A General Theory of Semi-Unification

Said Jahama *
Boston University
(jahama@cs.bu.edu)

A.J. Kfoury †
Boston University
(kfoury@cs.bu.edu)

December 20, 1993
Technical Report: bu-cs # 93-018

Abstract

Various restrictions on the terms allowed for substitution give rise to different cases of semi-unification. Semi-unification on finite and regular terms has already been considered in the literature. We introduce a general case of semi-unification where substitutions are allowed on non-regular terms, and we prove the equivalence of this general case to a well-known undecidable data base dependency problem , thus establishing the undecidability of general semi-unification.

We present a unified way of looking at the various problems of semi-unification. We give some properties that are common to all the cases of semi-unification. We also the principality property and the solution set for those problems. We prove that semi-unification on general terms has the principality property. Finally, we present a recursive inseparability result between semi-unification on regular terms and semi-unification on general terms.

*Partly supported by NSF grant CCR-9113196. Address: Department of Computer Science, Boston University, 111 Cummington St., Boston, MA 02215, USA.

†Partly supported by NSF grant CCR-9113196. Address: Department of Computer Science, Boston University, 111 Cummington St., Boston, MA 02215, USA.

1 Introduction

In this report we study the problem of semi-unification. Semi-unification has been studied in connection with various problems in computer science. Some of the areas in which the theory of semi-unification is applied are: proof theory, term rewriting systems, polymorphic type inference and natural language processing. Semi-unification generalizes the notion of unification and matching[1].

There are many definitions of semi-unification depending on the context in which it is used. To give the reader a flavor of the problem, we give here a standard definition of semi-unification; this definition will be modified in various ways in the following sections. Let Σ be a first-order signature consisting of exactly one function symbol and a countably infinite set of constant symbols. Let X be a countably infinite set of variables, and $\mathcal{T}$ the set of all finite terms over Σ and X. An *instance* Γ of semi-unification is a finite set of pairs:

$$\Gamma = \{(t_1, u_1), \ldots, (t_n, u_n)\}$$

where $t_i, u_i \in \mathcal{T}$. A *substitution* is a function $S : X \to \mathcal{T}$. Every substitution extends in a natural way to a Σ-homomorphism $S : \mathcal{T} \to \mathcal{T}$. A substitution S is a *solution* of the instance Γ iff there are substitutions $S_1, \ldots, S_n$ such that:

$$S_1(S(t_1)) = S(u_1)\ ,\ \ldots\ ,\ S_n(S(t_n)) = S(u_n) \tag{1}$$

The *semi-unification problem* is the problem of deciding whether an arbitrary instance Γ has a solution.

The preceding definition of the semi-unification problem allows variables in Γ to be substituted only by finite terms. We call this case of semi-unification *finite* semi-unification. Finite semi-unification and restrictions on it have been considered in connection with many research problems. We mention here some of these research areas. A sufficient condition for a rewrite rule to be non-terminating can be formulated as an instance of finite semi-unification with exactly one inequality (see [10, 19]). The type reconstruction problem for **ML** is known to be equivalent to a special case of finite semi-unification where the instance satisfies some acyclicity conditions [12].

[1] Given a pair of terms (t, u), unification is the problem of finding a substitution S such that $S(t) = S(u)$, while matching is the problem of finding a substitution S such that $S(t) = u$.

Furthermore, it has been shown in [8, 14] that the general form of finite semi-unification is equivalent to the type reconstruction problem for **ML** with polymorphic recursion. Finite semi-unification has also been studied in connection with Kriesel's conjecture on the length of proofs in Peano Arithmetic (see [18]).

If we generalize substitutions by allowing them to replace variables by terms that are not necessarily finite, we obtain other cases of semi-unification. For example, the case of *regular* semi-unification is obtained by extending the notion of a substitution by allowing it to replace a variable by a (possibly infinite) regular term [2]. We call such a substitution a regular substitution. If the substitutions $S, S_1, \ldots, S_n$ in equation 1 are allowed to be regular, we say that S is a regular solution for Γ (but note we still require that Γ be an instance defined on finite terms). Regular semi-unification has been considered in the literature in connection with the "clause satisfiability problem" in a "feature algebra" (this is a problem considered in computational linguistics, see [4]).

Both finite and regular semi-unification are undecidable. The proof in the regular case is by a direct reduction from the word problem of finite semi-groups, and is given in [4]. By contrast, finite semi-unification was long believed to be decidable. The undecidability result is given in [13]. The proof of this result is fairly technical and complicated, based on a specialized undecidability result about Turing machines established in the early 1960's [9].

Several special cases of finite semi-unification have been shown to be decidable. We list here some of those decidable special cases of finite semi-unification. Finite semi-unification on instances with exactly one inequality has been shown to be decidable [6, 10, 18]. Furthermore, in [10] a polynomial time decision procedure is given. Acyclic semi-unification is shown to be DEXPTIME-complete in the case of finite semi-unification [12]. In [11], an exponential procedure is given to show that left linear semi-unification is decidable, this result was improved in [7] where a polynomial procedure is given. Instances over 2 variables is also decidable in the finite case [15]. In addition, see [16] where another special case of finite semi-unification is shown to be decidable.

[2] Regular terms are terms with finitely many unequal subterms. In general, the same subterm occurs infinitely many times in a regular term.

In this report, we focus on the theory of semi-unification independent of the context in which it may be used. We first consider an original case of semi-unification. This case is obtained by allowing variables in Γ to be substituted by general unrestricted terms (the exact definition of general terms is given in Section 3). We call this case *general* semi-unification. As we show later in the report, the undecidability of regular semi-unification implies the existence of instances that do not have a regular solution , but have a solution in the general case. Notice that in the case of first-order unification, instances that have a solution in the general sense always have a regular solution. This presents a major difference between unification and semi-unification. It also provides some insight into the reason why semi-unification is undecidable.

In this report, we prove the equivalence of general semi-unification to a well known problem in database theory, the so-called functional and inclusion dependency problem. This dependency problem is undecidable (see [1, 2, 17]), and hence, we obtain the undecidability of general semi-unification.

We also present a metatheory that is common to the different cases of semi-unification. The core of this metatheory is a *redex* procedure used to transform instances of semi-unification. This redex procedure has the property that the each successive instance has a solution (finite, regular or general) iff the initial instance has a solution. Furthermore, in the case of finite semi-unification this redex procedure halts iff the initial instance has a finite solution[3]. In the general case the redex procedure will halt iff the initial instance has no general solution. The redex procedure might go on forever generating an infinite general solution for the instance. Hence, even though we allow substitutions of arbitrary terms in the general case, any instance that has a general solution has a solution that is generated by the redex procedure. We also give an underlying *equational theory* that can be used to describe various relationships in the desired solution for the instance.

We use this development of the metatheory of semi-unification to study properties that are common to all the three cases of semi-unification: finite, regular, and general. We show that each of the three cases of semi-unification on a signature with a single binary function symbol is equivalent to semi-unification of the same case on an arbitrary k-ary signature for $k \geq 2$. We also show that each of the three cases of semi-unification with only 2 inequalities

[3]In the finite case the redex procedure is the same as that given in [13].

is equivalent to semi-unification of the same case with $n \geq 2$ inequalities[4]. We also look at some restrictions that make (finite, regular and general) semi-unification decidable.

The report also studies some properties that distinguish the three cases of semi-unification. We study the principality property for each case. For the three cases of semi-unification, the principality property holds if instances that have a solution have a minimal solution such that every other solution for the instance can be obtained from the minimal solution by applying a substitution to it. In the finite case, it has been shown in [6, 13] that the principality property holds. We show that it holds also in the general case. However, in the regular case, it is still open whether this property holds. In the regular case, we show that even if the principality property holds, we cannot decide whether a given solution is principal. We also study the structure of the solution set for the various cases of semi-unification.

Two sets A and B are *recursively inseparable* if there is no recursive set containing A and disjoint from B [5]. Notice that recursive inseparability of A and B suffices to show the the nonrecursiveness of both A and B. In this report, we present a recursive inseparability result between the set of instances which have no general solution and the set of instances which have a regular solution. Consider a subset $\mathcal{T}'$ of general terms that contains the set of regular terms (for example $\mathcal{T}'$ could be the set of r.e. terms). Consider the case of semi-unification obtained by extending the notion of a regular substitution to allow a substitution to replace variables by members of $\mathcal{T}'$. We use the recursive inseparability result mentioned in this paragraph to show that the set of instances which have a solution (in the sense of any such extension) is recursively inseparable from the set of instances which have no general solution. See Section 7.2 for the details.

The report is organized as follows. Section 2 contains the definition of the finite semi-unification problem. It also gives a redex procedure originally described in [13]. It also gives the definition of the equational theory related to semi-unification. Most of the material in Section 2 is borrowed from [13]. Section 3 gives the definition of general semi-unification, and gives a modified redex procedure. It also relates this procedure to the equational theory. Section 4 gives the definition of the functional and inclusion dependency problem. It establishes the equivalence of an undecidable subcase of

[4]In the finite case the proof (unpublished) is due to Pudlack [18].

dependency to general semi-unification, which provides one way of proving the undecidability of general semi-unification. Section 5 gives the definition of regular semi-unification and relates it to clause satisfiability in a feature algebra. Section 5 also contains an outline of a reduction given in [4] from the word problem of semi-groups to the clause satisfiability problem. Section 6 contains a reduction from instances on a k-ary signature to instances on a binary signature. It also contains a reduction from instances with n inequalities to instances with 2 inequalities. Section 6 also contains some cases that are decidable in the three cases of semi-unification. Section 7 contains a study of the principality property and the solution set. It also contains the result that instances which do not have a general solution are recursively inseparable from instances which have a regular solution. Finally, Section 8 contains the conclusion of this report and a list of open problems.

2 Finite Semi-Unification

For the purpose of this section we will define a more restricted form of semi-unification. The first-order signature Σ now consists of exactly one binary function symbol. We denote this function symbol by $\rightarrow$ used in infix notation. For technical reasons, we now look at a pair (t, u) as an inequality $t \leq u$.

Let X be a countably infinite set of variables, and $\mathcal{T}$ the set of all *finite* terms over Σ and X. An *instance* Γ of semi-unification is a finite set of inequalities:

$$\Gamma = \{t_1 \leq u_1, \ldots, t_n \leq u_n\}$$

where $t_i, u_i \in \mathcal{T}$. A *substitution* is a function $S : X \rightarrow \mathcal{T}$.

A substitution S is a *solution* of the instance Γ iff there are substitutions $S_1, \ldots, S_n$ such that:

$$S_1(S(t_1)) = S(u_1)\ ,\ \ldots\ ,\ S_n(S(t_n)) = S(u_n) \tag{2}$$

The *finite semi-unification Problem* is the problem of deciding, for any such instance Γ, whether Γ has a solution. We sometimes refer to a solution by the whole vector $(S, S_1, \ldots, S_n)$ instead of just S. If S, S' are substitutions we sometimes write SS' for their composition $S(S'())$.

For a term t, define $L(t)$ and $R(t)$ (the left and right subterms of t). If t is a variable then $L(t)$ and $R(t)$ are undefined, otherwise:

$$L(t \to t') = t$$

$$R(t \to t') = t'$$

If $\Pi \in \{L, R\}^*$, say $\Pi = x_1 x_2 \ldots x_p$, the notation $\Pi(t)$ means $x_1(x_2(\cdots(x_p(t)\cdots))$.

We can assume that the inequalities in Γ are given in a fixed order, so Γ could be written as :

$$\{t_1 \leq_1 u_1, \ldots, t_n \leq_n u_n\}$$

Let $\overline{X}$ be an extension of X: elements of $\overline{X}$ are symbols of the form α_w where $\alpha \in X$ and $w \in \{1, 2, \ldots, n\}^*$. α_ε where ε is the empty string, is identified with α. Let $\overline{\mathcal{T}}$ be the set of finite terms over Σ and $\overline{X}$.

