(\forall V_2.\tau_2)$, it follows that $\theta(\forall V_1.\tau_1) \preceq \theta(\tau_2)$. Since the semantic domain has standard glb types, it follows that $\theta'(\tau_1) \preceq \theta(\tau_2)$ holds for some θ' that agrees with θ except possibly on V_1. Since no variable in V_1 occurs in τ_2, we know $\theta'(\tau_1) \preceq \theta'(\tau_2)$. By Lemma 3, it follows that θ' is a solution to $B(\{\tau_1 \preceq \tau_2\})$.

7. To show that for every assignment θ there is an assignment θ' such that $\theta(v) = \theta'(v)$ for all $v \notin V_2$ and $\theta'(t_i) \preceq \theta'(s_i)$ holds for all $i \leq n$, reverse the roles of τ_1 and τ_2. This argument, which is a variation of the previous case, relies on the fact that $B(\{\tau_2 \preceq \tau_1\})$ can be obtained from $B(\{\tau_1 \preceq \tau_2\})$ by reversing the direction of the $\preceq$ symbol.

■

One final technical lemma is required before we can show that the variable elimination procedure is complete. The intuition behind Lemma 6 is that if $B(\{\tau_1 \preceq \tau_2\})$ is a (V_1, V_2)-convertible system of constraints with bijection f (recall Definition 9), then $B(\{\tau_1 \preceq f(\tau_2)\}) = \emptyset$. This intuition is not quite correct, because there may be variables in τ_1 or τ_2 that are not in $V_1 \cup V_2$. In the following lemma, $\mathit{vars}(t)$ is the set of variables appearing in t.

LEMMA 6 Assume that

1. $B(S)$ is a subset of a (V_1, V_2)-convertible system of constraints with bijection f from V_1 to V_2.
2. For each constraint $s \preceq t \in S$, we have $vars(s) \cap vars(t) \cap (V_1 \cup V_2) = \emptyset$. In other words, any variables common to s and t are not in $V_1 \cup V_2$.

Define

$$F(x) = \begin{cases} f(x) & \text{if } x \in V_1 \\ x & \text{otherwise} \end{cases}$$

We extend F from variables to terms in the usual way. Define

$$S' = \{F(t) \preceq F(t') | t \preceq t' \in S\}$$

The claim is that $B(S') = \emptyset$.

Proof: The proof is by induction on the complexity of S, as defined in the proof of Lemma 3.

- $S = \emptyset$. Then $S' = \emptyset$ and $B(\emptyset) = \emptyset$.
- $S = \{t \preceq t\} \cup S_1$. By assumption (2), $vars(t) \cap (V_1 \cup V_2) = \emptyset$. By the definition of F it follows that $F(t) = t$. Using the definition of B, it is easy to see that because S satisfies assumptions (1) and (2) with bijection f that S_1 also satisfies assumptions (1) and (2) with the same bijection f. Now we have

$$\begin{array}{lll} & \emptyset & \\ = & B(S_1') & \text{by induction} \\ = & B(\{t \preceq t\} \cup S_1') & \text{definition of } B \\ = & B(\{F(t) \preceq F(t)\} \cup S_1') & F(t) = t \\ = & B(S') & \end{array}$$

- $S = \{t_1 \to t_2 \preceq s_1 \to s_2\} \cup S_1$. Let $T = \{s_1 \preceq t_1, t_2 \preceq s_2\} \cup S_1$.

  Using the definition of B, it is easy to check that T satisfies conditions (1) and (2) using the bijection f. By induction $B(T') = \emptyset$. Then

$$\begin{array}{lll} & \emptyset & \\ = & B(\{F(s_1) \preceq F(t_1), F(t_2) \preceq F(s_2)\} \cup S_1') & \text{by induction} \\ = & B(\{F(t_1) \to F(t_2) \preceq F(s_1) \to F(s_2)\} \cup S_1') & \text{definition of } B \\ = & B(\{F(t_1 \to t_2) \preceq F(s_1 \to s_2)\} \cup S_1') & \text{definition of } F \\ = & B(S') & \end{array}$$

- $S = \{s \preceq t\} \cup S_1$ and no previous case applies. Then $B(S) = \{s \preceq t\} \cup B(S_1)$. Since $B(S)$ is a subset of a (V_1, V_2)-convertible system of constraints, it follows that $s = \alpha$ and $t = \beta$ for some distinct variables α and β and that either $F(\alpha) = \beta$ and $F(\beta) = \beta$ or $F(\alpha) = \alpha$ and $F(\beta) = \alpha$. The rest is similar to the case for $t \preceq t$ above.

■

We are now ready to state and prove the first of the major theorems concerning completeness.

THEOREM 4 If the semantic domain has standard function types and standard glb types, then any two reduced quantified simple type expressions are equivalent iff they are α-equivalent.

Proof: The if-direction is clear and does not even require that the semantic domain have standard function types. To prove the only-if direction, let σ' and σ'' be two reduced quantified simple type expressions. If necessary, α-convert σ' to $\sigma_1 = \forall V_1.\tau_1$ and α-convert σ'' to $\sigma_2 = \forall V_2.\tau_2$ so that σ_1 and σ_2 are compatible. It suffices to show that σ_1 and σ_2 are α-equivalent.

By Lemma 5, $B(\{\tau_1 \preceq \tau_2\})$ is a (V_1, V_2)-miniscule system of constraints. By Lemma 4, $B(\{\tau_1 \preceq \tau_2\})$ is a (V_1, V_2)-convertible system of constraints; let f be corresponding bijection mapping variables in V_1 to V_2. Define

$$F(x) = \begin{cases} f(x) & \text{if } x \in V_1 \\ x & \text{otherwise} \end{cases}$$

Because σ_1 and σ_2 are compatible, $vars(\tau_1) \cap vars(\tau_2) \cap (V_1 \cup V_2) = \emptyset$. Then, by Lemma 6, we have

$$B(\{F(\tau_1) \preceq F(\tau_2)\}) = \emptyset$$

Since F is the identity on τ_2 it follows that

$$B(\{F(\tau_1) \preceq \tau_2\}) = \emptyset$$

from which it follows by the definition of B that $\tau_2 = F(\tau_1)$. This shows that σ_1 and σ_2 are α-equivalent as desired.

■

COROLLARY 2 If the semantic domain has standard function types, then no two different unquantified simple type expressions are equivalent.

Proof: Given a semantic domain $\mathcal{D}$ construct another semantic domain $\mathcal{D}'$ such that $\mathcal{D}_0 = \mathcal{D}'_0$ and $\mathcal{D}'$ has standard glb types using the construction in Example 1. Using the semantic domain $\mathcal{D}'$ suffices because the meaning of an unquantified type expression is always an element of $\mathcal{D}_0$ and $\mathcal{D}'_0 = \mathcal{D}_0$. If τ and τ' are equivalent, unquantified simple type expressions, then they are reduced and hence α-equivalent by Theorem 4. But since they have no quantifiers, α-equivalence implies that $\tau = \tau'$.

■

Finally, the following theorem states our main result.

THEOREM 5 If the semantic domain has standard function types and standard glb types, then a simple type expression is reduced iff it is irredundant.

Proof: The if direction follows from Theorem 3. To prove the only-if direction, let σ be a reduced simple type expression with the goal of proving that σ is irredundant. Let σ' be an irredundant type that is equivalent to σ. (Such a σ' can always be found by picking it to be a type expression equivalent to σ with the smallest possible number of quantified variables.) By Theorem 3, σ' is reduced. By Theorem 4, σ is α-equivalent to σ'. Therefore, it follows that σ and σ' have the same number of quantified variables. Hence, σ is irredundant as desired. ■

Theorem 5 shows that a syntactic test (reduced) is equivalent to a semantic test (irredundant). Theorem 5 requires that the semantic domain has standard function types. The following examples show that this assumption is necessary.

EXAMPLE: Consider the minimal semantic domain (Example 1). It is clear that $\forall\alpha.(\alpha \to \alpha) \equiv (\top \to \top)$ in the minimal semantic domain. Therefore, $\forall\alpha.(\alpha \to \alpha)$ is reduced but not irredundant. □

EXAMPLE: In the semantic domain used in [2], $x \to \top = y \to \top$ regardless of the values of x and y, because if the answer can be anything (i.e., $\top$), it does not matter what the domain is. In this case, $\forall\alpha.((\alpha \to \alpha) \to \top) \equiv \top \to \top$. Thus, $\forall\alpha.((\alpha \to \alpha) \to \top)$ is not irredundant even though it is reduced. □

Theorem 6 shows that the Variable Elimination Procedure (Procedure 1) is complete provided that the semantic domain has standard function types.

THEOREM 6 Let σ be a quantified simple type expression. If the semantic domain has standard function types and standard glb types, then $VEP(\sigma)$ is an irredundant simple type expression equivalent to σ.

Proof: Follows easily from Theorem 2 and Theorem 5. ■

To summarize, for simple type expressions the Variable Elimination Procedure that removes quantified variables occurring positively or negatively in a type produces an equivalent type with the minimum number of quantified variables. Furthermore, this type is unique up to the renaming and order of quantified variables.

A good feature of Theorem 6 is that the irredundant type expression produced by the Variable Elimination Procedure has no more arrows than the original type expression. This need not be the case if the semantic domain does not have standard function types.

EXAMPLE: Let $\mathcal{D}_0 = \mathcal{D}_1 = \{\bot, \top \to \bot, x, \bot \to \bot, \top \to \top, \bot \to \top, \top\}$ where x is a function type, $\top \to \bot$ is less than x, and x is less than the other three function types.

Let the set of type expressions be $\top$ and $\bot$ closed under $\to$. For this domain, we define the $\to$ operator as follows. The four possibilities for combining $\top$ and $\bot$

using $\rightarrow$ map to the corresponding elements of $\mathcal{D}_0$. For all other types $y_0 \rightarrow y_1$, either y_1 or y_2 (or both) is a function type. If either y_0 or y_1 is a function type, define $y_0 \rightarrow y_1 = x$.

In this domain $\theta(\forall\alpha.(\alpha \rightarrow \alpha)) = x$ for all assignments θ. This follows because $x \preceq \top \rightarrow \top$ and $x \preceq \bot \rightarrow \bot$ and for any other $y \rightarrow y$, we have $\theta(y \rightarrow y) = x$ (e.g., $\theta((\bot \rightarrow \bot) \rightarrow (\bot \rightarrow \bot)) = x$). Now we have that $\forall\alpha.(\alpha \rightarrow \alpha) \equiv (\bot \rightarrow \bot) \rightarrow \bot$ and, in fact, the quantified type is equivalent to exactly those unquantified types with a function type in one or both of the domain or range. Even though $\forall\alpha.(\alpha \rightarrow \alpha)$ has only one arrow, every irredundant type expression equivalent to $\forall\alpha.(\alpha \rightarrow \alpha)$ has at least two arrows. □

5. Recursive Type Expressions

This section extends the basic variable elimination algorithm to a type language with recursive types. The proofs of soundness and completeness parallel the structure of the corresponding proofs for the non-recursive case.

New issues arise in two areas. First, there is new syntax for recursive type equations, which requires corresponding extensions to the syntax-based algorithms (*Pos*, *Neg*, and *B*). Second, two new conditions on the semantic domain are needed. Roughly speaking, the two conditions are (a) that recursive equations have solutions in the semantic domain (which is needed to give meaning to recursive type expressions) and (b) that the ordering $\preceq$ satisfies a continuity property (which is required to guarantee correctness of the *Pos* and *Neg* computations). It is surprising that condition (b) is needed not just for completeness, but even for soundness. Fortunately, standard models of recursive types (including the ideal model and regular trees) satisfy both conditions.

5.1. Preliminaries

We begin by defining a type language with recursive types. We first require the technical notion of a *contractive equation*.

Definition 12. Let $\delta_1, \ldots, \delta_n$ be distinct type variables and let $\tau_1, \ldots, \tau_n$ be unquantified simple type expressions. A variable α is *contractive in an equation* $\delta_1 = \tau_1$ if every occurrence of α in τ_1 is inside a constructor (such as $\rightarrow$). A system of equations

$$\delta_1 = \tau_1 \wedge \ldots \wedge \delta_n = \tau_n$$

is *contractive* iff each δ_i is contractive in every equation of the system.