For $w \in \{1, 2, \ldots, n\}^*$, let $(\ )_w : \overline{\mathcal{T}} \to \overline{\mathcal{T}}$, be a homomorphism defined on variables in $\overline{X}$ by:

$$(\alpha_v)_w = \alpha_{vw}$$

2.1 Redex Procedure

We now present a procedure originally given in [13]. This procedure takes as input an instance Γ as defined above. The procedure halts iff Γ has a (finite) solution. If and when it halts it constructs a solution for Γ. The procedure consists of repeatedly reducing redexes which can be of two kinds. It halts when there are no more redexes left. In the following sections we will modify this procedure to make it suitable for the other cases of semi-unification.

- (*Redex I reduction*) Let $\xi \in \overline{X}$ and let $t' \notin \overline{X}$ be a term with the property that there is a path $\Pi \in \{L, R\}^*$ and $i \in \{1, \ldots, n\}$ such that if $t \leq_i u$ is the i-th inequality of Γ, then:

  $$\Pi(t) = t' \quad \text{and} \quad \Pi(u) = \xi$$

  The pair of terms $(\xi, (t')_i)$ is called a *redex I*. The result of reducing this redex consists in substituting $(t')_i$ for all occurrences of ξ throughout Γ.

- (*Redex II reduction*) Let $\xi \in \overline{X}$ and $u' \in \overline{\mathcal{T}}$ have the property that $\xi \neq u'$ and there are paths $\Pi, \Delta, \Sigma \in \{L, R\}^*$ and $i \in \{1, \ldots, n\}$ such that if $t \leq_i u$ is the i-th inequality in Γ, then:

  $$\Pi(t) = \Delta(t) \in \overline{X} \quad \text{and} \quad \Sigma\Pi(u) = \xi \quad \text{and} \quad \Sigma\Delta(u) = u'$$

  Such a pair (ξ, u') is called a *redex II*. The result of reducing this redex consists in substituting u' for all occurrences of ξ throughout Γ.

Notice that, in the redex procedure above, a semi-unification instance Γ is allowed to contain variables in $\overline{X}$. So, now a substitution is viewed as a function from $\overline{X}$ to $\overline{\mathcal{T}}$. The following lemma is taken from [13] and it illustrates the correctness of the redex procedure.

Lemma 1 *Let Γ be an instance of semi-unification with all variables in the initial X.*

1. *Γ has a solution iff the above procedure, when started on instance Γ, halts producing an instance Γ' without redexes.*

2. *If Γ has no solution, then the above procedure will keep producing instances Γ' which assign arbitrarily large terms to variables in X.*

If an instance of finite semi-unification has a solution, it will have a unique minimal solution in the sense that every other solution can be obtained by applying a substitution to this minimal solution. This minimal solution is called a *principal* solution. To make this notion precise, we first consider an ordering $\sqsubseteq_\Gamma$ on the set of all solutions for Γ:

Definition 2 *$S' \sqsubseteq_\Gamma S''$ iff there is a substitution $P : \overline{X} \to \overline{\mathcal{T}}$ such that $P(S'(\xi)) = S''(\xi)$ for every variable ξ occurring in Γ.*

A solution S of Γ is *principal* iff for every other solution S', $S \sqsubseteq_\Gamma S'$. The redex procedure given above, when it halts, produces the principal solution for Γ (see [13]).

2.2 Equational Theory

Each case of semi-unification we shall consider gives rise to an equational theory EQ, generalizing the concept of *Path Equations* given in [13].

Let Σ_{EQ} be a finite set of *unary* function symbols. Let X be a countably infinite set of variables. Let $\overline{X}$ be an extension of X, as defined in the previous section. Let $\overline{\mathcal{T}}_{EQ}$ be the set of finite terms over Σ_{EQ} and $\overline{X}$.

Note that any term $t \in \overline{\mathcal{T}}_{EQ}$ is of the form $\Pi\alpha_w$ where $\Pi \in \Sigma^*_{EQ}$ and $\alpha_w \in \overline{X}$.

An equation e is a string of the form $t = u$ where $t, u \in \overline{\mathcal{T}}_{EQ}$.

Definition 3 *An* **equational theory** EQ *consists of a finite set of axioms* E *(a set of equations over* $\overline{\mathcal{T}}_{EQ}$*), and* **Rules of Inference**.

Rules of Inference: Assume $t, u, v \in \overline{\mathcal{T}}_{EQ}$, $w \in \{1, \ldots, n\}^*$, $f \in \Sigma_{EQ}$.

(transitivity) $$\frac{t = u, \quad u = v}{t = v}$$

(symmetry) $$\frac{t = u}{u = t}$$

(instantiation) $$\frac{t = u}{t_w = u_w}$$

(subterm) $$\frac{t = u}{ft = fu}$$

We say $E \vdash e$ (E derives e) if e is obtained from E using only the above inference rules.

2.3 Finite Semi-Unification and Equational Theory

We now represent the finite semi-unification problem as a problem in an equational theory EQ. Notice that in this case the equational theory is exactly the same as path equations defined in [13].

Given an instance Γ we construct an instance E_Γ of equations. Σ_{EQ} consists of only two unary function symbols L and R, corresponding to the L and R functions defined in the previous section. The set of variables X is the same as the set of variables mentioned in Γ. We extend X to get $\overline{X}$, elements of $\overline{X}$ are symbols of the form α_w where $\alpha \in X$ and $w \in \{1, 2, \ldots, n\}^*$.

To construct E_Γ from the set of inequalities of Γ add the following set of equations for every inequality $t \leq_i u$, for $i \in \{1, \ldots, n\}$:

1. For every $\Pi, \Lambda \in \Sigma_{EQ}^*$ and $\alpha, \beta \in X$, such that $\Pi\Lambda(t) = \alpha$ and $\Lambda(u) = \beta$, add the following equation to E_Γ:

$$\alpha_i = \Pi\beta$$

2. For every $\Pi, \Lambda \in \Sigma_{EQ}^*$ and $\alpha, \beta \in X$, such that $\Pi\Lambda(u) = \alpha$ and $\Lambda(t) = \beta$, add the following equation E_Γ:

$$\alpha = \Pi\beta_i$$

The following lemma is a rephrasing of Lemma 4 in [13]. It illustrates the relationship between the equational theory and semi-unification.

Lemma 4 *Let Γ be an instance of semi-unification.*

1. *If e is derivable from E_Γ then it is satisfied by every solution of Γ.*

2. *Let Γ' be obtained from Γ by reducing some number of redexes, and let t be the term assigned to $\alpha \in X$ in Γ'. For every $\Pi \in \{L, R\}^*$ and $\beta_w \in \overline{X}$ such that $\Pi(t) = \beta_w$, $E_\Gamma \vdash (\Pi\alpha = \beta_w)$.*

Definition 5 *Let α be a variable in X.* $\mathrm{ext}(\alpha)$ *(the extent of α) is defined as:*

$$\mathrm{ext}(\alpha) = \{\Pi \in \Sigma_{EQ}^* \mid \textit{there is a } \beta_w \in \overline{X} \textit{ such that } E_\Gamma \vdash (\Pi\alpha = \beta_w)\}$$

Theorem 6 *An instance Γ of semi-unification has a finite solution iff for every variable $\alpha \in X$,* $\mathrm{ext}(\alpha)$ *is finite.*

Proof: For the "if" direction , if Γ has no solution then by Lemma 1 the redex procedure will assign arbitrary large terms for at least one variable $\alpha \in X$, so by Lemma 4 there are arbitrary many Π such that $\Pi\alpha = \beta_w$. Hence, $ext(\alpha)$ can not be finite.

For the "only if" direction, assume by contradiction that Γ has a finite solution S but there is a variable $\alpha \in X$ where $ext(\alpha)$ is not finite. Hence, there are arbitrary large $\Pi \in \Sigma_{EQ}^*$ such that, $\Pi\alpha = \beta_w$. But by Lemma 4, every solution must satisfy those equations. Hence, any solution must assign to α arbitrary large terms and S can not be finite, which is a contradiction. ■

Finite semi-unification is undecidable. Its undecidability was established in [13].

3 General Semi-Unification

The semi-unification problem described in Section 2 is restricted to substitutions on ordinary (finite) terms. The problem can be extended to allow substitutions on general terms (finite or infinite). If we allow such terms then every instance Γ over the first order signature consisting of one binary function symbol, as described in Section 2, will have a solution. In fact, instances with no constants always have an infinite solution (Corollary 18). This is no longer the case when the signature contains constants.

For the remainder of this report we need to look at terms in a different way.

Let Σ be a first-order signature consisting of a countably infinite set of constants C and one function symbol F of arity $k \geq 1$. Let X be a countably infinite set of variables. For definiteness, let C be $c^0, c^1, \ldots, c^i, \ldots$ and X be $x^0, x^1, \ldots, x^i, \ldots$, where $i \in \omega$.

Definition 7

- *$T \subseteq \{f^1, f^2, \ldots, f^k\}^*$ is called a (k-ary) tree iff it satisfies the following conditions:*

  1. *T is not empty.*
  2. *For all $\Pi \in T$ if $\Pi = \Delta_1\Delta_2$ then $\Delta_2 \in T$.*
  3. *For every i, j where $1 \leq i, j \leq k$ if $f^i\Pi \in T$ then $f^j\Pi \in T$.*

- *For any tree T, $\Pi \in$ exterior(T) iff $f^i\Pi \notin T$ where $1 \leq i \leq k$.*

- *A term t is a pair (T, φ) where T is a tree and φ : exterior$(T) \rightarrow (X \cup C)$.*

Every $\Pi \in T$ represents both a node in T and the path from the root of T to that node; the root of T is the empty string ε and Π as a path from the root is read from right to left.

Definition 8

- *Let $\mathcal{T}^*$ be the set of all terms (T, φ) .*

- *Let $\mathcal{T}_{fin}$ be the set of all terms (T, φ) such that T is a finite subset of $\{f^1, \ldots, f^k\}^*$. (Notice that $\mathcal{T}_{fin}$ is the set of terms defined in the Section 2 when k=2 and $C = \emptyset$). $\mathcal{T}_{fin}$ is a proper subset of $\mathcal{T}^*$.*

In this section, a substitution S is a function $S : X \rightarrow \mathcal{T}^*$. Every substitution S can be extended to a function $S : \mathcal{T}^* \rightarrow \mathcal{T}^*$, defined as follows. For every $t_1 = (T_1, \varphi_1) \in \mathcal{T}^*$, let $S(t_1) = t_2 = (T_2, \varphi_2)$ where:

$$T_2 = T_1 \cup \{\Pi\Delta | \Delta \in exterior(T_1), \varphi_1(\Delta) = x, S(x) = (T, \varphi), \Pi \in T\}$$

and for every $\Sigma \in exterior(T_2)$ such that $\Sigma = \Pi\Delta$ where $\Delta \in exterior(T_1)$

$$\varphi_2(\Sigma) = \begin{cases} \varphi_1(\Delta), & \text{if } \varphi_1(\Delta) = c \in C, \text{ (in this case } \Sigma = \Delta) \text{ ;} \\ \varphi(\Pi), & \text{if } \varphi_1(\Delta) = x \text{ and } S(x) = (T, \varphi) \text{ and } \Pi \in exterior(T) \text{ .} \end{cases}$$

An *instance* Γ of general semi-unification is a finite set of inequalities:

$$\Gamma = \{t_1 \leq u_1, \ldots, t_n \leq u_n\}$$

where $t_i, u_i \in \mathcal{T}_{fin}$.

A substitution S is a *solution* of the instance Γ iff there are substitutions $S_1, \ldots, S_n$ such that :

$$S_1(S(t_1)) = S(u_1), \ \ldots \ , \ S_n(S(t_n)) = S(u_n) \tag{3}$$

Observe that while S is a function from $\mathcal{T}_{fin}$ to $\mathcal{T}^*$, each S_i is in general a function from $\mathcal{T}^*$ to $\mathcal{T}^*$. The *general semi-unification problem* is the problem of deciding, for any such instance Γ, whether Γ has a solution. It is important to note that we do not require that the equalities in 3 be effectively tested. We can not impose such a requirement because substitutions here map variables to arbitrary terms (not necessarily finite, regular, or even recursive). On the other hand the general semi-unification problem is well posed, because an instance Γ is always a finite object and therefore can be effectively presented.