Contractiveness is a standard technical condition in systems with recursive types [13]. Contractiveness is necessary for equations to have unique solutions (e.g., an equation such as $\delta = \delta$ may have many solutions). The results of this section only apply to systems of contractive equations.

Definition 13. An *(unquantified) recursive type expression* is of the form: τ/E where E is a set of contractive equations and τ is an unquantified simple type expression.

Throughout this section, we use $\delta, \delta_1, \delta', \ldots$ for the *defined* variables that are given definitions in the set of equations E, and we use $\alpha, \alpha', \alpha_1, \ldots$ to indicate the *regular* variables, i.e., those that are not given definitions. To give meaning to recursive type expressions, the equations in a recursive type must have solutions in the semantic domain. The following definition formalizes this requirement.

Definition 14. A semantic domain has *contractive solutions* iff for every contractive system E of equations

$$\delta_1 = \tau_1 \wedge \ldots \wedge \delta_n = \tau_n$$

and for every assignment θ, there exists a unique assignment θ^E such that:

1. $\theta^E(\alpha) = \theta(\alpha)$ for all $\alpha \notin \{\delta_1, \ldots, \delta_n\}$
2. $\theta^E(\delta_i) = \theta^E(\tau_i)$ for all $i = 1, \ldots, n$.

Note that Definition 14 is well formed because assignments are applied only to unquantified simple type expressions, an operation that already has meaning (see Definition 3).

LEMMA 7 Let E and E' be contractive systems of equations and assume the semantic domain has contractive solutions.

1. Let θ be an assignment. If $E' \subseteq E$, then $(\theta^E)^{E'} = \theta^E$.
2. If E does not mention variables $\alpha_1, \ldots, \alpha_m$, then $(\theta[\alpha_1 \leftarrow d_1, \ldots, \alpha_m \leftarrow d_m])^E = \theta^E[\alpha_1 \leftarrow d_1, \ldots, \alpha_m \leftarrow d_m]$

Proof:

1. Immediate from the uniqueness of θ^E.
2. By repeated applications, it suffices to consider the case $m = 1$. If β is not defined by E, it is easy to see that $\theta^E[\alpha_1 \leftarrow d_1](\beta) = \theta[\alpha_1 \leftarrow d_1](\beta)$. If $\delta = \tau$ is a definition in E, then $\theta^E[\alpha_1 \leftarrow d_1](\delta) = \theta^E(\delta) = \theta^E(\tau) = \theta^E[\alpha \leftarrow d_1](\tau)$. By uniqueness of $(\theta[\alpha_1 \leftarrow d_1])^E$, it follows that $(\theta[\alpha_1 \leftarrow d_1])^E = \theta^E[\alpha_1 \leftarrow d_1]$.

■

An assignment is extended to (quantified) recursive type expressions as follows:

Definition 15.

1. $\theta(\tau/E) = \theta^E(\tau)$ for any unquantified simple type expression τ.
2. $\theta(\forall\alpha.\tau/E) = \sqcap\{\theta[\alpha \leftarrow x](\tau/E) | x \in \mathcal{D}_0\}$

Just as for simple type expressions, every unquantified simple type expression is assigned a meaning in $\mathcal{D}_0$ whereas quantified simple type expressions typically have meanings that are in $\mathcal{D}_1$ but not in $\mathcal{D}_0$. Lemma 8 shows that if a domain has contractive solutions, then definitions of "unused" variables can be dropped.

LEMMA 8 Assume the domain has contractive solutions. Let E by a set of equations and let $E' \subseteq E$. Assume that whenever a defined variable δ of E occurs in τ_0/E', then δ is a defined variable of E'. Then $\tau_0/E \equiv \tau_0/E'$.

Proof: Let $\delta_1, \ldots, \delta_m$ be the variables defined by E but not E'. Then we have:

$$
\begin{array}{lll}
 & \theta(\tau_0/E) & \\
= & \theta^E(\tau_0) & \\
= & (\theta^E)^{E'}(\tau_0) & \text{by part 1 of Lemma 7, since } E' \subseteq E \\
= & \theta[\delta_1 \leftarrow \theta^E(\delta_1), \ldots, \delta_m \leftarrow \theta^E(\delta_m)]^{E'}(\tau_0) & \\
= & \theta^{E'}[\delta_1 \leftarrow \theta^E(\delta_1), \ldots, \delta_m \leftarrow \theta^E(\delta_m)](\tau_0) & \text{by part 2 of Lemma 7} \\
= & \theta^{E'}(\tau_0) & \text{since } \delta_1, \ldots, \delta_m \text{ do not appear in } \tau_0 \\
= & \theta(\tau_0/E') &
\end{array}
$$

■

Surprisingly, even though contractive solutions guarantee that equations have unique solutions, this is not sufficient for soundness of the Variable Elimination Procedure. The crux of the problem is found in the reasoning that justifies using *Pos* and *Neg* (see Definition 18) as the basis for replacing variables by $\top$ or $\bot$ (Lemma 1). The *Pos* and *Neg* algorithms traverse a type expression to compute the set of positive and negative variables of the expression. In the case of recursive types, *Pos* and *Neg* can be regarded as using finite unfoldings of the recursive equations. We must ensure that these finite approximations correctly characterize the limit, which is the "infinite" unfolding of the equations. Readers familiar with denotational semantics will recognize this requirement as a kind of continuity property. Definition 17 defines *type continuity*, which formalizes the appropriate condition. Later in this section we give an example showing that type continuity is in fact necessary.

Definition 16. A *definable operator* is a function $F : \mathcal{D}_0 \to \mathcal{D}_0$ such that there is a recursive type expression $\tau_0/\bigwedge_{i=1}^m \delta_i = \tau_i$, an assignment θ, and a (regular) variable α such that $\tau_0 \neq \alpha$, α is contractive in all equations, and

$$F(d) = \theta[\alpha \leftarrow d](\tau_0/\bigwedge_{i=1}^{m} \delta_i = \tau_i)$$

holds for all $d \in \mathcal{D}_0$.

Definition 17. A semantic domain $\mathcal{D}$ has *type-continuity* iff for every monotonic, definable operator F and every $d', d'' \in \mathcal{D}_0$,

$$(F(d'') = d'' \ \wedge \ F(d') \preceq d') \ \Rightarrow \ d'' \preceq d'$$

The minimal semantic model (Example 1) has contractive solutions, type continuity, and standard glb types, but it lacks standard function types. The standard semantic model (Example 1) has standard glb types and standard function types but lacks contractive solutions (e.g., because the equation $\delta = \delta \to \delta$ has no solution). The standard model does have type continuity, but without contractive solutions type continuity is not very interesting; for the standard model, the only monotonic definable operators with a fixed point are constant functions. The standard semantic model can be extended to the usual regular tree model to provide contractive solutions without sacrificing the other properties.

LEMMA 9 The usual semantic domain of regular trees has contractive solutions, standard glb types, standard function types, and type continuity.

Proof: We briefly sketch the usual semantic domain $\mathcal{D}_0$ of regular trees. This discussion is not intended to give a detailed construction of the domain, but rather to highlight the important features. As usual, $\mathcal{D}_1$ consists of the non-empty upward closed subsets of $\mathcal{D}_0$. Therefore, the semantic domain has standard glb types.

A finite or infinite tree is *regular* if it has only a finite number of subtrees. The set $\mathcal{D}_0$ consists of the regular trees built from $\top$ and $\bot$ using the $\to$ operator. Thus, $\top$ and $\bot$ are elements of $\mathcal{D}_0$ and every other element x of $\mathcal{D}_0$ is equal to $x' \to x''$ for some $x', x'' \in \mathcal{D}_0$. Furthermore, x' and x'' are unique. It is well-known that such a domain has contractive solutions [6].

Let $x \preceq_0 y$ hold for all $x, y \in \mathcal{D}_0$. Let $x \preceq_{i+1} y$ hold iff $x = \bot$ or $y = \top$ or $x = x' \to x''$ and $y = y' \to y''$ and $x'' \preceq_i y''$ and $y' \preceq_i x'$. Notice that $\preceq_{i+1} \subseteq \preceq_i$. Then $x \preceq y$ holds iff $x \preceq_i y$ holds for all $i \geq 0$ [4].

First we check that $\preceq$ has standard function types.

$$\begin{array}{ll} & x' \to x'' \preceq y' \to y'' \\ \Leftrightarrow & \forall i (x' \to x'' \preceq_{i+1} y' \to y'') \\ \Leftrightarrow & \forall i (x'' \preceq_i y'' \text{ and } y' \preceq_i x') \\ \Leftrightarrow & \forall i (x'' \preceq_i y'') \text{ and } \forall i (y' \preceq_i x') \\ \Leftrightarrow & x'' \preceq y'' \text{ and } y' \preceq x' \end{array}$$

Thus $\preceq$ has standard function types.

Next we check that $\mathcal{D}_0$ has type continuity. Let $x =_i y$ stand for $x \preceq_i y$ and $y \preceq_i x$. Let F be a definable monotonic operator. For any $x, y \in \mathcal{D}_0$,

$$F^i(x) =_i F^i(y)$$

This fact follows by induction on i, using the fact that F is definable by a system of equations contractive in F's argument. To see this, note that in the base case $F^0(x) =_0 F^0(y)$, since every value is $=_0$ to every other value. Recall that $=_i$ means equal to depth i (where depth is the number of nested constructors) and that $F(z)$ is equivalent to a system of equations with occurrences of z embedded inside at least one constructor (contractiveness). Therefore, for the inductive step, it suffices to note that if $F^i(x) =_i F^i(y)$ (i.e., equal to a depth of i constructors) then $F(F^i(x)) =_{i+1} F(F^i(y))$ (i.e., equal to a depth of $i+1$ constructors).

Let d' and d'' be elements of $\mathcal{D}_0$ such that $F(d'') = d''$ and $F(d') \preceq d'$. It is easy to see by induction that $F^i(d'') = d''$ and, using monotonicity of F, that $F^i(d') \preceq d'$. Therefore, we have

$$d'' = F^i(d'') =_i F^i(d') \preceq d'$$

for all i. Hence, $d'' \preceq_i d'$ holds for all i. By definition of $\preceq$, it follows that $d'' \preceq d'$ and we conclude that the domain has type continuity. ■

5.2. Soundness

In this section, we extend variable elimination to recursive types. The first step is to extend the functions *Pos* and *Neg* to include type expressions that have defined variables (recall variables defined in E are denoted by δ):

Definition 18. Pos' and Neg' are the smallest sets of variables such that

1. If α is not defined in E, then $Pos'(\alpha/E) = \{\alpha\}$ and $Neg'(\alpha/E) = \emptyset$.
2. If $\delta = \tau$ is in E, then $Pos'(\delta/E) = Pos'(\tau/E) \cup \{\delta\}$ and $Neg'(\delta/E) = Neg'(\tau/E)$.
3. $Pos'(\bot/E) = Neg'(\bot/E) = \emptyset$
4. $Pos'(\top/E) = Neg'(\top/E) = \emptyset$
5. $Pos'(\tau_1 \to \tau_2/E) = Pos'(\tau_2/E) \cup Neg'(\tau_1/E)$
   and $Neg'(\tau_1 \to \tau_2/E) = Neg'(\tau_2/E) \cup Pos'(\tau_1/E)$

Let D be the set of E's defined variables.
Then define $Pos(\tau/E) = Pos'(\tau/E) - D$ and $Neg(\tau/E) = Neg'(\tau/E) - D$.

Note that *Pos* and *Neg* exclude defined variables while Pos' and Neg' include defined variables. Many functions satisfy these equations (so picking the smallest such sets is necessary to make Pos' and Neg' well-defined). For example, choosing

$$Pos(\delta/\delta = \delta \to \delta) = Neg(\delta/\delta = \delta \to \delta) = \{\alpha_4, \alpha_{29}\}$$

satisfies the equations, but the least solution is

$$Pos(\delta/\delta = \delta \to \delta) = Neg(\delta/\delta = \delta \to \delta) = \emptyset$$

Our results apply to the least solutions of the equations. It is easy to construct the least sets for *Pos* and *Neg* by adding variables only as necessary to satisfy the clauses of Definition 18.