We need to define functions on terms that allow us to select a subterm of a term t, similar to the L and R functions defined in Section 2. For a term t and $i \in \{1, \ldots, k\}$, $f^i(t)$ is a unary function which gives the i-th subtree of the k-ary tree representing t. More formally:

$$f^i(t) = f^i(T, \varphi) = (T_0, \varphi_0) \quad \text{where}$$

$$T_0 = \{\Pi : \Pi f^i \in T\}$$

$$\varphi_0(\Pi) = \varphi(\Pi f^i)$$

For $\Pi \in \{f^1, \ldots, f^k\}^*$, say $\Pi = f^{i_1} f^{i_2} \ldots f^{i_p}$ where $i_1, i_2, \ldots, i_p \in \{1, \ldots, k\}$, the notation $\Pi(t)$ means $f^{i_1}(f^{i_2}(\cdots(f^{i_p}(t)\cdots)))$.

Let $\overline{X}$ be an extension of X, as described previously, i.e. elements of $\overline{X}$ are symbols of the form α_w where $\alpha \in X$ and $w \in \{1, 2, \ldots, n\}^*$. For a term t , let $(t)_w = (T, \varphi)_w = (T, \varphi_w)$ where

$$\varphi_w(\Pi) = \begin{cases} x_w, & \text{if } \varphi(\Pi) = x \in X \ ; \\ c, & \text{if } \varphi(\Pi) = c \in C \ . \end{cases}$$

We also modify φ to allow variables in $\overline{X}$ to be used . More formally:

$$\varphi : exterior(T) \to (\overline{X} \cup C)$$

Let $\overline{\mathcal{T}}^*$ be the set of all terms $(t)_w$, where $t \in \mathcal{T}^*$ and $w \in \{1, 2, \ldots, n\}^*$. To make things more readable, when the the tree part of the term is $\{\varepsilon\}$ we sometimes omit it and only refer to the labeling function φ. For example, for a term t, $t = x$ is an abbreviation of $t = (\varepsilon, \varphi)$ and $\varphi(\varepsilon) = x$.

3.1 Redex Procedure

We now give a modified redex procedure that halts if Γ has no general solution. It also halts if Γ has a finite solution, and it does not halt iff Γ has an infinite solution. There are two main adjustments that need to be done to the redex procedure given earlier. The first one is that we now need to account for constants and for the possibility of concluding that Γ has no solution. The second adjustment is that we do not want the redex procedure to go forever if Γ has no solution. For this reason we need to develop a $fair$ strategy for the redex procedure. This could be done by forcing a timestamp ordering which ensures that redexes with old variables are reduced before redexes that have newer variables (generated as a result of another redex). We first give a version of the procedure without any bookkeeping (e.g. timestamps).

Modified Redex Procedure: Version 1

1. The input is an instance Γ where every term in Γ is in $\mathcal{T}_{fin}$.

2. (*Illegal Redexes*) If there is an inequality $t \leq_i u$ in Γ such that one of the two following cases occurs in the inequality, then the procedure stops and Γ has no solution.

   - There is a path $\Pi \in \{f^1, \ldots, f^k\}^*$ and $i \in \{1, \ldots, n\}$ such that $\Pi(t), \Pi(u) \in \overline{\mathcal{T}}^*$ and $\Pi(t), \Pi(u) \notin \overline{X}$, and either:
   $$\Pi(t) = c \quad \text{and} \quad \Pi(u) \neq c \text{ or}$$
   $$\Pi(u) = c \quad \text{and} \quad \Pi(u) \neq c \text{ where } c \in C.$$
   - There are paths $\Pi, \Delta, \Sigma \in \{f^1, \ldots, f^k\}^*$ and $\Sigma\Pi(u), \Sigma\Delta(u) \in \overline{\mathcal{T}}^*$ and $\Sigma\Pi(u), \Sigma\Delta(u) \notin \overline{X}$ and:
   $$\Pi(t) = \Delta(t) \in \overline{X} \quad \text{and} \quad \Sigma\Pi(u) = c \quad \text{and} \quad \Sigma\Delta(u) \neq c \quad \text{where} \quad c \in C$$

   If neither one of these cases occurs goto step 3.

3. Find a redex (ξ, t) as described in step 4 or 5 and reduce it. If there is no such redex, the procedure stops and Γ has a (finite) solution.

4. (*Redex I reduction*) Let $\xi \in \overline{X}$ and let $t' \notin \overline{X}$ be a term with the property that there is a path $\Pi \in \{f^1, \ldots, f^k\}^*$ and $i \in \{1, \ldots, n\}$ such that if $t \leq_i u$ is the i-th inequality of Γ, then:
$$\Pi(t) = t' \quad \text{and} \quad \Pi(u) = \xi$$

The pair of terms $(\xi, (t')_i)$ is called a *redex I*. The result of reducing this redex consists in substituting $(t')_i$ for all occurrences of ξ throughout Γ. Goto step 2.

5. (*Redex II reduction*) Let $\xi \in \overline{X}$ and u' be a term with the property that $\xi \neq u'$ and there are paths $\Pi, \Delta, \Sigma \in \{f^1, \ldots, f^k\}^*$ and $i \in \{1, \ldots, n\}$ such that if $t \leq_i u$ is the i-th inequality in Γ, then:

$$\Pi(t) = \Delta(t) \in \overline{X} \quad \text{and} \quad \Sigma\Pi(u) = \xi \quad \text{and} \quad \Sigma\Delta(u) = u'$$

   Such a pair (ξ, u') is called a *redex II*. The result of reducing this redex consists in substituting u' for all occurrences of ξ throughout Γ. Goto step 2.

For an instance Γ and a substitution S, $S(\Gamma)$ is the instance obtained by applying S to every term in Γ.

For the same reasons as in Section 2, we now consider a substitution to be from $\overline{X}$ to $\overline{\mathcal{T}}^*$. Each redex (ξ, t) determines a substitution $R : \overline{X} \to \overline{\mathcal{T}}^*$ such that $R(\xi) = t$ and $R(\eta) = \eta$ for $\eta \neq \xi$; let us call each such substitution a *basic substitution* of the redex procedure.

Since the redex procedure is non-deterministic, different runs of the procedure give rise to different sequences of basic substitutions (i.e. redexes). Consider a specific run of the procedure: let σ be the sequence of basic substitutions evaluated in the course of the run and R^i_σ the i-th basic substitution in σ. For simplicity, we refer to a run of the redex procedure by the sequence σ of basic substitutions it evaluates.

Let Γ be an instance of general semi-unification. A run σ of the redex procedure on Γ generates an illegal redex if there is a Γ' obtained from Γ, via σ , and Γ' has an illegal redex occurrence. Γ generates an illegal redex if there is a run σ of the redex procedure on Γ such that σ generates an illegal redex. There is no guarantee that the procedure above will detect that a given instance Γ generates an illegal redex, if it indeed does. The reason is that it could happen that the procedure might run forever choosing redexes that lead to instances that do not have an illegal redex, while some other choice of redexes leads to an instance that has an illegal redex. For this reason we have to ensure that every run of the redex procedure be fair in the sense that eventually all redexes will be accounted for.

We call redexes of the form (α, β), where α and β are both variables, *trivial* redexes. Notice that if instance Γ has a trivial redex (α, β) then it will also have a redex (β, α). A trivial redex could only be a Redex II . For technical reasons, in the second version of the redex procedure, we first reduce trivial redexes. Furthermore, among any such pair of trivial redexes we choose to reduce (i.e. rename) the variable that was introduced later in the procedure. In other words, if we have a choice between (α, β) and (β, α) we reduce the first redex if β was introduced before α, otherwise, we choose the second one.

We now give the precise definition of the notion of fairness.

Definition 9 *Let Γ be an instance of general semi-unification. A run σ of the redex procedure on Γ is* fair *if the following condition is satisfied:*

> *If σ does not generate an illegal redex then for every $k \geq 0$ and every redex (α, t) in $R_\sigma^k \cdots R_\sigma^1(\Gamma)$ there is a $j > k$ such that*
>
> 1. *If (α, t) is a trivial redex where $t = \beta$ then there is a variable $\gamma \in \overline{X}$ such that*
> $$R_\sigma^j \cdots R_\sigma^{k+1}(\alpha) = R_\sigma^j \cdots R_\sigma^{k+1}(\beta) = \gamma.$$
> 2. *If (α, t) is not a trivial redex then $R_\sigma^j \cdots R_\sigma^{k+1}(\alpha) = t'$ where t' is not a variable.*

This definition imposes a certain order of reduction on all redexes which have the same variable as their first component. For example, if both the trivial redex (β, α) where β was introduced later than α, and the non trivial redex (β, t) occur in $R_\sigma^k \cdots R_\sigma^1(\Gamma)$, then (β, α) has to be reduced before (β, t), assuming that there are no other trivial redexes occurring in $R_\sigma^k \cdots R_\sigma^1(\Gamma)$.

Modified Redex Procedure: Version 2

1. The input is an instance Γ where every term in Γ is in $\mathcal{T}_{fin}$. Set the initial value of p (the timestamp value) to 0 and every variable occurrence in Γ is assigned the current value of p. Let $\Gamma^0 = \Gamma$ and $q = 0$.

2. (*Illegal Redexes*) If there is an inequality $t \leq_i u$ in Γ^j such that one of the two following cases occurs in the inequality, then the procedure stops and Γ has no solution.

- There is a path $\Pi \in \{f^1, \ldots, f^k\}^*$ and $i \in \{1, \ldots, n\}$ such that $\Pi(t), \Pi(u) \in \overline{\mathcal{T}}^*$ and $\Pi(t), \Pi(u) \notin \overline{X}$, and either:

$$\Pi(t) = c \quad \text{and} \quad \Pi(u) \neq c \text{ or}$$
$$\Pi(u) = c \quad \text{and} \quad \Pi(u) \neq c \text{ where } c \in C.$$

- There are paths $\Pi, \Delta, \Sigma \in \{f^1, \ldots, f^k\}^*$ and $\Sigma\Pi(u), \Sigma\Delta(u) \in \overline{\mathcal{T}}^*$ and $\Sigma\Pi(u), \Sigma\Delta(u) \notin \overline{X}$ and:

$$\Pi(t) = \Delta(t) \in \overline{X} \quad \text{and} \quad \Sigma\Pi(u) = c \quad \text{and} \quad \Sigma\Delta(u) \neq c \quad \text{where} \quad c \in C$$

If neither one of these cases occurs goto step 3.

3. Find a trivial redex (ξ, η) as described in step 4, if any. If there is no such redex then goto step 5. If there is such a redex, choose it among all such redexes (ξ, η) such that if k is the smallest number where η occurs in Γ^k and ℓ is the smallest number where ξ occurs in Γ^ℓ then $k \leq \ell$ (i.e. η occurs before ξ). Reduce (ξ, η). Set $q = q + 1$. Goto step 2.

4. (*Trivial Redex reduction*) Let $\xi, \eta \in \overline{X}$ with the property that $\xi \neq \eta$ and there are paths $\Pi, \Delta, \Sigma \in \{f^1, \ldots, f^k\}^*$ and $i \in \{1, \ldots, n\}$ such that if $t \leq_i u$ is the i-th inequality in Γ^q, then:

$$\Pi(t) = \Delta(t) \in \overline{X} \quad \text{and} \quad \Sigma\Pi(u) = \xi \quad \text{and} \quad \Sigma\Delta(u) = \eta$$

The result of reducing the trivial redex (ξ, η) is a new instance Γ^{q+1} where Γ^{q+1} is obtained from Γ^q by substituting ξ for all occurrences of η throughout Γ^q. Every new occurrence of ξ in Γ^{q+1}, i.e. occurrences of ξ that are not in Γ^q, is assigned the timestamp of the occurrence of η it replaced.

5. Find a redex (ξ, t) as described in step 6 or 7, if any. If there is no such redex, the procedure stops and Γ has a (finite) solution. If there is such a redex (ξ, t), choose it among all such redexes so that the occurrence of ξ has minimum timestamp and reduce it. Set $q = q+1$ and $p = p+1$. Goto step 2.

6. (*Redex I reduction*) Let $\xi \in \overline{X}$ and let $t' \notin \overline{X}$ be a term with the property that there is a path $\Pi \in \{f^1, \ldots, f^k\}^*$ and $i \in \{1, \ldots, n\}$ such that if $t \leq_i u$ is the i-th inequality of Γ^q, then:

   $$\Pi(t) = t' \quad \text{and} \quad \Pi(u) = \xi$$

   The pair of terms $(\xi, (t')_i)$ is called a *redex I*. The result of reducing this redex is a new instance Γ^{q+1} where Γ^{q+1} is obtained from Γ^q by substituting $(t')_i$ for all occurrences of ξ throughout Γ^q. Every new variable occurrence in Γ^{q+1}, i.e. a variable occurrence that is in Γ^q, is assigned the timestamp $p+1$.

7. (*Redex II reduction*) Let $\xi \in \overline{X}$ and u' be a non variable term with the property that $\xi \neq u'$ and there are paths $\Pi, \Delta, \Sigma \in \{f^1, \ldots, f^k\}^*$ and $i \in \{1, \ldots, n\}$ such that if $t \leq_i u$ is the i-th inequality in Γ^q, then:

   $$\Pi(t) = \Delta(t) \in \overline{X} \quad \text{and} \quad \Sigma\Pi(u) = \xi \quad \text{and} \quad \Sigma\Delta(u) = u'$$

   Such a pair (ξ, u') is called a *redex II*. The result of reducing this redex is a new instance Γ^{q+1} where Γ^{q+1} is obtained from Γ^q by substituting u' for all occurrences of ξ throughout Γ^q. Every new variable occurrence in Γ^{q+1}, i.e. a variable occurrence that is not in Γ^q, is assigned the timestamp $p+1$.

Unless otherwise noted, when we refer to the redex procedure we are actually referring to Version 2 of the redex procedure.

Lemma 10 *Every run of the redex procedure on an instance Γ is fair.*

Proof: Given an instance Γ and a run σ of the redex procedure on Γ, for any integer $k \geq 0$, let Γ^k be $R_\sigma^k \cdots R_\sigma^1(\Gamma)$. If σ generates an illegal redex then, by definition, σ is fair. Otherwise, assume Γ_k has a trivial redex (ξ, η). observe that the redex procedure gives priority to trivial redexes to be reduced before other redexes. However, since Γ^k has finitely many potential trivial (renaming) redexes, there exists an integer $\ell \geq 0$ such that $\Gamma^{k+\ell}$ has no trivial redexes and $R_\sigma^{k+\ell}, \ldots, R_\sigma^{k+1}$ are trivial redexes . This implies that ξ and η have the same name in $\Gamma^{k+\ell}$, because otherwise $\Gamma^{k+\ell}$ will have the redex (ξ_1, η_1) where $R_\sigma^\ell \cdots R_\sigma^{k+1}(\xi) = \xi_1$ and $R_\sigma^\ell \cdots R_\sigma^{k+1}(\eta) = \eta_1$.

Assume Γ^k has a non trivial redex (ξ, t) and assume that the timestamp value of this occurrence of ξ is equal to r. Γ^k has finitely many variable occurrences with timestamp value $\leq r$. Let j be the number of such occurrences. Consider the instance $\Gamma^{k+\ell_1}$, as described above, where $\Gamma^{k+\ell_1}$ has no trivial redexes and $R_\sigma^{k+\ell_1}, \cdots, R_\sigma^{k+1}$ are trivial redexes and $\ell_1 \geq 0$. $\Gamma^{k+\ell_1}$ has a redex (ξ', t'), corresponding to the redex (ξ, t) in Γ^k, where $R_\sigma^{k+\ell_1}, \cdots, R_\sigma^{k+1}(\xi) = \xi'$ and $R_\sigma^{k+\ell_1}, \cdots, R_\sigma^{k+1}(t) = t'$. The timestamp value of ξ' is the same as the timestamp value of ξ which is r. Assume the non trivial redex $R_\sigma^{k+\ell_1+1}$ is of the form (η, t''). If $\eta = \xi'$ then we are done. Otherwise, the occurrence of η must have timestamp value $\leq r$, and thus, $\Gamma^{k+\ell_1+1}$ has $j-1$ variable occurrences with timestamp value $\leq r$.

The same argument applies for the instance $\Gamma^{k+\ell_1+\ell_2}$ where $\Gamma^{k+\ell_1+\ell_2}$ has no trivial redexes and $R_\sigma^{k+\ell_2}, \cdots, R_\sigma^{k+\ell_1+2}$ are all trivial redexes and $\ell_1, \ell_2 \geq 0$, and so on. Hence, we can go only finitely many steps before reducing a non trivial redex of the form (ξ'', t^0) because we will run out of variable occurrences with timestamp value $\leq r$. ∎

If the redex procedure runs for finitely many steps, then the associated sequence σ of basic substitutions is finite- say σ is $R_\sigma^1, R_\sigma^2, \cdots, R_\sigma^i$. For technical reasons, we want to view σ as an infinite sequence, and therefore define R_σ^j to be the identity substitution for every $j > i$.

Lemma 11 *Let Γ be an instance of general semi-unification with variables in $\overline{X}$. Let σ be a run of the redex procedure on Γ. Let Γ' be the instance obtained from Γ after evaluating the first $n \geq 1$ basic substitutions of the run σ. If S is a solution of Γ then there is a $P : \overline{X} \to \overline{\mathcal{T}}^*$ such that for any α occurring in Γ:*

$$S(\alpha) = PR_\sigma^n \cdots R_\sigma^1(\alpha)$$

Furthermore, P is a solution for Γ'.

Proof: By induction on n. The base case is when $n = 1$. Assume that for every P there is a $\xi \in \overline{X}$ such that

$$S(\xi) \neq PR_\sigma^1(\xi) \tag{4}$$

$R_\sigma^1(\xi)$ should be of the form (ξ, t^0) where $t^0 \neq \xi$. Notice that if t^0 was equal to ξ, then we are done because we can choose $P(\xi)$ to be the same as $S(\xi)$. So, the corresponding redex must be of the form (ξ, t^0) (we still call it R_σ^1). Let us assume that (ξ, t^0) was a redex I. This means that there is an inequality $t \leq_i u$ in Γ such that $\Pi(t) = t'$ and $\Pi(u) = \xi$ and $t'_i = t^0$. Equation 4 above is true iff for any P

$$S(\xi) \neq P(t^0). \tag{5}$$

Since t^0 is the same as t' with variable renaming, we can conclude that for any P:

$$S(\xi) \neq P(t'). \tag{6}$$

Since $\Pi(u) = \xi$ and $\Pi(t) = t'$. From 6 we can conclude, for any P

$$S(u) \neq P(t)$$

So, in particular

$$S(u) \neq S_i(S(t)).$$

So, S is not a solution for Γ. A similar argument can be used if R^1_σ is a redex II. So, there is a substitution P such that for every α occurring in Γ

$$S(\alpha) = PR^1_\sigma(\alpha).$$

We can also conclude that for every inequality $t \leq_i u$ in Γ

$$S(t) = P(R^1_\sigma(t)) \quad \text{and } S(u) = P(R^1_\sigma(u))$$

Hence, P is a solution for Γ', where Γ' is the result of applying R^1_σ to Γ.

For the induction hypothesis, assume the lemma is true for every $n \leq m$, we need to show that it is true for $n = m+1$. By hypothesis, if S is a solution for Γ, then for any α occurring in Γ there is a P' such that

$$S(\alpha) = P'R^m_\sigma \cdots R^1_\sigma(\alpha). \tag{7}$$

Let Γ' and Γ'' be the instances obtained from Γ after m and $m+1$ redexes of run σ respectively. If R^{m+1}_σ is the identity substitution then the result follows directly, otherwise R^{m+1}_σ is a redex occurring in Γ'. By hypothesis, P' is a solution of Γ' and there is a P'' such that for every α occurring in Γ'

$$P'(\alpha) = P''R^{m+1}_\sigma(\alpha) \tag{8}$$

By Hypothesis, P'' is a solution for Γ''. From equations 7 and 8 we can conclude, for every $\alpha \in X$:

$$S(\alpha) = P''R^{m+1}_\sigma R^m_\sigma \cdots R^1_\sigma(\alpha)$$

which proves the statement for all n. ∎

Definition 12 *Let Γ be an instance of general semi-unification with all variables in the initial X. Let σ be a run of the redex procedure on Γ. For every $\alpha \in X$, define $S^{\omega}_{\sigma}(\alpha)$ as follows. Let $R^{i}_{\sigma} \cdots R^{1}_{\sigma}(\alpha)$ be the term (T^i, φ^i) for every $i \geq 1$. Observe that $T^1 \subseteq T^2 \subseteq \cdots$. If $T^{\omega} = \bigcup T^i$ and $\Pi \in \text{exterior}(T^{\omega})$ then there is an i such that, for every $\Pi \in \text{exterior}(T^j)$ and for every $j \geq i$, $\varphi^j(\Pi) = \varphi^i(\Pi)$. Define the map φ^{ω} on $\text{exterior}(T^{\omega})$ by setting, for every $\Pi \in \text{exterior}(T^{\omega})$*

$$\varphi^{\omega}(\Pi) = \varphi^{i}(\Pi)$$

where i is the smallest integer such that for every $j \geq i$, $\varphi^j(\Pi) = \varphi^i(\Pi)$. Define $S^{\omega}_{\sigma}(\alpha) = (T^{\omega}, \varphi^{\omega})$.

The following Lemma is a direct result from the definition of S^{ω}_{σ}.

Lemma 13 *Let Γ be an instance of general semi-unification with all variables in the initial X. Let σ be a run of the redex procedure on instance Γ.*

1. *For every $\alpha \in X$. If there is a $\Pi \in \{f^1, \ldots, f^k\}^*$ such that $\Pi(S^{\omega}_{\sigma}(\alpha))$ is defined then there is an n such that $\Pi(R^{n}_{\sigma} \cdots R^{1}_{\sigma}(\alpha))$ is defined*

2. *For every $\alpha \in X$. If there is a $\Pi \in \{f^1, \ldots, f^k\}^*$ such that $\Pi(S^{\omega}_{\sigma}(\alpha)) \in \overline{X} \cup C$ then there is an n such that*

$$\Pi(R^{n}_{\sigma} \cdots R^{1}_{\sigma}(\alpha)) = \Pi(S^{\omega}_{\sigma}(\alpha))$$

Lemma 14 *Let Γ be an instance of general semi-unification with all variables in the initial X. Let σ be a run of the redex procedure on instance Γ. If the run σ does not generate an illegal redex, then S^{ω}_{σ} is a solution for Γ.*

Proof: Let Γ^0 be the same as Γ and, for every $n \geq 1$, let Γ^n be $R^{n}_{\sigma}(\Gamma^{n-1})$. Let Γ^{ω} be the instance obtained by applying S^{ω}_{σ} to every variable in Γ. We show that S^{ω}_{σ} is a solution for Γ by establishing that: for every inequality $t \leq_i u$ in Γ^{ω}, *none* of the following cases occurs. Once this is established, it is straightforward to see that S^{ω}_{σ} is a solution for Γ.

1. There is a path $\Pi \in \{f^1, \ldots, f^k\}^*$ and $\Pi(t), \Pi(u) \in \overline{\mathcal{T}}^*$ and $\Pi(t), \Pi(u) \notin \overline{X}$, and

$$\Pi(t) = c \quad \text{and} \quad \Pi(u) \neq c \quad \text{where} \quad c \in C$$

(and conversely for u). If there is such a Π, then by Lemma 13 part 2 there is an l such that the i-th inequality of Γ^l is of the form $t' \leq_i u'$ and $\Pi(t') = c$. If $\Pi(u) = c^1 \in C$ then, by Lemma 13 part 2, there is an m such that the i-th inequality of Γ^m is of the form $t'' \leq_i u''$ and $\Pi(u'') = c^1$. Hence, Γ^k where k is the maximum of (l, m) has an illegal redex, which contradicts the assumption of the lemma. Otherwise, $f^1\Pi(u)$ is defined, by Lemma 13 part 1, there is an m such that the i-th inequality of Γ^m is of the form $t'' \leq_i u''$ and $\Pi(u'') \neq c$. Hence, Γ^k where k is the maximum of (l, m) has an illegal redex, which contradicts the assumption of the lemma.