At this point we digress to discuss the complexity of computing *Pos* and *Neg* sets for recursive types. Let the print representation of a system of type equations have size n and let the system have m type variables. Observe that the problem can be factored into m independent subproblems, one for each type variable. Focusing on a single variable α, the problem is to compute two bits for each subexpression E: whether $\alpha \in Pos'(E)$ and whether $\alpha \in Neg'(E)$. This subproblem can be solved in time linear in n, so to solve all m subproblems is $\mathcal{O}(mn)$.

We now explain how to decide $\alpha \in Pos'(E)$ and $\alpha \in Neg'(E)$ for every subexpression E in linear time. Define a graph with one node for each subexpression of the type and the associated system of equations. The graph has the following directed edges:

- There is an edge from each node for $X \to Y$ to the node for X. These edges are called *crossing* edges.
- There is an edge from each node for $X \to Y$ to the node for Y.
- For each node for δ there is an edge to the node for τ where $\delta = \tau$ is an equation of the system.

This graph has $\mathcal{O}(n)$ nodes and $\mathcal{O}(n)$ edges.

Let a be the node for variable α and let b be the node for some expression B. Assume there is a path from b to a in the graph. A path with an even number of crossings is *positive*; a path with an odd number of crossings is *negative*. It is easy to show $\alpha \in Pos'(B)$ if there is a positive path from b to a and $\alpha \in Neg'(B)$ if there is a negative path from b to a.

To compute the property for every subexpression efficiently, we reverse all the edges in the graph and perform a modified depth-first search from a, marking each node along the way as either positive or negative or both according to the marks of its predecessor and whether the edge being traversed is a crossing. Each edge may be visited at most twice (once for a positive path and once for a negative path) so the overall complexity is linear. This concludes the discussion of the complexity of computing *Pos* and *Neg* sets.

The following relationship between defined and regular variables is easy to show using Definition 18. The intuition is that if α is positive (resp. negative) in the definition of δ, then α is positive (resp. negative) in any position where δ appears positively, and negative (resp. positive) in any position where δ appears negatively. Note that since α is not a defined variable (and hence in *Pos* iff it is in Pos' and in *Neg* iff it is in Neg'), we could replace Pos' by *Pos* and Neg' by *Neg* in the following lemma. This remark does not apply to δ as defined variables are never in *Pos* and *Neg*.

LEMMA 10 If $\alpha \in Pos'(\delta/E)$ then

$$\delta \in Neg'(\tau/E) \;\Rightarrow\; \alpha \in Neg'(\tau/E)$$
$$\delta \in Pos'(\tau/E) \;\Rightarrow\; \alpha \in Pos'(\tau/E)$$

If $\alpha \in Neg'(\delta/E)$ then

$$\delta \in Neg'(\tau/E) \;\Rightarrow\; \alpha \in Pos'(\tau/E)$$
$$\delta \in Pos'(\tau/E) \;\Rightarrow\; \alpha \in Neg'(\tau/E)$$

LEMMA 11 Let τ be a recursive type expression and let $d_1, d_2 \in \mathcal{D}_0$ where $d_1 \preceq d_2$. Let θ be any assignment. If the semantic domain has contractive solutions and type continuity, then

1. if $\alpha \notin Pos(\tau)$, then $\theta[\alpha \leftarrow d_2](\tau) \preceq \theta[\alpha \leftarrow d_1](\tau)$.
2. if $\alpha \notin Neg(\tau)$, then $\theta[\alpha \leftarrow d_1](\tau) \preceq \theta[\alpha \leftarrow d_2](\tau)$.

Proof: Let $\tau = \tau_0/E$. The result is proven by induction on the number of equations in E with a sub-induction on the structure of τ_0. The sub-induction on τ_0's structure proceeds as in Lemma 1. The interesting case is the new base case where τ_0 is a defined variable δ_1 with $\delta_1 = \tau_1$ in E.

Assume $\tau = \delta_1/E$ where $\delta_1 = \tau_1$ is an equation in E. If $\alpha \in Pos(\delta_1/E)$ and $\alpha \in Neg(\delta_1/E)$, then the result is vacuously true. If $\alpha \notin Pos(\delta_1/E)$ and $\alpha \notin Neg(\delta_1/E)$, then let E' be those equations in E that do not contain α and do not (recursively) refer to a defined variable that contains α in its definition. Using Lemma 8, it can be shown that $\delta_1/E' \equiv \delta_1/E$, so it suffices to prove the result for δ_1/E'. Notice that α does not occur in δ_1/E'. It is easy to check that $\theta[\alpha \leftarrow d](\delta_1/E') = \theta(\delta_1/E')$ holds for all $d \in \mathcal{D}_0$ and the result follows.

Assume $\alpha \notin Pos(\delta_1/E)$ and $\alpha \in Neg(\delta_1/E)$. We claim $\delta_1 \notin Neg'(\tau_1/E)$. To see this, note that

	$\delta_1 \in Neg'(\tau_1/E) \wedge \alpha \in Neg'(\delta_1/E)$	
$\Rightarrow$	$\alpha \in Pos'(\tau_1/E)$	by Lemma 10
$\Rightarrow$	$\alpha \in Pos'(\delta_1/E)$	by Definition 18
$\Rightarrow$	$\alpha \in Pos(\delta_1/E)$	by Definition 18

which violates the assumption $\alpha \notin Pos(\delta_1/E)$. Therefore $\delta_1 \notin Neg'(\tau_1/E)$. Let E' be E with the equation $\delta_1 = \tau_1$ deleted. Now

	$\delta_1 \notin Neg'(\tau_1/E)$	
$\Rightarrow$	$\delta_1 \notin Neg'(\tau_1/E')$	see below
$\Rightarrow$	$\delta_1 \notin Neg(\tau_1/E')$	since $Neg(\tau_1/E') \subseteq Neg'(\tau_1/E')$

The second line follows because deleting equations from E can only decrease the least solutions of the equations for Pos' and Neg'.

Fix an assignment θ. For each $d_0 \in \mathcal{D}_0$, define $F_{d_0}(d) = \theta[\alpha \leftarrow d_0][\delta_1 \leftarrow d](\tau_1/E')$. It is clear that F_{d_0} is a definable operator. By the induction hypothesis, F_{d_0} is a monotonic operator. It is easy to see that $\alpha \notin Pos(\tau_1/E')$, so it also follows from the induction hypothesis that F is anti-monotonic in its subscript. More formally, if $d_1 \preceq d_2$, then $F_{d_2}(d) \preceq F_{d_1}(d)$ holds for every $d \in \mathcal{D}_0$.

Define a function h on $\mathcal{D}_0$ by

$$h(d_0) = \theta[\alpha \leftarrow d_0](\delta_1/E)$$

Now we have

$$\begin{array}{lll} & F_{d_0}(h(d_0)) & \\ = & \theta[\alpha \leftarrow d_0][\delta_1 \leftarrow h(d_0)](\tau_1/E') & \\ = & (\theta[\alpha \leftarrow d_0][\delta_1 \leftarrow h(d_0)])^{E'}(\tau_1) & \\ = & ((\theta[\alpha \leftarrow d_0])^{E})^{E'}(\tau_1) & \text{see below} \\ = & (\theta[\alpha \leftarrow d_0])^{E}(\tau_1) & \text{by part 1 of Lemma 7, since } E' \subseteq E \\ = & (\theta[\alpha \leftarrow d_0])^{E}(\delta_1) & \\ = & \theta[\alpha \leftarrow d_0](\delta_1/E) & \\ = & h(d_0) & \end{array}$$

To check that the fourth line follows, we would like to check that

$$(\theta[\alpha \leftarrow d_0][\delta_1 \leftarrow h(d_0)])^{E'} = ((\theta[\alpha \leftarrow d_0]^{E})^{E'}$$

It suffices to check that

$$\theta[\alpha \leftarrow d_0][\delta_1 \leftarrow h(d_0)](v) = (\theta[\alpha \leftarrow d_0])^{E}(v)$$

for all v not given definitions in E'. If v is not given a definition by E, then

$$\theta[\alpha \leftarrow d_0][\delta_1 \leftarrow h(d_0)](v) = \theta[\alpha \leftarrow d_0](v) = (\theta[\alpha \leftarrow d_0])^{E}(v)$$

If v is given a definition by E (but not by E'), then $v = \delta_1$; in this case,

$$\begin{array}{ll} & \theta[\alpha \leftarrow d_0][\delta_1 \leftarrow h(d_0)](v) \\ = & \theta[\alpha \leftarrow d_0][\delta_1 \leftarrow h(d_0)](\delta_1) \\ = & h(d_0) \\ = & \theta[\alpha \leftarrow d_0](\delta_1/E) \\ = & (\theta[\alpha \leftarrow d_0])^{E}(\delta_1) \\ = & (\theta[\alpha \leftarrow d_0])^{E}(v) \end{array}$$

Hence, the fourth line follows.

Thus, $F_{d_0}(h(d_0)) = h(d_0)$ holds for every $d_0 \in \mathcal{D}_0$. Let $d_1 \preceq d_2$. $F_{d_2}(h(d_2)) = h(d_2)$. $F_{d_2}(h(d_1)) \preceq F_{d_1}(h(d_1)) = h(d_1)$. By type continuity, $h(d_2) \preceq h(d_1)$ which is the desired result.

If $\alpha \in Pos(\delta_1/E)$ and $\alpha \notin Neg(\delta_1/E)$, then the proof is omitted since it is similar to the case where $\alpha \notin Pos(\delta_1/E)$ and $\alpha \in Neg(\delta_1/E)$. Like the previous case, F is monotonic in its argument; unlike the previous case, F is monotonic in its subscript. ■

COROLLARY 3 Assume the semantic domain has type continuity and contractive solutions.

1. If $\alpha \notin Pos(\tau/E)$, then $\theta((\tau/E)[\alpha \leftarrow \top]) \preceq \theta(\tau/E) \preceq \theta((\tau/E)[\alpha \leftarrow \bot])$ holds for all assignments θ.
2. If $\alpha \notin Neg(\tau/E)$, then $\theta((\tau/E)[\alpha \leftarrow \bot]) \preceq \theta(\tau/E) \preceq \theta((\tau/E)[\alpha \leftarrow \top])$ holds for all assignments θ.

The rest of the soundness results proceed as before. In particular, the Variable Elimination Procedure remains unaffected, except that it uses the new definitions of *Pos* and *Neg*. Just as in Section 4, we extend *Pos* and *Neg*:

$$\begin{aligned} Pos(\forall\alpha.\sigma) &= Pos(\sigma) - \{\alpha\} \\ Neg(\forall\alpha.\sigma) &= Neg(\sigma) - \{\alpha\} \end{aligned}$$

LEMMA 12 If σ is a quantified recursive type expression and the semantic domain has type continuity and contractive solutions, then

$$\begin{aligned} \alpha \notin Neg(\sigma) &\Rightarrow \forall\alpha.\sigma \equiv \sigma[\alpha \leftarrow \bot] \\ \alpha \notin Pos(\sigma) &\Rightarrow \forall\alpha.\sigma \equiv \sigma[\alpha \leftarrow \top] \end{aligned}$$

Proof: Same as the proof for Lemma 2. ■

EXAMPLE: Let $\sigma = (\delta_3 \to \delta_3) \to (\delta_2 \to \delta_1)/\delta_1 = \alpha_1 \to \delta_1 \wedge \delta_2 = \alpha_2 \to \delta_2 \wedge \delta_3 = \alpha_3 \to \delta_3$. Note that $Pos(\sigma) = \{\alpha_2, \alpha_3\}$ and $Neg(\sigma) = \{\alpha_1, \alpha_3\}$. Assuming that the semantic domain has contractive solutions and type continuity, Lemma 12 allows us to conclude that

$$\begin{aligned} & \forall\alpha_1\forall\alpha_2\forall\alpha_3.((\delta_3 \to \delta_3) \to (\delta_2 \to \delta_1)/\ \delta_1 = \alpha_1 \to \delta_1 \wedge \delta_2 = \alpha_2 \to \delta_2 \wedge \delta_3 = \alpha_3 \to \delta_3) \\ \equiv\ & \forall\alpha_3.((\delta_3 \to \delta_3) \to (\delta_2 \to \delta_1)/\ \delta_1 = \top \to \delta_1 \wedge \delta_2 = \bot \to \delta_2 \wedge \delta_3 = \alpha_3 \to \delta_3) \end{aligned}$$

□

The next example shows that the assumption of type continuity is needed in the proof of Lemma 12.

EXAMPLE: Consider the type expression $\forall\alpha.(\delta/\delta = \alpha \to \delta)$. If Lemma 12 holds, then we have

$$\begin{aligned} & \delta/\delta = \top \to \delta \\ \equiv\ & \forall\alpha.(\delta/\delta = \alpha \to \delta) && \text{by Lemma 12} \\ \preceq\ & (\delta/\delta = \bot \to \delta) && \text{since } \bot \text{ is an instance of } \alpha \end{aligned}$$

Let $\delta_0 = \top \to \delta_0$ and $\delta_1 = \bot \to \delta_1$ be elements of the semantic domain. Any semantic domain in which it is *not* the case that $\delta_0 \preceq \delta_1$ serves as a counterexample to the conclusion of Lemma 12.

Take the semantic domain to be the set of regular trees and define $x \preceq'_0 y$ to hold iff $x = y$. Let $x \preceq'_{i+1} y$ hold iff $x = \bot$, $y = \top$, or $\exists x_1, x_2, y_1, y_2(x = x_1 \to x_2 \wedge y = y_1 \to y_2 \wedge y_1 \preceq'_i x_1 \wedge x_2 \preceq'_i y_2)$. Let $x \preceq' y$ hold iff $x \preceq'_i y$ for some i. Next, notice that $\delta_0 \preceq'_{i+1} \delta_1$ iff $\top \to \delta_0 \preceq'_{i+1} \bot \to \delta_1$ iff $\bot \preceq'_i \top \wedge \delta_0 \preceq'_i \delta_1$. It is easy to see by induction that $\delta_0 \not\preceq'_i \delta_1$ is true for all i. Hence, $\delta_0 \not\preceq' \delta_1$. Thus, the conclusion of Lemma 12 does not hold for this semantic domain.

This semantic domain has contractive solutions, standard function types, and standard glb types. What it lacks is type continuity, and it is instructive to see why. Consider the two definable operators:

$$\begin{aligned} F_\bot(d) &= [\alpha \leftarrow d](\bot \to \alpha) \\ F_\top(d) &= [\alpha \leftarrow d](\top \to \alpha) \end{aligned}$$

Let $\top \to \top \to \ldots$ be the infinite regular tree where $\top$ appears in the domain of every "$\to$". Observe that

$$F_\top(\top \to \top \to \ldots) = \top \to \top \to \ldots$$

Note that for all d we have $F_\top(d) \preceq' F_\bot(d)$. In particular,

$$F_\top(\bot \to \bot \to \ldots) \preceq' F_\bot(\bot \to \bot \to \ldots) = \bot \to \bot \to \ldots$$

If the domain had type continuity, it would follow that

$$\top \to \top \to \ldots \preceq' \bot \to \bot \to \ldots$$

As shown above, this relation does not hold, so therefore the domain does not have type continuity.

□

As discussed at the beginning of this section, type continuity is needed to guarantee that the finite computation performed by *Pos* and *Neg* is consistent with the orderings on all finite and infinite trees. Example 18 shows how the problem arises when $x \not\preceq' y$, but $x \preceq_i y$ for all i (where $\preceq_i$ is the relation used in Lemma 9 to define the usual ordering on regular trees). Thus, in contrast to the case of simple expressions where no additional assumptions on the semantic domain are needed for soundness, type continuity is needed to prove soundness for recursive type expressions.

We remark that the definition of a reduced quantified recursive type expression is the same as a reduced quantified simple type expression (Definition 6).

THEOREM 7 Let σ be any quantified recursive type expression. If the semantic domain has type continuity and contractive solutions, then $\sigma \equiv VEP(\sigma)$ and $VEP(\sigma)$ is a reduced recursive type expression.

Proof: Follows easily from Lemma 12. ■

As with simple type expressions, an irredundant quantified type expression is one such that all equivalent quantified type expressions have at least as many quantified variables. Note that this definition does not say anything about the number of defined variables. It is conceivable (although we will see that this is not the case under our usual assumptions) that an irredundant type might require many more defined variables.

THEOREM 8 If the semantic domain has type continuity and contractive solutions, then every irredundant recursive type expression is reduced.

Proof: Same as the proof of Theorem 3. ■

5.3. Completeness

In this section, we face concerns similar to those found in Section 4.2.1.

Definition 19. Let S be a set of constraints over unquantified recursive type expressions. Define B to be a function on sets of constraints such that $B(S)$ is the smallest set of constraints where the following all hold. These clauses are to be applied in order, with the earliest one that applies taking precedence. The variables s and t refer to unquantified simple type expressions.

1. $B(\emptyset) = \emptyset$
2. If t is $\top$, $\bot$, or a regular variable, then $B(\{t/E \preceq t/E'\} \cup S) = B(S)$.
3. $B(\{s_1 \to s_2/E \preceq t_1 \to t_2/E'\} \cup S) = B(\{t_1/E' \preceq s_1/E, s_2/E \preceq t_2/E'\} \cup S)$.
4. If $\delta = \tau$ is in E, then $B(\{\delta/E \preceq t\} \cup S) = B(\{\tau/E \preceq t\} \cup S)$.
5. If $\delta = \tau$ is in E', then $B(\{s \preceq \delta/E'\} \cup S) = B(\{s \preceq \tau/E'\} \cup S)$.
6. Otherwise, $B(\{s/E \preceq t/E'\} \cup S) = \{s \preceq t\} \cup B(S)$.

LEMMA 13 Assume that $\mathcal{D}$ is a semantic domain with contractive solutions, standard function types, and standard glb types. If θ is a solution of $\{t_1/E \preceq t_2/E'\}$, then it is a solution of $B(\{t_1/E \preceq t_2/E'\})$.

Proof: The proof is very similar to the proof of Lemma 3 and so is omitted. The most important new case is Part 6 of Definition 19. In this clause, note that s and t must be regular variables of E and E' respectively. Thus $B(S)$ does not mention any defined variables. This observation is needed to show that if $\theta^E(t_1) \preceq \theta^{E'}(t_2)$ then θ is a solution of $B(\{t_1/E \preceq t_2/E'\})$. ■

LEMMA 14 Let $\forall V_1.\tau_1/E_1$ and $\forall V_2.\tau_2/E_2$ be compatible recursive type expressions. If the semantic domain has contractive solutions, standard function types, and standard glb types, then $B(\{\tau_1/E_1 \preceq \tau_2/E_2\})$ is a (V_1, V_2)-miniscule system of constraints.

Proof: Let $B(\{\tau_1/E_1 \preceq \tau_2/E_2\}) = \{s_1 \preceq t_1, \ldots, s_n \preceq t_n\}$. We show that $B(\{\tau_1/E_1 \preceq \tau_2/E_2\})$ satisfies the conditions of Definition 10.

1. By compatibility V_1 and V_2 are disjoint sets of variables.
2. By Part 3 of Definition 19 at most one of s_i and t_i is a $\to$ expression.
3. Consider a constraint $t \preceq t$. Constraints of the form $\bot \preceq \bot$, $\top \preceq \top$, and $\alpha \preceq \alpha$ are eliminated by Part 2 of Definition 19, constraints $t \to t' \preceq t \to t'$ are eliminated by Part 3, and constraints $\delta \preceq \delta$ are eliminated by Parts 4 and 5. Therefore, for all t, we have $t \preceq t$ is not in $B(\{\tau_1/E_1 \preceq \tau_2/E_2\})$.
4. Same as Part 4 of the proof of Lemma 5 (but using Definition 19).

5. Proof similar to the previous step.

6. Let θ be any assignment. We must show that there is an assignment θ' such that $\theta(v) = \theta'(v)$ for all $v \notin V_1$ and $\theta'(s_i) \preceq \theta'(t_i)$ holds for all $i \leq n$. Since

   $$\theta(\forall V_1.\tau_1/E_1) \preceq \theta(\forall V_2.\tau_2/E_2),$$

   it follows that

   $$\theta(\forall V_1.\tau_1/E_1) \preceq \theta(\tau_2/E_2).$$

   Since the semantic domain has standard glb types, it follows that $\theta'(\tau_1/E_1) \preceq \theta(\tau_2/E_2)$ holds for some θ' that agrees with θ except possibly on V_1. Since no variable in V_1 occurs in τ_2/E_2, we know $\theta'(\tau_1/E_1) \preceq \theta'(\tau_2/E_2)$. By Lemma 13, it follows that θ' is a solution to $B(\tau_1/E_1 \preceq \tau_2/E_2)$.

7. Similar to the previous step with the roles of τ_1 and τ_2 reversed.

■

THEOREM 9 If the semantic domain has contractive solutions, standard glb types, and standard function types, then any two equivalent reduced recursive type expressions have the same number of quantified variables.

Proof: Let σ' and σ'' be two equivalent reduced recursive type expressions. If necessary, α-convert σ' to $\sigma_1 = \forall V_1.\tau_1/E$ and α-convert σ'' to $\sigma_2 = \forall V_2.\tau_2/E'$ in such a way that σ_1 and σ_2 are (V_1, V_2) compatible. By Lemma 14, $B(\tau_1/E \preceq \tau_2/E')$ is a (V_1, V_2)-miniscule system of constraints. By Lemma 4, $B(\tau_1/E \preceq \tau_2/E')$ is a (V_1, V_2)-convertible system of constraints, which implies that $|V_1| = |V_2|$.

■

Unlike the case of simple expressions, two equivalent reduced types need not be α-equivalent. For example, consider a semantic domain that has contractive solutions. Let $\delta_0 = \delta_0 \to \delta_0$ and $\delta_1 = (\delta_1 \to \delta_1) \to (\delta_1 \to \delta_1)$. These two types exist since the domain has contractive solutions. By substituting $(\delta_0 \to \delta_0)$ in for δ_0, we obtain $\delta_0 = (\delta_0 \to \delta_0) \to (\delta_0 \to \delta_0)$. Since the domain has contractive solutions, it follows that $\delta_0 = \delta_1$. Clearly, the type expressions $\delta_0/\delta_0 = \delta_0 \to \delta_0$ and $\delta_1/\delta_1 = (\delta_1 \to \delta_1) \to (\delta_1 \to \delta_1)$ are not α-equivalent.

THEOREM 10 If the semantic domain has contractive solutions, type continuity, standard function types, and standard glb types, then a recursive type expression is reduced iff it is irredundant.

Proof: The if-direction follows from Theorem 8. To prove the only-if direction, let σ be a reduced recursive type expression with the goal of proving that σ is irredundant. Let σ' be an irredundant recursive type expression that is equivalent to σ. (Such a σ' can always be found by picking it to be a type expression equivalent

to σ with the smallest possible number of quantified variables.) By Theorem 8, σ' is reduced. By Theorem 9, σ and σ' have the same number of quantified variables. Hence, σ is irredundant as desired. ■

THEOREM 11 Let σ be quantified recursive type expression. If the semantic domain has contractive solutions, type continuity, standard glb types, and standard function types, then $VEP(\sigma)$ is an irredundant recursive type expression equivalent to σ.

Proof: Follows easily from Theorem 7 and Theorem 10. ■

6. Intersection and Union Types

In this section we extend our results to type languages with union and intersection types. This is the first point at which the technique of eliminating variables that appear solely in monotonic or anti-monotonic positions is sound but not complete.

6.1. Preliminaries

As a first step union and intersection types are added to simple type expressions.

Definition 20. Extended type expressions are generated by the grammar

$$\tau ::= \alpha \mid \top \mid \bot \mid \tau_1 + \tau_2 \mid \tau_1 \cdot \tau_2 \mid \tau_1 \to \tau_2$$

Extended quantified types are adapted in the obvious way to use extended type expressions instead of simple type expressions. The operations $+$ and $\cdot$ are interpreted as least-upper bound and greatest-lower bound, respectively. To give meaning to extended type expressions an assumption is needed about the upper and lower bounds that exist in the domain.

Definition 21. A semantic domain $\mathcal{D} = (\mathcal{D}_0, \mathcal{D}_1, \preceq, \sqcap)$ has *standard upper and lower bounds* if every pair of elements $\tau_1, \tau_2 \in \mathcal{D}_0$ have a least upper bound $\tau_1 \sqcup \tau_2$ and a greatest lower bound $\tau_1 \sqcap \tau_2$ in $\mathcal{D}_0$.