2. There are paths $\Pi, \Delta, \Sigma \in \{f^1, \ldots, f^k\}^*$ such that $\Sigma\Pi(u), \Sigma\Delta(u) \in \overline{\mathcal{T}}^*$ and $\Sigma\Pi(u), \Sigma\Delta(u) \notin \overline{X}$ and:

   $\Pi(t) = \Delta(t) \in \overline{X}$ and $\Sigma\Pi(u) = c$ and $\Sigma\Delta(u) \neq c$ where $c \in C$.

   By Lemma 13 part 2, if $\Sigma\Pi(u) = c$ then there is an l such that the i-th inequality of Γ^l is of the form $t' \leq_i u'$ and $\Sigma\Pi(u') = c$. If $\Delta\Pi(u) = c^1 \in C$ then, by Lemma 13 part 2, there is an m such that the i-th inequality of Γ^m is of the form $t'' \leq_i u''$ and $\Delta\Pi(u'') = c^1$. Hence, Γ^k where k is the maximum of (l, m) has an illegal redex, which contradicts the assumption of the lemma. Otherwise, if $\Sigma\Pi(u) \notin \overline{X} \cup C$, then $f^1\Sigma\Pi(u)$ is defined. By Lemma 13 part 1, there is an m such that the i-th inequality of Γ^m is of the form $t'' \leq_i u''$ and $\Sigma\Pi(u) \notin \overline{X} \cup C$. Hence, Γ^k where k is the maximum of (l, m) has an illegal redex, which contradicts the assumption of the lemma.

3. There is a $\xi \in \overline{X}$ and a term $t' \notin \overline{X}$ with the property that there is a path $\Pi \in \{f^1, \ldots, f^k\}^*$ and:

   $$\Pi(t) = t' \quad \text{and} \quad \Pi(u) = \xi.$$

   By Lemma 13 part 2, if $\Pi(u) = \xi$ then there is an l such that the i-th inequality of Γ^l is of the form $t' \leq_i u'$ and $\Pi(u') = \xi$. If $\Pi(t) = c^1 \in C$ then, by Lemma 13 part 2, there is an m such that the i-th inequality of Γ^m is of the form $t'' \leq_i u''$ and $\Pi(t'') = c^1$. Hence, Γ^k where k is the maximum of (l, m) has a redex (ξ, c^1). But, since σ is fair and by the assumption of the lemma that σ does not generate an illegal redex, ξ should be replaced by c^1 in some $\Gamma^{l'}$, which is a contradiction. Otherwise, $f^1\Pi(t)$ is defined, by Lemma 13 part 1, there is an m such that the i-th inequality of Γ^m is of the form $t'' \leq_i u''$ and $f^1\Pi(u'')$ is defined. Hence, Γ^k where k is the maximum of l, m has a redex (ξ, v). But, since σ is fair and by the assumption of the lemma that σ does not generate an illegal redex, ξ should be replaced by some term v' in some $\Gamma^{l'}$, which is a contradiction.

4. There is a $\xi \in \overline{X}$ and a term u' with the property that $\xi \neq u'$ and there are paths $\Pi, \Delta, \Sigma \in \{f^1, \ldots, f^k\}^*$ and:

$$\Pi(t) = \Delta(t) \in \overline{X} \quad \text{and} \quad \Sigma\Pi(u) = \xi \quad \text{and} \quad \Sigma\Delta(u) = u'.$$

   By Lemma 13 part 2, if $\Sigma\Pi(u) = \xi$ then there is an l such that the i-th inequality of Γ^l is of the form $t' \leq_i u'$ and $\Sigma\Pi(u') = \xi$. If $\Delta\Pi(u) = \beta \in \overline{X} \cup C$ then, by Lemma 13 part 2, there is an m such that the i-th inequality of Γ^m is of the form $t'' \leq_i u''$ and $\Delta\Pi(u'') = \beta$. Hence, Γ^k where k is the maximum of (l, m) has a redex (ξ, β). But, since σ is fair and by the assumption of the lemma that σ does not generate an illegal redex, ξ should be replaced by β in some $\Gamma^{l'}$, which is a contradiction. Otherwise, if $\Delta\Pi(u) \notin \overline{X} \cup C$, then $f^1\Delta\Pi(u)$ is defined. By Lemma 13 part 1, there is an m such that the i-th inequality of Γ^m is of the form $t'' \leq_i u''$ and $\Sigma\Pi(u'') \notin \overline{X} \cup C$. Hence, Γ^k where k is the maximum of (l, m) has a redex (ξ, v). But, since σ is fair and by the assumption of the lemma, σ does not generate an illegal redex, hence, ξ should be replaced by some term v in some $\Gamma^{l'}$, which is a contradiction. ∎

Theorem 15 *An instance Γ of general semi-unification has no solution iff Γ generates an illegal redex occurrence.*

Proof: The "only if" part is direct from Lemma 14, for any run σ, S_σ^ω is a solution.

For the "if" part, assume first that the initial Γ has an illegal redex occurrence. This could be the case iff there is an inequality of the form $t \leq_i u$ in Γ and one of the following occurs:

1. There is a path $\Pi \in \{f^1, \ldots, f^k\}^*$ and $\Pi(t), \Pi(u) \notin \overline{X}$ such that

$$\Pi(t) = c \in C \quad \text{and} \quad \Pi(u) \neq c.$$

S cannot substitute for a constant, and in this case Γ has no solution. The case when

$$\Pi(u) = c \in C \quad \text{and} \quad \Pi(t) \neq c$$

is treated similarly.

2. There are $\Pi, \Delta, \Sigma \in \{f^1, \ldots, f^k\}^*$ and $\Sigma\Pi(u), \Sigma\Delta(u) \notin \overline{X}$ such that:

$$\Pi(t) = \Delta(t) = x \in \overline{X} \quad \text{and} \quad \Sigma\Pi(u) = c \in C \quad \text{and} \quad \Sigma\Delta(u) \neq c.$$

S_i can substitute for x only one value, so whatever S_i substitutes for x $S_i(S(t))$ will not be equal to $S(u)$.

Now, assume that Γ has no illegal redex occurrence and Γ' was obtained from Γ after n steps of the redex procedure but Γ' has an illegal redex occurrence. From Lemma 11 we can conclude that if Γ has a solution then Γ' will have a solution. But, Γ' has an illegal redex occurrence which means that it has no solution. Hence, Γ has no solution. ∎

We now extend the ordering $\sqsubseteq_\Gamma$ given in Section 2 (Definition 2) and consider it on the set of all general solutions for an instance Γ. $\sqsubseteq_\Gamma$ is now defined as follows:

Definition 16 $S' \sqsubseteq_\Gamma S''$ *iff there is a substitution* $P : \overline{X} \to \overline{\mathcal{T}}^*$ *such that* $P(S'(\xi)) = S''(\xi)$ *for every variable* ξ *occurring in* Γ *.*

A solution S of Γ is *principal* iff for every other solution S', $S \sqsubseteq_\Gamma S'$.

Corollary 17 *Let* Γ *be an instance of semi-unification. Let* σ *be a run of the redex procedure on* Γ*. If* Γ *has a solution then* S^ω_σ *is a principal solution for* Γ*.*

Proof: From Theorem 15 and Lemma 14, we can conclude that if an instance Γ has a solution then S^ω_σ is also a solution for Γ. Assume that there is a solution S' for Γ such that there is an $\alpha \in X$ and for any P

$$S'(\alpha) \neq P(S^\omega_\sigma(\alpha))$$

hence, there is a $\Pi \in \{f^1, \ldots, f^k\}^+$ such that $\Pi(P(S^\omega_\sigma(\alpha)))$ is defined and

$$\Pi(S'(\alpha)) \neq \Pi(P(S^\omega_\sigma(\alpha)))$$

where either $\Pi(S'(\alpha)) \in \overline{X} \cup C$ or $\Pi(P(S^\omega_\sigma(\alpha)) \in \overline{X} \cup C$. If $\Pi(P(S^\omega_\sigma(\alpha))) \in \overline{X} \cup C$ then, by Lemma 13 part 2, there is an n such that

$$\Pi(S'(\alpha)) \neq \Pi(P(R^n_\sigma \cdots R^1_\sigma(\alpha)))$$

Hence, for any P

$$S'(\alpha) \neq P(R^n_\sigma \cdots R^1_\sigma(\alpha))$$

which contradicts Lemma 11. Otherwise, if $\Pi(P(S^\omega_\sigma(\alpha))) \notin \overline{X} \cup C$ and $\Pi(S'(\alpha)) \in \overline{X} \cup C$ then, by Lemma 13 part 1, there is an n such that

$$\Pi(S'(\alpha)) \neq \Pi(P(R^n_\sigma \cdots R^1_\sigma(\alpha)))$$

Hence, for any P

$$S'(\alpha) \neq P(R^n_\sigma \cdots R^1_\sigma(\alpha))$$

which contradicts Lemma 11. Hence, S^ω_σ is a principal solution. ∎

Corollary 18 *An instance Γ with no constants always has a (general) solution.*

Proof: If an instance Γ has no constants, then Γ does not generate an illegal redex. Hence, any such Γ has a solution. ∎

3.2 Equational Theory for General Semi-Unification

We represent general semi-unification as a problem in the equational theory as defined in Section 2.2, where Σ_{EQ} is now an arbitrary finite set of at least 2 unary function symbols, $\Sigma_{EQ} = \{f^1, f^2, \ldots, f^k\}$ and $k \geq 2$. Let Γ be an instance of semi-unification, F a function symbol of arity k. Let X_Γ and C_Γ be the set of variables and constants occurring in Γ respectively. We consider an operation on Γ, denoted $|\Gamma|$, the result of which is another instance Γ_s. Γ_s has no constants, so it always has a general solution. The variables occurring in Γ_s are $X_\Gamma \cup C_\Gamma$. Γ_s is obtained from Γ as follows. First we define a term t_Γ which mentions every member in C_Γ and mentions no members of X.

Assume $C_\Gamma = \{c_1, \ldots, c_j\}$. A suitable choice for t_Γ is given by setting for every $i \in \{1, \ldots, j\}$

$$f^1 \underbrace{f^2, \ldots, f^2}_{i-1}(t_\Gamma) = c_i$$

and

$$f^3 \underbrace{f^2 \cdots f^2}_{i-1}(t_\Gamma) = \cdots = f^k \underbrace{f^2 \cdots f^2}_{i-1}(t_\Gamma) = c_1$$

and setting $\underbrace{f^2 \cdots f^2}_{j}(t_\Gamma) = c_1$.

Let x be a fresh variable mentioned no where in Γ. If $t_i \leq u_i$ was the ith inequality in Γ then the ith inequality of Γ_s is $t'_i \leq u'_i$ where t'_i is defined as follows

$$f^1(t'_i) = t_i \quad \text{and} \quad f^2(t'_i) = t_\Gamma \quad \text{and for } 2 < j \leq k \quad f^j(t'_i) = x$$

and u'_i is defined as

$$f^1(u'_i) = u_i \quad \text{and} \quad f^2(u'_i) = t_\Gamma \quad \text{and for } 2 < j \leq k \quad f^j(u'_i) = x$$

We sometimes refer to Γ_s as the stripped version of Γ.

Lemma 19 *Let σ be a run of the redex procedure on Γ_s. Γ has a solution iff for every $c \in C_\Gamma$ either $S^\omega_\sigma(c) = c_w$ or $S^\omega_\sigma(c) = \beta$, where $w \in \{1, 2, \ldots, n\}^*$ and $\beta \in \overline{X_\Gamma}$.*

Proof: The proof follows from the following 3 claims.

- Claim 1: Γ has an illegal redex occurrence iff Γ_s has a redex of the form (c, t) where $c \in C_\Gamma$ and $t \neq c_w$ and $t \notin \overline{X}_\Gamma$.