Note that requiring $\tau_1 \sqcap \tau_2$ exist is different from having standard glb types, as standard glb types are glb's of (potentially) infinite sets in $\mathcal{D}_1$.

PROPOSITION 2 The Standard Model (Example 1) and Regular Tree Model (Lemma 9) both have standard upper and lower bounds. Furthermore, in both these models, $x_1 \to y_1 \sqcup x_2 \to y_2 = x_1 \sqcap x_2 \to y_1 \sqcup y_2$ and $x_1 \to y_1 \sqcap x_2 \to y_2 = x_1 \sqcup x_2 \to y_1 \sqcap y_2$.

Proof: First note that for every $x \in \mathcal{D}_0$ it is the case that $x = \top$, $x = \bot$, or $x = x_1 \to x_2$ for some x_1 and x_2. Also note that both models have standard function types.

We must show that $x \sqcup y$ and $x \sqcap y$ exist for all $x, y \in \mathcal{D}_0$. It is easy to check that the following equations cover all possibilities:

$$\begin{array}{rcl} \top \sqcup x & = & \top \\ \bot \sqcup x & = & x \\ \bot \sqcap x & = & \bot \\ \top \sqcap x & = & x \\ x_1 \to y_1 \sqcup x_2 \to y_2 & = & x_1 \sqcap x_2 \to y_1 \sqcup y_2 \end{array} \qquad \begin{array}{rcl} x \sqcup \top & = & \top \\ x \sqcup \bot & = & x \\ x \sqcap \bot & = & \bot \\ x \sqcap \top & = & x \\ x_1 \to y_1 \sqcap x_2 \to y_2 & = & x_1 \sqcup x_2 \to y_1 \sqcap y_2 \end{array}$$

The eight equations on the first four lines are easy to verify. To justify the equation $x_1 \to y_1 \sqcup x_2 \to y_2 = x_1 \sqcap x_2 \to y_1 \sqcup y_2$, note that

$$\begin{array}{rcl} x_1 \sqcap x_2 & \preceq & x_1 \\ x_1 \sqcap x_2 & \preceq & x_2 \\ y_1 & \preceq & y_1 \sqcup y_2 \\ y_2 & \preceq & y_1 \sqcup y_2 \end{array}$$

from which it follows that $x_1 \sqcap x_2 \to y_1 \sqcup y_2$ is an upper bound of both $x_1 \to y_1$ and $x_2 \to y_2$. Let $a \to b$ be any other upper bound of $x_1 \to y_1$ and $x_2 \to y_2$. Since the domain has standard function types, we have $a \preceq x_1$ and $a \preceq x_2$, so $a \preceq x_1 \sqcap x_2$. Similarly $y_1 \sqcup y_2 \preceq b$. Therefore, $x_1 \sqcap x_2 \to y_1 \sqcup y_2$ is the least upper bound. The justification of the last equation is symmetric.

Finally, we need to show that the above argument is sufficient. In the case of the Standard Model, the above is sufficient to push the lower and upper bounds to the leaves where they can be eliminated. In the Regular Tree Model, when taking bounds for α and α' with associated definitions E and E' respectively, create a new variable for each bound of a defined variable from E and a defined variable from E'. For each such newly created variable, form its definition by taking the bound of the right hand sides and moving the bound inside according to the above procedure. This forms a new set of equations. The details of this construction are left to the reader. ■

Given an assignment θ, the meanings of the new type operations are:

$$\begin{array}{rcl} \theta(\tau_1 + \tau_2) & = & \theta(\tau_1) \sqcup \theta(\tau_2) \\ \theta(\tau_1 \cdot \tau_2) & = & \theta(\tau_1) \sqcap \theta(\tau_2) \end{array}$$

6.2. Soundness for Non-Recursive Types

We first extend *Pos* and *Neg* to include the new operations.

$$\begin{array}{rcl} Pos(\tau_1 + \tau_2) & = & Pos(\tau_1) \cup Pos(\tau_2) \\ Pos(\tau_1 \cdot \tau_2) & = & Pos(\tau_1) \cup Pos(\tau_2) \\ Neg(\tau_1 + \tau_2) & = & Neg(\tau_1) \cup Neg(\tau_2) \\ Neg(\tau_1 \cdot \tau_2) & = & Neg(\tau_1) \cup Neg(\tau_2) \end{array}$$

We can now restate the basic lemma needed to prove soundness for the non-recursive case.

LEMMA 15 Let τ be any extended simple type expression. Let $d_1, d_2 \in \mathcal{D}_0$ where $d_1 \preceq d_2$. If the domain has standard upper and lower bounds, then

1. If $\alpha \notin \mathit{Pos}(\tau)$, then $\theta(\tau[\alpha \leftarrow d_2]) \preceq \theta(\tau[\alpha \leftarrow d_1])$ holds for all assignments θ.
2. If $\alpha \notin \mathit{Neg}(\tau)$, then $\theta(\tau[\alpha \leftarrow d_1]) \preceq \theta(\tau[\alpha \leftarrow d_2])$ holds for all assignments θ.

Proof: This proof is by induction on the structure of τ and is an easy extension of the proof of Lemma 1. Let $d_1, d_2 \in \mathcal{D}_0$ where $d_1 \preceq d_2$, and let θ be any assignment. There are two new cases:

- Let $\tau = \tau_1 + \tau_2$. Assume $\alpha \notin \mathit{Pos}(\tau)$. By the definition of *Pos*, we know

  $$\alpha \notin \mathit{Pos}(\tau_1) \cup \mathit{Pos}(\tau_2)$$

  and therefore

  $$\begin{aligned} \theta[\alpha \leftarrow d_2](\tau_1) &\preceq \theta[\alpha \leftarrow d_1](\tau_1) \\ \theta[\alpha \leftarrow d_2](\tau_2) &\preceq \theta[\alpha \leftarrow d_1](\tau_2) \end{aligned}$$

  follow by induction. The relationships still hold if the right-hand sides are made larger, so

  $$\begin{aligned} \theta[\alpha \leftarrow d_2](\tau_1) &\preceq \theta[\alpha \leftarrow d_1](\tau_1) \sqcup \theta[\alpha \leftarrow d_1](\tau_2) \\ \theta[\alpha \leftarrow d_2](\tau_2) &\preceq \theta[\alpha \leftarrow d_1](\tau_1) \sqcup \theta[\alpha \leftarrow d_1](\tau_2) \end{aligned}$$

  Combining these two inequalities we get

  $$\theta[\alpha \leftarrow d_2](\tau_1) \sqcup \theta[\alpha \leftarrow d_2](\tau_2) \preceq \theta[\alpha \leftarrow d_1](\tau_1) \sqcup \theta[\alpha \leftarrow d_1](\tau_2)$$

  The proof for the subcase $\alpha \notin \mathit{Neg}(\tau)$ is similar.

- Let $\tau = \tau_1 \cdot \tau_2$. This case is very similar to the previous one, with $\sqcap$ substituted for $\sqcup$.

■

An inspection of the results from Section 4.1 shows that the proofs of Lemma 2 and Theorem 2 depend only on Lemma 1 and not on a particular language of type expressions. Therefore, by Lemma 15, it is immediate that Procedure 1 is a sound variable elimination procedure for extended simple types in domains with standard upper and lower bounds.

While variable elimination is sound for extended simple type expressions, it is not complete.

EXAMPLE: In either the Standard Model or Regular Tree Model we have

$$\forall\alpha.(\alpha \to \alpha) + \top \equiv \top$$

Clearly, the first type is reduced and not irredundant. □

Similarly, $\forall\alpha.(\alpha \to \alpha)\cdot \perp \equiv \perp$. In general, the *Pos* and *Neg* computations overestimate the set of positive and negative variables for expressions $\tau_1 + \tau_2$ where $\theta(\tau_1) \preceq \theta(\tau_2)$ for all θ (and similarly for $\cdot$).

A subtler source of incompleteness arises from interaction between universal quantification and unions and intersections.

EXAMPLE:

$$\begin{aligned}
& \forall\alpha,\beta.\alpha\cdot\beta \to \alpha\cdot\beta \\
= \; & \sqcap\{\theta[\alpha \leftarrow x_1, \beta \leftarrow x_2](\alpha\cdot\beta \to \alpha\cdot\beta) | x_1, x_2 \in \mathcal{D}_0\} \\
& \text{by Proposition 1} \\
= \; & \sqcap\{x_1 \sqcap x_2 \to x_1 \sqcap x_2) | x_1, x_2 \in \mathcal{D}_0\} \\
= \; & \sqcap\{x \to x | x \in \mathcal{D}_0\} \\
& \text{since } \{x_1 \sqcap x_2 | x_1, x_2 \in \mathcal{D}_0\} = \mathcal{D}_0 \\
= \; & \forall\alpha.\alpha \to \alpha
\end{aligned}$$

□

Note that there is no explicit relationship between α and β in the type. The relationship follows from the fact that the variables are always used together and the universal quantification.

6.3. Improvements

We do not know a complete version of the Variable Elimination Procedure in the presence of union and intersection types. In this section we briefly illustrate some heuristic improvements that have been useful in practice [1, 9]. As illustrated in Section 6.2, redundant intersections and unions are a significant source of incompleteness. This suggests the following procedure:

Procedure 12 (Extended Variable Elimination Procedure (EVEP)) Let σ be an extended quantified type.

1. Let σ_1 be the result of replacing any subexpression $\tau_1 + \tau_2$ in σ by τ_2 if $\theta(\tau_1) \preceq \theta(\tau_2)$ for all assignments θ.
2. Let σ_2 be the result of replacing any subexpression $\tau_1 \cdot \tau_2$ in σ' by τ_2 if $\theta(\tau_2) \preceq \theta(\tau_1)$ for all assignments θ.
3. Let $\sigma_3 = VEP(\sigma_2)$.

4. Halt if no variables are eliminated in (3); the result is σ_3. Repeat (1)-(3) on σ_3 otherwise.

Note that deciding whether a type is equivalent to $\top$ or $\bot$ in all assignments is not necessarily easy, depending on the expressiveness of the type language under consideration.

The interesting aspect of Procedure 12 is that iterating the elimination of intersections, unions, and variables is necessary, as the following example shows:

$$\begin{array}{ll} & \forall\alpha,\beta.\beta \to (\alpha \cdot \beta) \\ \equiv & \forall\beta.\beta \to (\bot \cdot \beta) \qquad \text{since } \alpha \text{ is not negative} \\ \equiv & \forall\beta.\beta \to \bot \qquad \text{since } \bot \cdot \tau = \bot \\ \equiv & \top \to \bot \qquad \text{since } \beta \text{ is not positive} \end{array}$$

Since each iteration but the last of the Extended Variable Elimination Procedure eliminates at least one variable, the complexity is at worst the product of the number of quantified variables $\mathcal{O}(m)$ and the cost of recomputing the *Pos* and *Neg* sets $\mathcal{O}(mn)$, for a total cost of $\mathcal{O}(m^2 n)$.

6.4. Soundness and Incompleteness for Recursive Types

In this section we consider extended recursive types.

Definition 22. An *extended recursive type* has the form

$$\tau / \bigwedge_{1 \le i \le n} \delta_i = \tau_i$$

where the equations are contractive and $\tau, \tau_1, \ldots, \tau_n$ are extended type expressions (i.e., with unions and intersections; see Definition 20).

It will come as no surprise that, in addition to standard upper and lower bounds, the domain must have contractive solutions and type continuity for variable elimination to be sound for extended recursive types. An inspection of the statement and proof of Lemma 11 shows that it does not depend on a particular definition of type, but only on type continuity and soundness of the non-recursive case. Thus, adding the hypothesis that the domain has standard upper and lower bounds to Lemma 11, and substituting Lemma 15 for Lemma 1 in the proof of the lemma, gives a proof of soundness for variable elimination on extended recursive types.

Because extended simple types are a subset of the extended recursive types and the VEP is incomplete for extended simple types, it follows that variable elimination is incomplete for extended recursive types.

7. Constrained Type Expressions

This section presents results for types with polymorphism and subtyping constraints, which is also called *bounded polymorphism* or *constrained types*. This language is the most general that we consider.

7.1. Preliminaries

We begin by defining a type language with subsidiary subtyping constraints. We present the definitions and proofs as though "$\rightarrow$" were the only constructor but the results apply more generally and we leave it to the reader to fill in the details for other constructors.

Definition 23. An *(unquantified) constrained type expression* has the form τ_0/C where C is a finite set of constraints of the form

$$\begin{array}{ccc} \tau_1 & \preceq & \tau_1' \\ \ldots & & \ldots \\ \tau_n & \preceq & \tau_n' \end{array}$$

where τ_i and τ_i' are unquantified simple type expressions for all $1 \leq i \leq n$.

Unlike the case of recursive types, note that Definition 23 makes no distinction between "regular" and "defined" variables—all variables are regular.

Definition 24. Let θ be any assignment. Then

1. $\theta(\tau/C) = \theta(\tau)$ provided that θ is a solution of C.

2. $\theta(\forall\delta_1,\ldots,\delta_n.\tau/C) =$

$$\sqcap\{\theta[\delta_1 \leftarrow x_1,\ldots,\delta_n \leftarrow x_n](\tau/C) | x_1,\ldots,x_n \in \mathcal{D}_0 \text{ and } \theta[\delta_1 \leftarrow x_1,\ldots,\delta_n \leftarrow x_n] \text{ is a solution of } C\} \quad (1)$$

The meaning of an unquantified constrained type τ/C under assignment θ is undefined unless θ is a solution of C. Furthermore, the $\sqcap$ operation in the meaning of a quantified constrained type under assignment θ is restricted to those modifications of θ that satisfy the constraints. It is easy to see that constrained types are a generalization of recursive types, because any recursive type

$$\forall\alpha_1,\ldots,\alpha_m.\tau/\ \delta_1 = \tau_1 \rightarrow \tau_1' \wedge \ldots \wedge \delta_n = \tau_n \rightarrow \tau_n'$$

can be written as a constrained type

$$\forall\alpha_1,\ldots,\alpha_m.\tau/\ \delta_1 \preceq \tau_1 \to \tau_1' \wedge \delta_1 \succeq \tau_1 \to \tau_1' \wedge \ldots \wedge \delta_n \preceq \tau_n \to \tau_1' \wedge \delta_n \succeq \tau_n \to \tau_1'$$

It is also worth noting that it is well-defined for a quantified constrained type to have an inconsistent system of constraints. For example, if $C = \top \preceq \delta \preceq \bot$, then

$$\begin{aligned} & \theta(\forall\delta.\tau/C) \\ = \ & \sqcap\{\theta[\delta \leftarrow x](\tau/C) | x \in \mathcal{D}_0 \text{ and } \theta[\delta \leftarrow x] \text{ is a solution of } C\} \\ = \ & \sqcap\{\} \\ = \ & \top \end{aligned}$$

An important feature of constrained types is that the constraints may have multiple lower (or upper) bounds for a single variable, such as

$$\tau_1 \to \tau_2 \preceq \alpha \ \wedge \ \tau_3 \to \tau_4 \preceq \alpha \ \wedge \ \beta \preceq \gamma \ \wedge \ \beta \preceq \tau_5 \to \tau_6$$

In any solution of these constraints, α must be an upper bound of $\tau_1 \to \tau_2$ and $\tau_3 \to \tau_4$, and β must be a lower bound of γ and $\tau_5 \to \tau_6$.

To give algorithms for eliminating quantified variables from constrained types, it is necessary to characterize the solutions of constraints. To minimize the number of new concepts needed to explain the algorithms in the case of constrained types, we build on the results of Section 5 by characterizing solutions of constraints in terms of equations.

Definition 25. A system C of constraints is *fully closed* iff

$$\begin{aligned} \tau_1 \preceq \tau_0 \in C \wedge \tau_0 \preceq \tau_2 \in C \ &\Rightarrow \ \tau_1 \preceq \tau_2 \in C \\ \tau_1 \to \tau_2 \preceq \tau_3 \to \tau_4 \in C \ &\Rightarrow \ \tau_3 \preceq \tau_1 \in C \wedge \tau_2 \preceq \tau_4 \in C \\ \tau_1 \to \tau_2 \preceq \bot \in C \ &\Rightarrow \ \top \preceq \bot \in C \\ \top \preceq \tau_1 \to \tau_2 \in C \ &\Rightarrow \ \top \preceq \bot \in C \end{aligned}$$

A system is *closed* iff it can be obtained from a fully closed system by the deletion of some subset of the trivial constraints $\tau \preceq \tau$. A closed system C is *consistent* iff $\top \preceq \bot \notin C$.

For example, the system $\{\delta_1 \to \delta_2 \preceq \delta_3, \delta_3 \preceq \delta_1 \to \delta_2, \delta_1 \to \delta_2 \preceq \delta_1 \to \delta_2, \delta_1 \preceq \delta_1, \delta_2 \preceq \delta_2, \delta_3 \preceq \delta_3\}$ is fully closed (and hence closed) whereas the system $\{\delta_1 \to \delta_2 \preceq \delta_3, \delta_3 \preceq \delta_1 \to \delta_2\}$ is closed but not fully closed.

Definition 25 is taken from [8]. Intuitively, closing a system of constraints C is equivalent to solving C, and if the closed system has no inconsistent constraints, then it has solutions. Instead of asserting that closed consistent systems have solutions directly, we characterize those solutions in terms of equations. In this section, we do not assume the existence of standard function types. However, if a system is not closed, standard function types are generally required for the closure to be equivalent.

Definition 26. We use Σ and Π to abbreviate multiple type unions and intersections, respectively. Let C be a closed consistent system of constraints. Let the variables of C be $\delta_1, \ldots, \delta_n$. For each variable δ_i appearing in C, define

$$\begin{aligned} L^C_{\delta_i} &= \bot + \Sigma\{\tau | \tau \preceq \delta_i \in C \text{ and if } \tau \text{ is a variable } \delta_j, \text{ then } j < i\} \\ U^C_{\delta_i} &= \top \cdot \Pi\{\tau | \tau \succeq \delta_i \in C \text{ and if } \tau \text{ is a variable } \delta_j, \text{ then } j < i\} \end{aligned}$$

Let $\alpha_1, \ldots, \alpha_n$ be fresh variables. Define a system of equations E_C for C:

$$\bigwedge_{1 \le i \le n} \delta_i = L^C_{\delta_i} + (\alpha_i \cdot U^C_{\delta_i})$$

The intuition behind Definition 26 is that any solution for the equation for δ_i ranges between $L^C_{\delta_i}$ (when $\alpha_i = \bot$) and $U^C_{\delta_i}$ (when $\alpha_i = \top$). For example, consider the system C of constraints

$$\delta_1 \preceq \delta_2 \ \wedge \ \top \to \bot \preceq \delta_1 \ \wedge \ \delta_1 \preceq \bot \to \top$$

Closing this system gives

$$\begin{array}{l} \delta_1 \preceq \delta_2 \ \wedge \ \top \to \bot \preceq \delta_1 \ \wedge \ \delta_1 \preceq \bot \to \top \ \wedge \\ \top \to \bot \preceq \delta_2 \ \wedge \ \top \to \bot \preceq \bot \to \top \ \wedge \ \bot \preceq \top \end{array} \tag{2}$$

which is consistent. The equations E_C are

$$\begin{aligned} \delta_1 &= (\top \to \bot) + (\alpha_1 \cdot (\bot \to \top)) \\ \delta_2 &= ((\top \to \bot) + \delta_1) + (\alpha_2 \cdot \top) \end{aligned}$$

This example shows that the equations E_C are not necessarily contractive, since δ_1 appears outside of a constructor in the equation for δ_2. However, E_C is always equivalent to a contractive system of equations.

LEMMA 16 Let E_C be a system of equations for a closed consistent system of constraints C. Then there is a system of constraints E'_C that is contractive such that E_C and E'_C have the same solutions.

Proof: Examination of L_{δ_i} and U_{δ_i} in Definition 26 shows that if δ_j occurs outside of a $\to$ expression in the equation for δ_i, then $j < i$. We show by induction on i how to construct a contractive equation for δ_i. The equation for δ_1 has no variable δ_k outside of a $\to$ expression, so the equation for δ_1 is already contractive. Assume that $\delta_1, \ldots, \delta_{i-1}$ have contractive equations. Any variable δ_j outside of a $\to$ expression in the equation for δ_i can be eliminated by substituting the right-hand side of an equation for δ_j. Because $j < i$, we can choose a contractive equation for δ_j, in which case the resulting equation for δ_i is also contractive. ■

Applying Lemma 16 to the example system of equations above, the contractive system is

$$\begin{aligned}\delta_1 &= (\top \to \bot) + (\alpha_1 \cdot (\bot \to \top)) \\ \delta_2 &= ((\top \to \bot) + (\top \to \bot) + (\alpha_1 \cdot (\bot \to \top))) + (\alpha_2 \cdot \top)\end{aligned}$$

To prove that a consistent closed system has a solution, it is helpful to define two additional notions. A *visible expression* in a system of constraints is one of the top-level "$\to$" arms of either an upper or lower bound. A system of constraints is *flat* iff every visible expressions is $\top$, $\bot$, or a variable. Thus, for example, in the constraint system $\{\delta_0 \to \delta_1 \preceq (\delta_2 \to \delta_3) \to \delta_4, \delta_5 \preceq \delta_6 \to \top\}$, there are six visible expressions: δ_0, δ_1, $\delta_2 \to \delta_3$, δ_4, δ_6, and $\top$; this system is not flat.

LEMMA 17 Assume the domain has standard upper and lower bounds. Let C be a flat, closed, consistent system of constraints. If θ is a solution of E_C, then θ is a solution of C.

Proof: Let the variables that occur in C be $\delta_1, \ldots, \delta_n$. In this proof, we use t to denote $\top$, $\bot$, or a variable; τ is used to denote an arbitrary expression. Note that if τ is a term in $L^C_{\delta_i}$, then $\theta(\tau) \preceq \theta(L^C_{\delta_i}) \preceq \theta(L^C_{\delta_i} + \alpha_i \cdot U^C_{\delta_i}) = \theta(\delta_i)$; hence, $\theta(\tau) \preceq \theta(\delta_i)$. Thus, to show that $\theta(\tau) \preceq \theta(\delta_i)$ it suffices to show that τ is a term in $L^C_{\delta_i}$. To show that $\theta(\delta_i) \preceq d$ for some d in the semantic domain, it suffices to show that $\theta(L^C_{\delta_i}) \preceq d$ (for which it suffices to show that $\theta(\tau) \preceq d$ for every term τ in $L^C_{\delta_i}$) and $U^C_{\delta_i} \preceq d$ (for which it suffices to show that τ is a factor in $U^C_{\delta_i}$ for some τ such that $\theta(\tau) = d$). Without loss of generality, we can assume that C is fully closed since we can add all constraints of the form $t_1 \preceq t_1$ and $t_1 \to t_2 \preceq t_1 \to t_2$ (where t_1 and t_2 are either $\bot$, $\top$, or a variable δ_i for some $i \leq n$) which makes C fully closed but leaves it finite, flat, and consistent. Notice that adding these trivial constraints does not change $L^C_{\delta_i}$ or $U^C_{\delta_i}$ in any way nor does it affect whether θ is a solution of the constraints. There are four cases to consider since each upper and lower bound must either be of the form t or $t_1 \to t_2$.

1. If $t_1 \preceq t_2 \in C$, then $\theta(t_1) \preceq \theta(t_2)$.

   There are three subcases since t_1 can be $\bot$, $\top$, or a variable δ_i.

   (A) If $\bot \preceq t_2 \in C$, then $\theta(\bot) \preceq \theta(t_2)$.
   This holds since $\bot$ is the least element in the semantic domain.

   (B) If $\top \preceq t_2 \in C$, then $\theta(\top) \preceq \theta(t_2)$.
   Since C is consistent, either $t_2 = \top$ (in which case $\theta(\top) \preceq \theta(t_2)$ holds trivially) or $t_2 = \delta_i$ for some i (in which case $\theta(\top) \preceq \theta(t_2)$ holds since $\top$ is a term in $L^C_{\delta_i}$).