  For the "only if" part, Γ has an illegal redex iff there is an inequality of the form $t \leq_i u$ in Γ and one of the following occurs:

  1. There is a path $\Pi \in \{f^1, \ldots, f^k\}^*$ and $\Pi(t), \Pi(u) \notin \overline{X}_\Gamma$ such that

     $$\Pi(t) = c \in C \quad \text{and} \quad \Pi(u) \neq c.$$

     In this case Γ_s will have a redex Π (c, t') where $t' = \Pi(u)$.

2. There is a path $\Pi \in \{f^1, \ldots, f^k\}^*$ and $\Pi(t), \Pi(u) \notin \overline{X}_\Gamma$ such that

$$\Pi(u) = c \in C_\Gamma \quad \text{and} \quad \Pi(t) \neq c.$$

   Here, Γ_s will have a redex I (c, t') where $t' = (\Pi(u))_i$.

3. There are $\Pi, \Delta, \Sigma \in \{f^1, \ldots, f^k\}^*$ and $\Sigma\Pi(u), \Sigma\Delta(u) \notin \overline{X}_\Gamma$ such that:

$$\Pi(t) = \Delta(t) = x \in \overline{X}_\Gamma \quad \text{and} \quad \Sigma\Pi(u) = c \in C_\Gamma \quad \text{and} \quad \Sigma\Delta(u) \neq c.$$

   Again, Γ_s will have a redex II (c, t') where $t' = \Sigma\Delta(u)$.

For the "if" of the claim, it is easy to check that any redex (c, t') where $c \in C_\Gamma$ and $t \neq c_w$ and $t \notin \overline{X}_\Gamma$ matches one of the cases for illegal redexes above.

We now introduce a special strategy of picking a redex from a stripped instance. Notice that a sequence of redexes picked using this strategy corresponds to a sequence of redexes obtained using a run of the redex procedure. The strategy tries to pick a redex using step 1 first until there are no redexes that satisfy the condition of step 1 then it proceeds to pick a redex using step 2. The strategy consists of the following steps:

1. If there is a redex of the form (c_i, c), where $c \in C_\Gamma$ and $i \in \{1, 2, \ldots, n\}$, pick it to be reduced.
2. If a redex was of the form (c, β) where $c \in C_\Gamma$ and $\beta \in \overline{X}_\Gamma$ then pick its symmetric redex to be reduced, i.e. the redex (β, c).

Notice that such a strategy will only affect the renaming of variables and hence, it still has the properties as a run of the redex procedure, which proves the correctness of our next claim.

- Claim 2: For any run of the redex procedure on Γ_s, $S^\omega_\sigma(c) = t$, where $t \neq c_w$ and $t \notin \overline{X}_\Gamma$ iff there is a sequence σ' of redexes obtained using the above strategy such that $R^n_{\sigma'} \cdots R^1_{\sigma'}(\Gamma_s) = \Gamma'_s$ and Γ'_s has a redex (c, t') such that $t' \neq c_w$ and $t' \notin \overline{X}_\Gamma$.

- Claim 3: Γ generates an illegal redex occurrence, i.e. there is a Γ' obtained from Γ such that Γ' has an illegal redex iff $|\Gamma'| = \Gamma'_s$ and

$R^n_{\sigma'} \cdots R^1_{\sigma'}(\Gamma_s) = \Gamma'_s$ where σ' is a sequence of redexes obtained using the above strategy.

To proof this claim we will show that for any run σ of the redex procedure on Γ there is a corresponding run σ' on Γ_s, using the above strategy, and vice versa, such that

$$R^n_{\sigma} \cdots R^1_{\sigma}(\Gamma) = \Gamma'$$

and

$$R^m_{\sigma'} \cdots R^1_{\sigma'}(\Gamma_s) = \Gamma'_s$$

and $|\Gamma'| = \Gamma'_s$. Let us start with Γ, assume R^1_{σ} was of the form (ξ, t). Γ_s initially has no redexes satisfying the conditions of step 1 of the above strategy. So, we pick $R^1_{\sigma'}$ to be the same as R^1_{σ}. Now, we apply the strategy until we obtain an instance Γ^0_s such that Γ^0_s has no redexes satisfying step 1 of the strategy. It is not hard to see that $|R^1_{\sigma}(\Gamma)| = \Gamma^0_s$. We repeat the same procedure for R^2_{σ} and so on.

For the other direction, we start from the stripped version Γ_s. Again, Γ_s has no redexes satisfying the conditions of step 1 of the above strategy. So, we pick $R^1_{\sigma'}$ to be of the form (ξ, t) where $\xi \in \overline{X}_{\Gamma}$. Now, we apply the strategy until we obtain an instance Γ^0_s such that Γ^0_s has no redexes satisfying step 1 of the strategy. Γ has the same redex (ξ, t) so we pick R^1_{σ} to be this redex. $|R^1_{\sigma}(\Gamma)| = \Gamma^0_s$. So, we repeat the same step for R^2_{σ}. ∎

We now construct a set of equations E_{Γ_s} based on the instance Γ_s, similar to the construction in Section 2.3. The set of variables in E_{Γ_s} is the same as the set of variables occurring in Γ_s which is $X_{\Gamma} \cup C_{\Gamma}$. We extend $X_{\Gamma} \cup C_{\Gamma}$ to get $\overline{X_{\Gamma} \cup C_{\Gamma}}$ whose elements are symbols of the form α_w where $\alpha \in X_{\Gamma} \cup C_{\Gamma}$ and $w \in \{1, 2, \ldots, n\}^*$. To construct E_{Γ_s} from the set of inequalities of Γ_s add the following set of equations for every inequality of the form $t \leq_i u$:

1. For every $\Pi, \Lambda \in \Sigma^*_{EQ}$ and $\alpha, \beta \in X_{\Gamma} \cup C_{\Gamma}$, such that $\Pi\Lambda(t) = \alpha$ and $\Lambda(u) = \beta$, add the following equation to E_{Γ_s}:

$$\alpha_i = \Pi\beta$$

2. For every $\Pi, \Lambda \in \Sigma_{EQ}^*$ and $\alpha, \beta \in X \cup C_\Gamma$, such that $\Pi\Lambda(u) = \alpha$ and $\Lambda(t) = \beta$, add the following equation to E_{Γ_s}:

$$\alpha = \Pi\beta_i$$

Notice that this construction is identical to the construction in Section 2.3, the only difference being that Σ_{EQ} is not restricted here to two elements. This does not alter the proof of Lemma 4. So, we now give an extended version of Lemma 4.

Lemma 20 *Let Γ be an instance of general semi-unification with no constants.*

1. *If e is derivable from E_Γ then it is satisfied be every solution of Γ.*
2. *Let Γ' be obtained from Γ by reducing some number of redexes, and let t be the term assigned to $\alpha \in X$ in Γ'. For every $\Pi \in \{L, R\}^*$ and $\beta_w \in \overline{X}$ such that $\Pi(t) = \beta_w$, $E_\Gamma \vdash (\Pi\alpha = \beta_w)$.*

Proof: The proof of the second part of this lemma is the same as Lemma 4, the proof of the first part is a routine induction on the length of a derivation of an equation. Again, we refer the reader to the proof of Lemma 4 in [13] for the details of the proof.

Theorem 21 *Let Γ be an instance of general semi-unification, and let Γ_s be its stripped version. Γ has no solution iff one of the following is true.*

1. *There is a $c \in C_\Gamma$, $\Pi \in \Sigma_{EQ}^+$ and $\beta_w \in \overline{X_\Gamma \cup C_\Gamma}$ such that*

$$E_{\Gamma_s} \vdash (\Pi c = \beta_w)$$

2. *There are two distinct $c, c^1 \in C$ such that*

$$E_{\Gamma_s} \vdash (c = c^1)$$

Proof: For the "if" part, assume that an equation of the form $(\Pi c = \beta_w)$ was derived from E_{Γ_s}, by Lemma 20, this equation must be satisfied by every solution of Γ_s. In particular $S^\omega(c) \neq c$ and is not a variable. Hence, by Lemma 19 Γ has no solution. The same reasoning applies for the second case.

For the "only if" part, assume Γ has no solution. Then, by Lemma 19, there is a constant $c \in C_\Gamma$ such that for any run σ of the redex procedure on Γ_s $S^\omega_\sigma(c) = t$ where t could either be a constant c^1 or a term t and there is a $\Pi \in \Sigma^+_{EQ}$ and $\beta_w \in \overline{X_\Gamma \cup C_\Gamma}$ such that $\Pi(t) = \beta_w$. By part 2 of Lemma 20, E_{Γ_s} derives an equation of the form $\Pi c = \beta_w$ or an equation of the form $c = c^1$. ∎

4 Database Dependencies

We now consider Database dependencies and their relation to general semi-unification. In what follows we adapt the definitions of [3]. A relation scheme is an object $R[U]$, where R is the name of the relation scheme and U a finite set of attributes $\{A_1, A_2, \cdots, A_n\}$. A tuple t over U is a function from U to the elements of some basic domain. t could be restricted to some subset X of U, which is denoted by $t[X]$. A relation r over U is a set of tuples over U. For U, X, r as above $r[X] = \{t[X] : t \in r\}$ is the projection of r on X.

Let $U_i \subseteq U$ for $1 \leq i \leq q$, a *database scheme* Δ is a set of relation schemes $\{R_1, R_2, \cdots, R_q\}$, and a *database* $d = \{r_1, r_2, \cdots, r_q\}$ associates each relation scheme $R_i[U_i]$ with a nonempty relation r_i over U_i.

Assume $R[U]$ is a relation scheme. Let A and B be nonempty subsets of U. $R : A \rightarrow B$ is a *functional dependency* (FD). A relation r satisfies $R : X \rightarrow Y$ if, whenever t_1, t_2 are tuples of r with $t_1[A]{=}t_2[A]$, then $t_1[B] = t_2[B]$. When A and B are sets with one attribute each, we say $R : A \rightarrow B$ is a *unary* FD.

Let $R[U]$ and $S[U']$ be two relation schemes over Δ. $\{(A_1, B_1), \cdots, (A_n, B_n)\}$ is a set of ordered pairs of attributes, with A_i from U and B_i from U'. $R.A_1A_2\cdots A_n \subseteq S.B_1B_2\cdots B_n$ is called an *inclusion dependency* (IND). Relations r over $R[U]$ and s over $S[U']$ satisfy this dependency if for each $t \in r$ there is a $t' \in s$ with $t[A_i] = t'[B_i], 1 \leq i \leq n$. Let A and B be two attributes over Δ.We write $A \equiv B$ as an abbreviation of the dependency $R : AB \subseteq AA$ and we call such a dependency *equivalence dependency*.

Let Δ be a database scheme, Γ a set of dependencies over Δ, and γ a dependency over Δ. Γ implies γ , denoted by $\Gamma \models \gamma$, if whenever a database d over a scheme Δ satisfies Γ it also satisfies γ. In what follows we assume that Δ has only one relation scheme, and we omit the relation name from the dependencies.

Given an instance Γ of restricted dependency Δ, we construct a set E_Γ of equations for an equational theory EQ in the following way. Let Σ_E be a finite set of unary function symbols. Each element of Σ corresponds to one unary functional dependency in Γ. Let n be the number of inclusion dependencies in Γ. Let X_E be a finite set of variables. Each element of X_E corresponds to one attribute in Δ. We extend X_E to get $\overline{X_E}$, and elements of $\overline{X_E}$ are of the form α_w where $\alpha \in X_E$ and $w \in \{1, 2 \cdots, n\}^*$. Let $\overline{\mathcal{T}}$ be the set of finite terms over Σ and $\overline{X_E}$. E_Γ is a set of equations defined over $\overline{\mathcal{T}}$. The set of equations E_Γ consists of the following equations

- For every functional dependency FD_i of the form

  $$A \rightarrow B,$$

  add $f^i A = B$.

- For every inclusion dependency IND_i of the form

  $$A^1 \dots A^m \subseteq B^1 \dots B^m,$$

  add m equations $A_i^1 = B^1$ and $\dots$ and $A_i^m = B^m$.