   (C) If $\delta_i \preceq t_2 \in C$, then $\theta(\delta_i) \preceq \theta(t_2)$.
   There are three subcases since t_2 can be $\bot$, $\top$, or a variable δ_j.

      i. If $\delta_i \preceq \bot \in C$, then $\theta(\delta_i) \preceq \theta(\bot)$.
      Note that $\theta(U^C_{\delta_i}) = \bot$ since $\bot$ is a factor in $U^C_{\delta_i}$. Also note that $\theta(L^C_{\delta_i}) = \bot$ since every term in $L^C_{\delta_i}$ is either $\bot$ or $\delta_{i'}$ for some $i' < i$ (so the

induction hypothesis applies since $\delta_{i'} \preceq \bot \in C$ and $\theta(\delta_{i'}) \preceq \bot$ follows). Hence, $\theta(\delta_i) \preceq \bot$ as desired.

ii. If $\delta_i \preceq \top \in C$, then $\theta(\delta_i) \preceq \theta(\top)$.
This result holds since $\top$ is the greatest element in the semantic domain.

iii. If $\delta_i \preceq \delta_j \in C$, then $\theta(\delta_i) \preceq \theta(\delta_j)$.
This result is proven by induction on i. If $i < j$, then $\theta(\delta_i) \preceq \theta(\delta_j)$ holds since δ_i is a term in $L^C_{\delta_j}$. If $i = j$, then $\theta(\delta_i) \preceq \theta(\delta_j)$ holds trivially. So we may assume that $j < i$. Let τ_1 be a term in $L^C_{\delta_i}$. Hence $\tau_1 \preceq \delta_i \in C$ and $\delta_i \preceq \delta_j \in C$, from which it follows that $\tau_1 \preceq \delta_j \in C$. Either τ_1 is a term in $L^C_{\delta_j}$ (and so $\theta(\tau_1) \preceq \theta(\delta_j)$) or $\tau_1 = \delta_{i'}$ where $j \leq i' < i$ (and so $\theta(\tau_1) \preceq \theta(\delta_j)$ holds by induction). Thus, $\theta(L^C_{\delta_i}) \preceq \theta(\delta_j)$. Since δ_j is a factor in $U^C_{\delta_i}$, it follows that $\theta(U^C_{\delta_i}) \preceq \theta(\delta_j)$. Hence, $\theta(\delta_i) \preceq \theta(\delta_j)$ as desired.

2. If $t_1 \to t_2 \preceq t_3 \in C$, then $\theta(t_1 \to t_2) \preceq \theta(t_3)$.
Notice that $t_3 \neq \bot$ since otherwise C is not consistent.
If $t_3 = \top$, the result holds since $\top$ is the greatest element in the domain.
If $t_3 = \delta_i$, the result follows since $t_1 \to t_2$ is a term of $L^C_{\delta_i}$.

3. If $t_1 \to t_2 \preceq t_3 \to t_4 \in C$, then $\theta(t_1 \to t_2) \preceq \theta(t_3 \to t_4)$.
Since C is fully closed, $t_3 \preceq t_1 \in C$ and $t_2 \preceq t_4 \in C$. It follows from case 1 that $\theta(t_3) \preceq \theta(t_1)$ and $\theta(t_2) \preceq \theta(t_4)$ hold. By Property 5 of domains (Definition 1), it follows that $\theta(t_1 \to t_2) \preceq \theta(t_3 \to t_4)$ as desired.

4. If $t_1 \preceq t_2 \to t_3 \in C$, then $\theta(t_1) \preceq \theta(t_2 \to t_3)$.
Notice that $t_1 \neq \top$ since otherwise C is not consistent.
If $t_1 = \bot$, the result holds since $\bot$ is the least element in the domain.
If $t_1 = \delta_i$, the result $\theta(\delta_i) \preceq \theta(t_2 \to t_3)$ is proven by induction on i.
Let τ be a term in $L^C_{\delta_i}$. Note that $\tau \neq \top$ or else C is not consistent. If $\tau = \bot$, then $\theta(\tau) \preceq \theta(t_2 \to t_3)$ holds since $\bot$ is the least element in the semantic domain. If $\tau = \delta_{i'}$ for some $i' < i$, then $\theta(\tau) \preceq \theta(t_2 \to t_3)$ holds by induction. If $\tau = t_4 \to t_5$, then $\theta(\tau) \preceq \theta(t_2 \to t_3)$ holds by case 3. Hence, $\theta(\tau) \preceq \theta(t_2 \to t_3)$ holds for every term τ in $L^C_{\delta_i}$. Thus, $\theta(L^C_{\delta_i}) \preceq \theta(t_2 \to t_3)$. Since $t_2 \to t_3$ is a factor in $U^C_{\delta_i}$, it follows $\theta(U^C_{\delta_i}) \preceq \theta(t_2 \to t_3)$. Thus, $\theta(\delta_i) \preceq \theta(t_2 \to t_3)$ as desired.

Thus, θ is a solution to C as desired. ■

Definition 27. A domain is *adequate* if

1. the domain has contractive solutions, and

2. the domain has standard upper and lower bounds,

By Proposition 2 and Lemma 9, the Regular Tree Model (Lemma 9) is adequate. Theorem 13 shows the relationship between solutions of C and solutions of E_C.

THEOREM 13 Assume the domain is adequate and let C be a closed, consistent system of constraints. Then the following all hold:

1. If θ is a solution of E_C, then θ is a solution of C.

2. If θ is a solution of C, then $\theta[\ldots, \alpha_i \leftarrow \theta(\delta_i), \ldots]$ is a solution of E_C.

3. C has a solution

Proof:

1. This result is proven by induction on the complexity of C where the complexity of a system is the pair (md, nvt) under the lexicographical ordering where md represents the maximum nesting depth of "$\rightarrow$" in the visible terms and nvt represents the number of visible terms with that maximum nesting depth. The base case consists of the flat systems, i.e. those in which $md = 0$. The result for the base case follows from Lemma 17.

   We now proceed with the induction. Let θ be an arbitrary solution of E_C. Assume there are n variables $\delta_1, \ldots, \delta_n$ mentioned in the constraints of C. Let τ_0 be any one of the visible expressions with maximum depth. Since the system is not flat (because we are in the inductive case), $\tau_0 = \tau_0' \rightarrow \tau_0''$. Let δ_{n+1} be a fresh variable (distinct from $\delta_1, \ldots, \delta_n, \alpha_1, \ldots, \alpha_n$). Let R be a syntactic function on expressions that replaces each occurrence of τ_0 by δ_{n+1}. Thus, for example, $R((\tau_0 \rightarrow \tau_0) \rightarrow \delta_1) = (\delta_{n+1} \rightarrow \delta_{n+1}) \rightarrow \delta_1$. Let C' be the system of constraints defined by $C' = \{R(\tau_1) \preceq R(\tau_2) | \tau_1 \preceq \tau_2 \in C\} \cup \{\tau_0 \preceq \delta_{n+1}, \delta_{n+1} \preceq \tau_0\} \cup \{R(\tau_1) \preceq \tau_0 | \tau_1 \preceq \tau_0 \in C\} \cup \{\tau_0 \preceq R(\tau_1) | \tau_0 \preceq \tau_1 \in C\}$. Since any solution of C' satisfies $\delta_{n+1} = \tau_0$ and hence satisfies $R(\tau) = \tau$ for all expressions τ, it is clear that any solution of C' is also a solution of C. It is easy to check that C' is a closed, consistent system of constraints. If V is the set of visible expressions of C, then $\{R(\tau) | \tau \in V\} \cup \{\tau_0', \tau_0''\}$ is the set of visible expressions of C'; since τ_0 was eliminated as a visible expression, it is clear that C' is less complex than C and so the induction hypothesis applies to C'. Let E_0 (respectively, E_1) be the contractive system of constraints equivalent to E_C (respectively, $E_{C'}$) that is guaranteed to exist by Lemma 16. Since the domain has contractive solutions, θ^{E_1} solves E_1 and hence $E_{C'}$. By the induction hypothesis, θ^{E_1} solves C'. Since θ^{E_1} is a solution of C', $\theta^{E_1}(\delta_{n+1}) = \theta^{E_1}(\tau_0)$. Thus, $\theta^{E_1}(R(\tau)) = \theta^{E_1}(\tau)$ holds for all expressions τ. Let $1 \leq i \leq n$. If $\tau_0 \preceq \delta_i \in C$, then $L^C_{\delta_i} = L_i + \tau_0$ where L_i is $L^C_{\delta_i}$ with the τ_0 term removed and $L^{C'}_{\delta_i} = R(L_i) + \tau_0$; hence it follows that $\theta^{E_1}(L^{C'}_{\delta_i}) = \theta^{E_1}(R(L_i) + \tau_0) = \theta^{E_1}(L_i + \tau_0) = \theta^{E_1}(L^C_{\delta_i})$. If $\tau_0 \preceq \delta_i \notin C$, then $L^{C'}_{\delta_i} = R(L^C_{\delta_i})$ and hence, $\theta^{E_1}(L^{C'}_{\delta_i}) = \theta^{E_1}(R(L^C_{\delta_i})) = \theta^{E_1}(L^C_{\delta_i})$. In either case, $\theta^{E_1}(L^{C'}_{\delta_i}) = \theta^{E_1}(L^C_{\delta_i})$. Similarly, $\theta^{E_1}(U^{C'}_{\delta_i}) = \theta^{E_1}(U^C_{\delta_i})$. Thus, $\theta^{E_1}(\delta_i) = \theta^{E_1}(L^{C'}_{\delta_i} + \alpha_i \cdot U^{C'}_{\delta_i}) = \theta^{E_1}(L^C_{\delta_i} + \alpha_i \cdot U^C_{\delta_i})$. Thus, θ^{E_1} is a solution to E_C and hence to E_0. Since the system E_0 is contractive, this implies that $(\theta^{E_1})^{E_0} = \theta^{E_1}$.

$$
\begin{array}{lll}
 & \theta^{E_1} & \\
= & (\theta^{E_1})^{E_0} & \text{by the above} \\
= & (\theta[\ldots,\delta_{n+1} \leftarrow \theta^{E_1}(\delta_{n+1})])^{E_0} & \text{by definition of } \theta^{E_1} \\
= & (\theta[\delta_{n+1} \leftarrow \theta^{E_1}(\delta_{n+1})])^{E_0} & \text{since } \delta_1 \ldots \delta_n \text{ are redefined by } E_0 \\
= & \theta^{E_0}[\delta_{n+1} \leftarrow \theta^{E_1}(\delta_{n+1})] & \text{by Part 2 of Lemma 7} \\
= & \theta[\delta_{n+1} \leftarrow \theta^{E_1}(\delta_{n+1})] & \text{since } \theta \text{ is a solution of } E_C \text{ (and hence } E_0)
\end{array}
$$

Since θ^{E_1} is a solution of C' and hence of C, it follows that $\theta[\delta_{n+1} \leftarrow \theta^{E_1}(\delta_{n+1})]$ is also a solution of C. But δ_{n+1} does not occur in C. Therefore, θ is a solution of C as desired.

2. Let $\theta' = \theta[\ldots,\alpha_i \leftarrow \theta(\delta_i),\ldots]$.

$$
\begin{array}{lll}
 & \theta'(\delta_i) & \\
= & \theta(\delta_i) & \\
= & \theta(L^C_{\delta_i}) \sqcup \theta(\delta_i) & \text{since } \theta(L^C_{\delta_i}) \preceq \theta(\delta_i) \\
= & \theta(L^C_{\delta_i}) \sqcup (\theta(\delta_i) \sqcap \theta(U^C_{\delta_i})) & \text{since } \theta(\delta_i) \preceq \theta(U^C_{\delta_i}) \\
= & \theta'(L^C_{\delta_i}) \sqcup (\theta'(\alpha_i) \sqcap \theta'(U^C_{\delta_i})) & \text{since no } \alpha_j \text{ occurs in } L^C_{\delta_i} \text{ or } U^C_{\delta_i} \\
= & \theta'(L^C_{\delta_i} + \alpha_i \cdot U^C_{\delta_i}) &
\end{array}
$$

This proves that θ' is a solution of E_C as desired.

3. By Part 1 of Definition 27, contractive equations always have solutions in an adequate domain. By Lemma 16, E_C is equivalent to a contractive system of equations. Thus, E_C has a solution. The result follows from Part 1.

■

COROLLARY 4 Let C be a closed, consistent system of constraints over an adequate domain. Let E'_C be the corresponding contractive set of equations (see Lemma 16). Then $\forall\delta_1,\ldots,\delta_n.\tau/C \equiv \forall\alpha_1,\ldots,\alpha_n.\tau/E'_C$ provided $\alpha_1,\ldots,\alpha_n$ do not appear in $\forall\delta_1,\ldots,\delta_n.\tau/C$.

Proof: Fix an assignment θ.

Let $x_1,\ldots,x_n$ be elements of $\mathcal{D}_0$ such that $\theta[\ldots,\delta_i \leftarrow x_i,\ldots]$ solves C.
By Part 2 of Theorem 13, it follows that $\theta[\ldots,\delta_i \leftarrow x_i,\ldots][\ldots,\alpha_i \leftarrow x_i,\ldots]$ solves E_C.
By Lemma 16, it follows that $\theta[\ldots,\delta_i \leftarrow x_i,\ldots][\ldots,\alpha_i \leftarrow x_i,\ldots]$ solves E'_C.
Since E'_C is contractive, it follows that $(\theta[\ldots,\delta_i \leftarrow x_i,\ldots][\ldots,\alpha_i \leftarrow x_i,\ldots])^{E'_C}$
$= \theta[\ldots,\delta_i \leftarrow x_i,\ldots][\ldots,\alpha_i \leftarrow x_i,\ldots]$.

$$
\begin{array}{lll}
 & \theta(\forall\alpha_1,\ldots,\alpha_n.\tau/E'_C) & \\
\preceq & \theta[\ldots,\alpha_i \leftarrow x_i,\ldots](\tau/E'_C) & \\
= & (\theta[\ldots,\alpha_i \leftarrow x_i,\ldots])^{E'_C}(\tau) & \\
= & (\theta[\ldots,\delta_i \leftarrow x_i,\ldots][\ldots,\alpha_i \leftarrow x_i,\ldots])^{E'_C}(\tau) & \text{since } E'_C \text{ redefines } \delta_1,\ldots,\delta_n \\
= & \theta[\ldots,\delta_i \leftarrow x_i,\ldots][\ldots,\alpha_i \leftarrow x_i,\ldots](\tau) & \text{by the above} \\
= & \theta[\ldots,\delta_i \leftarrow x_i,\ldots](\tau) & \text{since no } \alpha_i \text{ appears in } \tau \\
= & \theta[\ldots,\delta_i \leftarrow x_i,\ldots](\tau/C) & \text{since } \theta[\ldots,\delta_i \leftarrow x_i,\ldots] \text{ solves } C
\end{array}
$$

Hence it follows that $\theta(\forall\alpha_1,\ldots,\alpha_n.\tau/E'_C) \preceq \theta(\forall\delta_1,\ldots,\delta_n.\tau/C)$ for arbitrary θ.

Let $x_1, \ldots, x_n$ be arbitrary elements of $\mathcal{D}_0$. Let $\theta' = \theta[\ldots, \alpha_i \leftarrow x_i, \ldots]$. Notice that $(\theta')^{E'_C}$ solves E'_C and hence solves E_C by Lemma 16. By Part 1 of Theorem 13, it follows that $(\theta')^{E'_C}$ solves C. Since $(\theta')^{E'_C}$ solves C, it follows from Definition 24 that $(\theta')^{E'_C}(\forall\delta_1, \ldots, \delta_n.\tau/C) \preceq (\theta')^{E'_C}(\tau)$.

$$
\begin{array}{lll}
 & \theta(\forall\delta_1, \ldots, \delta_n.\tau/C) & \\
= & \theta'(\forall\delta_1, \ldots, \delta_n.\tau/C) & \text{since no } \alpha_i \text{ appears in } \tau/C \\
= & (\theta')^{E'_C}(\forall\delta_1, \ldots, \delta_n.\tau/C) & \text{since no } \delta_i \text{ is free in } \forall\delta_1, \ldots, \delta_n.\tau/C \\
\preceq & (\theta')^{E'_C}(\tau) & \text{by the above} \\
= & \theta'(\tau/E'_C) & \\
= & \theta[\ldots, \alpha_i \leftarrow x_i, \ldots](\tau/E'_C) &
\end{array}
$$

Hence it follows that $\theta(\forall\delta_1, \ldots, \delta_n.\tau/C) \preceq \theta(\forall\alpha_1, \ldots, \alpha_n.\tau/E'_C)$ for arbitrary θ.

The result follows since $\theta(\forall\delta_1, \ldots, \delta_n.\tau/C) = \theta(\forall\alpha_1, \ldots, \alpha_n.\tau/E'_C)$ holds for all assignments θ. ■

A solution procedure for constraints over regular trees is given in [15].

7.2. Soundness

The equivalence between constraints and equations in an adequate domain suggests that variable elimination can be performed by first translating from constraints to equations, applying the results of Section 6.4 to eliminate variables, and then translating back (if desired) to constraints. We can improve on this procedure with a modified Extended Variable Elimination Procedure that takes advantage of the structure of constraint systems.

Procedure 14 Let $\sigma = \forall\delta_1, \ldots, \delta_n.\tau/C$ be a quantified constrained type. Let

$$\sigma' = \forall\alpha_1, \ldots, \alpha_n.\tau/E_C$$

be the corresponding extended quantified type. Perform the following steps on σ':

1. Let σ_1 be the result of replacing any subexpression $\tau_1 + \tau_2$ in σ' by τ_2 if $\theta(\tau_1) \preceq \theta(\tau_2)$ for all assignments θ. In particular, if any equation of E_C has the form

   $$\delta_i = L^C_{\delta_i} + \top \cdot U^C_{\delta_i}$$

   then replace the equation by

   $$\delta_i = U^C_{\delta_i}$$

   since $\theta(L^C_{\delta_i}) \preceq \theta(U^C_{\delta_i})$ for any solution θ of the constraints.

2. Let σ_2 be the result of replacing any subexpression $\tau_1 \cdot \tau_2$ in σ' by τ_2 if $\theta(\tau_2) \preceq \theta(\tau_1)$ for all assignments θ. In particular, if any equation of E_C has the form

   $$\delta_i = L^C_{\delta_i} + \bot \cdot U^C_{\delta_i}$$

then replace the equation by

$$\delta_i = L^C_{\delta_i}$$

3. Let $\sigma_3 = VEP(\sigma_2)$.

4. Halt if no variables are eliminated in (3); the result is σ_3. Repeat (1)-(3) on σ_3 otherwise.

We remark that the "in particular" parts of steps 1 and 2 will only become relevant if some α_i is set to either $\bot$ or $\top$ by the elimination of that variable α_i during step 3.

EXAMPLE: Consider the type

$$\forall \delta_1, \delta_2.\ \delta_1 \to \delta_2 /\ \delta_1 \preceq \delta_2$$

The extended quantified type is

$$\forall \alpha_1, \alpha_2 . \delta_1 \to \delta_2 / E \quad \text{where } E \text{ is } (\delta_1 = \bot + \alpha_1 \cdot \delta_2 \ \ \wedge \ \ \delta_2 = \bot + \alpha_2 \cdot \top)\ .$$

Now $Pos(\delta_1 \to \delta_2/E) = \{\alpha_2\}$ and $Neg(\delta_1 \to \delta_2/E) = \{\alpha_1, \alpha_2\}$. Thus α_1 can be set to $\top$; performing this substitution and simplifying gives:

$$\forall \alpha_2 . \delta_1 \to \delta_2 /\ \ \delta_1 = \delta_2 \ \ \wedge \ \ \delta_2 = \alpha_2$$

Substituting α_2 for the other variables gives:

$$\forall \alpha_2 . \alpha_2 \to \alpha_2$$

□

Soundness is easy to prove for the procedure given above.

LEMMA 18 Let $\sigma = \forall \delta_1, \ldots, \delta_n . \tau / C$ and assume that C is closed and consistent. If the domain is adequate and has type continuity, then Procedure 14 is sound for σ.

Proof: Follows from Corollary 4 and Lemma 11. ■

We note that it is easy to give an algorithm that implements the effect of Procedure 14 on the constraints directly, without requiring translations to and from equations. We have chosen to present the constrained types in terms of equations to build on the previous sections.

8. Conclusions

Polymorphic types with subtyping have rich structure. In this paper, we have shown that for simple non-recursive types and recursive types, it is possible to compute an optimal representation of a polymorphic type in the sense that no other equivalent type has fewer quantified variables. Thus, the optimal representation can be interpreted as having the minimum polymorphism needed to express the type.

In more complex type languages, in particular in languages with union and intersection types, the same methods are sound but incomplete. The completeness results for the simpler type languages show that the source of incompleteness is in fact union and intersection types in these languages. The problem of whether there is a sound and complete variable elimination procedure for languages with intersection and union types remains open.

We have also given a sound variable elimination procedure for polymorphic constrained types. Variable elimination is critically important in implementations of type systems using constrained types [9], and in fact the desire to better understand variable elimination in this setting was the original motivation for this work. However, the problem of whether there is a sound and complete procedure for eliminating variables in polymorphic constrained types also remains open.

Notes

1. The source code for the Illyria system and an interactive demo are available at URL `http://www.cs.berkeley.edu/~aiken/Illyria-demo.html`.

References

1. A. Aiken and B. Murphy. Implementing regular tree expressions. In *Proceedings of the 1991 Conference on Functional Programming Languages and Computer Architecture*, pages 427–447, August 1991.
2. A. Aiken and E. Wimmers. Type inclusion constraints and type inference. In *Proceedings of the 1993 Conference on Functional Programming Languages and Computer Architecture*, pages 31–41, Copenhagen, Denmark, June 1993.
3. A. Aiken, E. Wimmers, and T.K. Lakshman. Soft typing with conditional types. In *Twenty-First Annual ACM Symposium on Principles of Programming Languages*, pages 163–173, Portland, Oregon, January 1994.
4. Roberto M. Amadio and Luca Cardelli. Subtyping recursive types. *ACM Transactions on Programming Languages and Systems*, 15(4):575–631, 1993. Also in Proc. POPL'91.
5. L. Cardelli and P. Wegner. On understanding types, data abstraction and polymorphism. *Computing Surverys*, 17(4):471–522, December 1985.
6. Bruno Courcelle. Infinite trees in normal form and recursive equations having a unique solution. *Mathematical Systems Theory*, 13:131–180, 1979.
7. Pavel Curtis. Constrained quantification in polymorphic type analysis. Technical Report CSL-90-1, Xerox Parc, February 1990.
8. J. Eifrig, S. Smith, and V. Trifonov. Sound polymorphic type inference for objects. In *OOPSLA '96*, 1995.
9. M. Fähndrich and A. Aiken. Making set-constraint program analyses scale. In *CP96 Workshop on Set Constraints*, August 1996.

10. Fritz Henglein and Christian Mossin. Polymorphic binding-time analysis. In Donald Sannella, editor, *Proceedings of European Symposium on Programming*, volume 788 of *Lecture Notes in Computer Science*, pages 287–301. Springer-Verlag, April 1994.
11. Stefan Kaes. Type inference in the presence of overloading, subtyping and recursive types. In *1992 ACM Conference on Lisp and Functional Programming. San Francisco, California. LISP Pointers V, 1*, pages 193–204, June 1992.
12. A. Koenig. An anecdote about ML type inference. In *Proceedings of the USENIX 1994 Symposium on Very High Level Languages*, October 1994.
13. D. MacQueen, G. Plotkin, and R. Sethi. An ideal model for recursive polymophic types. In *Eleventh Annual ACM Symposium on Principles of Programming Languages*, pages 165–174, January 1984.
14. J. C. Mitchell and R. Harper. The essence of ML. In *Fifteenth Annual ACM Symposium on Principles of Programming Languages*, pages 28–46, January 1988.
15. J. Palsberg and P. O'Keefe. A type system equivalent to flow analysis. *ACM Transactions on Programming Languages and Systems*, 17(4):576–599, July 1995. Preliminary version in Proc. POPL'95, 22nd Annual SIGPLAN–SIGACT Symposium on Principles of Programming Languages, pages 367–378, San Francisco, California, January 1995.
16. F. Pottier. Simplifying subtyping constraints. In *Proceedings of the 1996 ACM SIGPLAN International Conference on Functional Programming*, pages 122–133, May 1996.
17. Geoffrey S. Smith. Principal type schemes for functional programs with overloading and subtyping. *Science of Computer Programming*, 23:197–226, 1994.