Theorem 22 *Let Γ be an instance of restricted dependency,*

1. $\Gamma \models A \equiv B$ *iff* $E_\Gamma \vdash A = B$.

2. $\Gamma \models A \rightarrow B$ *iff* $E_\Gamma \vdash \Pi A = B$ *where* $\Pi \in \Sigma^+$.

3. $\Gamma \models A^1 \dots A^m \subseteq B^1 \dots B^m$ *iff* $E_\Gamma \vdash (A_w^1 = B^1)$*and*$\dots$ *and* $E_\Gamma \vdash (A_w^m = B^m)$ *where* $w \in \{1, 2 \dots, n\}^+$.

Proof: This follows directly from theorem 1 in [2]. When only unary FDs are considered, the equational theory described there is equivalent to the equational theory EQ described here. Terms of the equational theory in [2] could be represented as terms as described here. Unary functions corresponding to FDs are the same. Unary functions corresponding to INDs are represented as indecies here. For example given a term ia where i is an IND and a is a variable it is represented as a_i here. This representation eliminates the need for the *commutativity rule* given in [2]. Now, given this correspondence between the two equational theories, it is readily checked that the inference rules in both generate an equivalent set of equations. ∎

For the remainder of this section we only consider equivalence implication problems, i.e. we only consider implication problems of the form $\Gamma \models A \equiv B$, for some attributes A and B. Notice that Γ still contains INDs and unary FDs. We still call this case restricted dependency. This special case of dependency is shown to be undecidable in [2]. We show that restricted dependency is equivalent to general semi-unification. To accomplish this we transfer a set of equations E in EQ that correspond to an instance Γ of restricted dependency to a set of equations E' that correspond to an instance Γ' of general semi-unification and vice versa.

4.1 From Restricted Dependency to General Semi-Unification

Given an instance Γ of restricted dependency, assume that we are given an implication problem of the form $\Gamma \models A \equiv B$, for some attributes A and B. Construct the set E_Γ of equations as described above. By Theorem 21, the above implication is true iff

$$E_\Gamma \vdash A = B.$$

Construct another set of equations E'_Γ where $\Sigma_{E'_\Gamma} = \Sigma_E$, and $X_{E'_\Gamma} = X_E \cup \{c^1, c^2\}$ where $c^1, c^2 \notin X_E$, by adding the following equations to E_Γ :

$$A_{n+1} = c^1$$

$$B_{n+1} = c^2$$

Lemma 23 *$\Gamma \models A \equiv B$ iff $E'_\Gamma \vdash c^1 = c^2$.*

Proof: Since c^1 and c^2 only occurred in the two equations above, $E'_\Gamma \vdash c^1 = c^2$ iff $E'_\Gamma \vdash A_{n+1} = B_{n+1}$ iff $E'_\Gamma \vdash A = B$ iff $E_\Gamma \vdash A = B$. ∎

From E'_Γ we construct another set of equations E''_Γ. E''_Γ will look exactly like a set of equations obtained from an instance of general semi-unification. To construct E''_Γ add the following equations to E'_Γ:

1. For every variable $B \in X_{E'_\Gamma}$, if there is an equation of the form $fB = D$ in E'_Γ, add an equation of the form $gB = N^g$ for every $g \in \Sigma_{E'_\Gamma}$, where N^g is a fresh variable. Notice that this only adds redundant equations to E''_Γ.

2. Replace every equation e in E''_Γ of the form $fB = D$ by the following two equations:

$$fB = R^e_j$$
$$R^e_j = D$$

   where $j = n + 2$ and R^e is a fresh variable.

3. for every $1 \leq i \leq n + 2$ add the equations:

$$c^1_i = c^1 \quad \text{and} \quad c^2_i = c^2.$$

Lemma 24 $\Gamma \models A \equiv B$ *iff* $E''_\Gamma \vdash c^1 = c^2$.

Proof: A similar reasoning as in Lemma 23. Given a set of equations E''_Γ, we can construct an instance Γ' of general semi-unification that will correspond to E''_Γ. Γ' will have $n + 2$ inequalities. It will have only two constants c^1, c^2. The set of variables occurring in Γ' is $X_{E''_\Gamma} - \{c^1, c^2\}$.

From the above we can conclude the following lemma.

Lemma 25 $\Gamma \models A \equiv B$ *iff the corresponding instance of general semi-unification* Γ' *has no solution.*

Lemma 26 [2] *Let* Γ *be an instance of restricted dependency. Given two attributes* A *and* B *in* Γ*, it is undecidable whether:*

$$\Gamma \models A \equiv B.$$

This directly leads us to the undecidability result:

Theorem 27 *It is undecidable whether an instance of general semi-unification has a solution.*

4.2 From General Semi-Unification to Restricted Dependency

Given an instance Γ of general semi-unification, construct the set of equations E_Γ as described in Section 3. Recall that $X_{E_\Gamma} = X_\Gamma \cup C_\Gamma$. By Theorem 10, Γ has no solution iff one of the following is true.

1. There is a $c \in C_\Gamma$, $\Pi \in \Sigma_{EQ}^{+}$ and $\beta_w \in \overline{X_\Gamma \cup C_\Gamma}$ such that

$$E_\Gamma \vdash (\Pi c = \beta_w).$$

2. There are two distinct $c, c^1 \in C_\Gamma$ such that

$$E_\Gamma \vdash (c = c^1).$$

Construct another set of equations E'_Γ from E_Γ by adding the following equations:

- For every variable $\alpha \in X_{E_\Gamma}$ add an equation:

  $$g\alpha = c'$$

  where $g \notin \Sigma_{E_\Gamma}$ and c' is a new variable.

- For every $c \in C_\Gamma$ and $f \in \Sigma_{E_\Gamma}$ add an equation:

  $$fc = b$$

  where b is a new variable.

- For every $f \in \Sigma_{E_\Gamma}$ add an equation:

  $$fb = b$$

- Add the equation:

  $$gb = c''$$

  where c'' is a new variable.

- For every $1 \leq i \leq n$, where n is the number of inequalities in Γ add the equations:

  $$c'_i = c' \quad \text{and} \quad c''_i = c''$$

Lemma 28 *Γ has no solution iff one of the following is true.*

1. $E'_\Gamma \vdash c' = c''$.

2. *There are two distinct $c, c^1 \in C_\Gamma$ such that*

$$E'_\Gamma \vdash (c = c^1).$$

Proof: The second part of the lemma is direct: $E'_\Gamma \vdash (c = c^1)$ iff $E_\Gamma \vdash (c = c^1)$. For the first part, $E'_\Gamma \vdash c' = c''$ iff $E'_\Gamma \vdash g\beta_w = c''$, where $\beta_w \in \overline{X_\Gamma \cup C_\Gamma}$ iff $E'_\Gamma \vdash g\beta_w = gb$ iff there is a $\Pi \in \Sigma^+_{EQ}$ such that $E'_\Gamma \vdash g\Pi c = g\beta_w$ iff $E'_\Gamma \vdash \Pi c = \beta_w$ iff $E_\Gamma \vdash \Pi c = \beta_w$. ∎

Given the set of equations E'_Γ, we simplify it as follows:

1. Let $\Pi \in \Sigma^*_{E'_\Gamma}$, replace every equation of the form $\Pi\alpha_i = \beta$ by the two equations

$$\alpha_i = x \quad \text{and} \quad \Pi x = \beta$$

where x is a new variable

2. Let $\Pi \in \Sigma^*_{E'_\Gamma}$, $f \in \Sigma_{E'_\Gamma}$, $\alpha, \beta \in X_{E'_\Gamma}$, replace every equation of the form $f\Pi\alpha = \beta$. by the following two equations

$$y = \Pi\alpha \quad \text{and} \quad fy = \beta.$$

where y is a new variable. Repeat step 2 until no longer applicable.

Theorem 29 *The restricted dependency problem is equivalent to general semi-unification.*

Proof: By Lemma 25 we can reduce an instance of restricted dependency to general semi-unification. Let $E_{\Gamma''}$ be the set of equations obtained from $E_{\Gamma'}$ by applying the above simplifications to $E_{\Gamma'}$. $E_{\Gamma''}$ looks exactly like a set of equations obtained from an instance of restricted dependency. So given such equations we can easily go back to an instance Γ' of restricted dependency. Γ has no solution iff there are $c^1, c^2 \in C_\Gamma \cup \{c', c''\}$ such that $\Gamma' \models c^1 \equiv c^2$. ∎

5 Regular Semi-Unification

In this case of semi-unification we are still allowed to substitute an infinite term for a variable, however, such an infinite term must be regular. Let Σ be the first-order signature described in Section 3 consisting of a countably infinite set of constants C and one function symbol F of arity $k \geq 1$. Recall that a term t is a pair (T, φ) as defined in Section 3. Let $sub(t)$ be the set of subterms of t. More formally, $sub(t)$ is defined as:

$$sub(t) = \{\Pi(t) | \Pi \in \{f^1, \ldots, f^k\}^*\}$$

Definition 30

- *Let* $\mathcal{T}_{reg}$ *be the set of all terms* (T, φ) *such that* $\mathrm{sub}(t)$ *is finite.*

In this section, a substitution S is a function $S : X \to \mathcal{T}_{reg}$. As in Section 3, S can be extended to a function $S : \mathcal{T}_{reg} \to \mathcal{T}_{reg}$. An *instance* Γ of *regular* semi-unification is a finite set of inequalities:

$$\Gamma = \{t_1 \leq u_1, \ldots, t_n \leq u_n\}$$

where $t_i, u_i \in \mathcal{T}_{fin}$. A substitution S is a *solution* of the instance Γ iff there are substitutions $S_1, \ldots, S_n$ such that:

$$S_1(S(t_1)) = S(u_1), \ \ldots \ , \ S_n(S(t_n)) = S(u_n)$$

The *regular semi-unification problem* is the problem of deciding, for any such instance Γ, whether Γ has a solution.

5.1 Feature Algebra

Feature Algebra is a special type of Algebra considered in computational linguistics. The problem relevant to our discussion is, given a set of constraints containing a subsumption preorder, does there exist a *finite* feature algebra satisfying these constraintsΓ

This problem has been shown to be undecidable [4] by a reduction from the word problem for finite semi-groups [5]. We now give an outline of the steps of this reduction. This outline is needed to prove a result later in the report.

- Given a finite alphapet Δ, let E be a finite set of equations $s_i = t_i, i = 1, \cdots, n$ where s_i and t_i are nonempty strings over Δ. Let $s = t$ be another such equation.

- Consider the class Λ of all finite semi groups finitely generated by Δ. The word problem for finite semi-groups is to determine whether every finite semi-group $\in \Lambda$ satisfying all the equations in E also satisfies the equation $s = t$.

- Given a set of equations E and another equation e, construct a set of constraints C_E (the details of constructing C_E are not relevant here, we refer the reader to [4] for the complete construction).

- C_E has a solution in a finite feature algebra iff there is a finite semi-group satisfying all the equations in E but not the equation e.

Also, in [4] the problem of whether a set of constraints has a solution in a finite feature algebra is shown to be equivalent to regular semi-unification[5]. This proves the undecidability of regular semi-unification.

In the above problem, if we eliminate the restriction that the feature algebra is finite, i.e. the problem now becomes, given a set of constraints containing a subsumption preorder, does there exist a feature algebra satisfying these constraintsΓ This problem can also be shown undecidable by a reduction from the word problem of semi-groups using the same construction in [4]. Similarly, this problem can be shown to be equivalent to general semi-unification, again, using the same reduction given in [4] and outlined above. This provides an alternative way of proving the undecidability of general semi-unification. This also presents a way of proving a recursive inseparability result between regular semi-unification and general semi-unification, see Section 7.1 of this report for the details.

5.2 Regular Semi-Unification and the Redex Procedure

Observe that regular terms are a subset of general terms, in other words $\mathcal{T}_{reg} \subset \mathcal{T}^*$. Hence, given an instance Γ of regular semi-unification we can

5

In [4] they refer to regular semi-unification as semi-unification on rational trees run the redex procedure, presented in Section 3, on Γ. The redex procedure still has a property similar to that given by Lemma 11. This means that the redex procedure has the property that the initial instance Γ has a regular solution iff the instance obtained after m steps Γ_m has a regular solution. The problem is that we can not guarantee that the redex procedure in its present form will stop if Γ has a regular solution. Also, we can not guarantee that it will stop in all cases in which Γ has no regular solution. In its present form, the redex procedure stops if Γ has no general solution. It also stops if Γ has a finite solution.

On the other hand, we can give a procedure that stops if Γ has a regular solution. This procedure consists of enumerating regular substitutions in a certain order and at the same time testing whether the current substitution is a solution for Γ. If the substitution is a solution then the procedure stops, otherwise it enumerates another substitution. We do not give details here, but it is readily checked that this procedure is effective. This means that, by contrast to the general case, we can not modify our redex procedure to make it stop if Γ has no regular solution. Otherwise, regular semi-unification would be decidable for the reason that we can run the redex procedure on Γ and at the same time run the procedure defined in this paragraph ; eventually one of the two procedures would stop (of course, assuming that the redex procedure could be modified in this way). Nevertheless, it would be interesting to find an instance that does not stop on either of these procedures. This instance would provide an example of an instance which has a general solution but not a regular one. We were not able to come up with such an instance. However, we know that there is such an instance. We refer the reader to Sections 7.1 and 7.2 of this report for results related to the discussion here.

6 Common Properties of Semi-Unification

For simplicity, in finite semi-unification, we chose to give a standard definition of terms. This made an instance of finite semi-unification different from an instance of general (or regular) semi-unification. An instance of finite semi-unification can be viewed as an instance of general semi-unification and vice versa, and likewise for an instance of regular semi-unification. Therefore, we can simply talk about an instance Γ of semi-unification (without qualification) and ask whether Γ has a finite,regular or general solution; corresponding to the 3 Problems of semi-unification discussed in this report. In this section we discuss some of the properties that are common to the the three problems of semi-unification.

6.1 Instances with a Binary Function Symbol

We have defined an instance of semi-unification on a signature with an arbitrary k-ary function symbol. By a standard construction we can reduce an instance Γ_k with arity $k \geq 2$ to an instance Γ_2 with a binary function symbol, such that Γ_k has a general (regular, finite) solution iff Γ_2 has a general (resp. regular, finite) solution.

Notice that from an instance Γ_k, with a k-ary function symbol, we can construct another instance Γ_{k+1}, with a $k+1$-ary function symbol such that Γ_k has a solution iff Γ_{k+1} has a solution. One possible way to construct Γ_{k+1} is by having for every inequality of the form $t \leq_i u$ in Γ_k a corresponding inequality $t' \leq_i u'$ in Γ_{k+1} where:

$$f^j(t') = f^j(t) \quad \text{for } 1 \leq j \leq k \quad \text{and} \quad f^{k+1} = x$$

for some fresh variable x. So, we assume, without loss of generality that $k = 2^j$ for some j.

For a number $0 \leq i < k$, i is represented by f^{i+1}. Assume we are counting to the base k. It is standard that any digit i could be uniquely represented by a binary number $B(i)$ on the digits 0,1 (i.e. f^1, f^2 in our representation); the length of such a number is j. Let $t = (T, \varphi)$ be a term on a signature with $k > 2$. We first consider an operation on the tree part of t, the result of which denoted $[T]_k$ is a binary tree. More formally $[\ ]_k : \{f^1, \ldots, f^k\}^* \to \{f^1, f^2\}^*$ defined recursively as follows.

$$[\varepsilon]_k = \varepsilon$$

$$[\Pi f^i]_k = [\Pi]_k B(f^i), \qquad \text{for } i = 1, \ldots, k$$

Let $[T]_k$ be the least binary tree containing the set $\{[\Pi]_k | \Pi \in T\}$.

Lemma 31 *The operation* $[\ ]_k$ *is a bijection from* exterior(T) *to* exterior($[\mathrm{T}]_\mathrm{k}$).

Proof: Any $\Pi \in \mathit{exterior}(T)$ is uniquely represented by $[\Pi]_k$. Hence, $[\ ]_k$ is one-to-one. Notice that any $\Pi \in [T]_k$ could be written as $\Pi = \Pi_2[\Pi_1]_k$ where

$|\Pi_2| < j$ and $\Pi_1 \in T$. Using the definition of $[T]_k$, if $|\Pi_2| > 0$ then for every $i \in \{1, \ldots, k\}$, $f^i\Pi_1 \in T$, otherwise $[T]_k$ will not be minimal. Hence, there is a Π_3 such that $\Pi_3\Pi_2 = f^i$ for some $i \in \{1, \ldots, k\}$. In this case Π cannot be in $exterior([T]_k)$. So, if $\Pi \in exterior([T]_k)$ it has to be of the form $\Pi = [\Pi_1]_k$ where $\Pi_1 \in T$ and $\Pi_1 \in exterior(T)$. Hence, $[\,]_k$ is a bijection. ■

Define $[\varphi]_k : exterior([T]_k) \to X \cup C$ by setting $[\varphi]_k(\Pi) = \varphi([\Pi]_k)$. For a term $t = (T, \varphi)$ define $[t]_k = ([T]_k, [\varphi]_k)$. For an instance Γ of k-ary semi-unification, we define $[\Gamma]_k$ as follows. If $t \leq_i u$ is the i-th inequality in Γ then $[t]_k \leq_i [u]_k$ is the i-th inequality in $[\Gamma]_k$. Note that $[\Gamma]_k$ is an instance of binary semi-unification. Given a substitution S on Γ, we define a substitution $[S]_k$ on $[\Gamma]_k$ such that if $S(\alpha) = t$ then $[S]_k(\alpha) = [t]_k$. The following lemma illustrates the relation between S and $[S]_k$.

Lemma 32 *Let S be a substitution on a k-ary signature and let $[S]_k$ be the corresponding substitution on a binary signature. For any term t over the k-ary signature*

$$[S(t)]_k = [S]_k([t]_k).$$

Proof: Let $t = (T, \varphi)$, $S(t) = t_1 = (T_1, \varphi_1)$, using the definition of a substitution in Section 3 we have:

$$T_1 = T \cup \{\Pi\Delta | \Delta \in exterior(T), \varphi(\Delta) = x, S(x) = (T_2, \varphi_2), \Pi \in T_2\}$$

Notice that $\Pi \in exterior(T_1)$ iff either $\Pi \in exterior(T)$ and $\varphi(\Pi) = c$, or $\Pi = \Sigma\Delta$ such that $\Delta \in exterior(T)$, $\varphi(\Delta) = x$, $S(x) = (T_2, \varphi_2)$ and $\Sigma \in exterior(T_2)$. By Lemma 31, $\Pi \in exterior([T_1]_k)$ iff $\Pi = [\Pi_1]_k$ and either $\Pi_1 \in exterior(T)$ and $\varphi(\Pi_1) = c$, or $\Pi_1 = \Sigma\Delta$ such that $\Delta \in exterior(T)$, $\varphi(\Delta) = x$, $S(x) = (T_2, \varphi_2)$ and $\Sigma \in exterior(T_2)$. Again by Lemma 31 and by definition of φ_k and the definition of $[S]_k$ we have, $\Pi \in exterior([T_1]_k)$ iff $\Pi = [\Pi_1]_k$ and either $[\Pi_1]_k \in exterior([T]_k)$ and $\varphi_k([\Pi_1]_k) = c$, or $\Pi_1 = \Sigma\Delta$ such that $[\Delta] \in exterior([T]_k)$, $\varphi_k([\Delta]_k) = x$, $[S(x)]_k = ([T_2]_k, [\varphi_2)]_k$ and $[\Sigma]_k \in exterior([T_2]_k)$. If $[S]_k([t]_k) = (T', \varphi')$, we conclude that $[T_1]_k = T'$ because they have the same exterior. Using a similar argument, we conclude that $(\varphi_1)_k = \varphi'$. Hence, $[S(t)]_k = [S]_k([t]_k)$. ■

Lemma 33 *Let Γ be an instance of k-ary semi-unification. $(S, S^1, \ldots, S^n)$ is a general (regular, finite) solution for Γ iff $([S]_k, [S^1]_k, \ldots, [S^n]_k)$ is a general (resp. regular, finite) solution for $[\Gamma]_k$.*

Proof: For any two terms t, u over the k-ary signature $t = u$ iff $[t]_k = [u]_k$. $(S, S^1, \ldots, S^n)$ is a general (regular,finite) solution for Γ iff for any inequality $t^i \leq u^i$ in Γ, $S^i(S(t^i)) = S(u^i)$ iff $[S^i(S(t^i))]_k = [S(u^i)]_k$, using Lemma 32, iff $[S^i]_k([S]_k([t^i]_k)) = [S]_k([u^i]_k)$ iff $([S]_k, [S^1]_k, \ldots [S^n]_k)$ is a general (resp. regular,finite) solution for $[\Gamma]_k$. ∎

A term t on a binary signature is in *k-standard form* if it is in the range of the function $[\ ]_k$. For a variable α, we define a special kind of variables, associated with α, such a variable is of the form α^{Π} where $\Pi \in \{f^1, f^2\}^*$, $\alpha^{\varepsilon} = \alpha$. A similar type of constants is introduced for a constant c.

Given a binary term $t = (T, \varphi)$ and a $k = 2^j$, we consider an operation *k- standard*(t) the result of which is a binary term $t' = (T', \varphi')$ where T' is the least binary tree (it is unique) in k-standard form such that $T' \subseteq T$. And for every $\Pi = \Pi_2\Pi_1 \in exterior(T')$ where $\Pi_1, \Pi_2 \in \{f^1, f^2\}^*$ and $\Pi_1 \in exterior(T)$ and $\varphi(\Pi_1) = a$

$$\varphi'(\Pi) = a^{\Pi_2}$$

Given a substitution S from variables to binary terms, we define *k-standard*(S) such that if $S(\alpha) = t$ then *k-standard*(S)(α)= *k-standard* (t)

Lemma 34 *If $[\Gamma]_k$ has a general (regular, finite) solution then $[\Gamma]_k$ has a general (resp. regular, finite) solution S' such that for any variable α, $S'(\alpha) = [t]_k$ where t is a term over the k-ary signature.*

Proof: In the case where $[\Gamma]_k$ has a general or finite solution the proof is straightforward. Recall that if Γ has a general or finite solution then it will have one generated by the redex procedure, see Section 3 and 2 respectively. Hence, the proof follows from the following claim. Any substitution generated by the redex procedure is of the form (α, t') where t' is in k-standard form. To show the correctness of this claim, first observe that initially any inequality in $[\Gamma]_k$ is of the form $t \leq_i u$ where both t and u are in k-standard form. A term t' is in k-standard form if $t' = a$, where a is a variable or a constant, or for every Π where $|\Pi| = |j|$ (recall that $k = 2^j$), $\Pi(t')$ is defined and in k-standard form. So, for every term v in $[\Gamma]_k$ (i.e. $v = t^i$ or $v = u^i$ for some inequality i), if $\Pi(v) = \alpha$ for some variable α then, for every other term v' in $[\Gamma]_k$, $\Pi(v')$ is in k-standard form.

The case when $[\Gamma]_k$ has a regular solution follows from the following claim:

> If S is a regular solution for $[\Gamma]_k$ then *k-standard*(S) is a regular solution for $[\Gamma]_k$.

To show that k-$standard(S)$ is regular, it suffices to show that, if t is a regular term then k-$standard(t)$ is a regular term. By definition, a term t is regular if $sub(t)$ is finite where:

$$sub(t) = \{\Pi(t) | \Pi \in \{f^1, f^2\}^*\}$$

But, by a standard argument, $sub(t)$ is finite iff the set $sub_j(t)$ is finite where:

$$sub_j(t) = \{\Pi(t) | \Pi \in \{f^1, f^2\}^* \text{ and } |\Pi| \text{ is a multiple of } j\}$$

From the definition of the operation k-$standard()$ we can conclude that $\Pi(k$-$standard(t))$= k-$standard(\Pi(t))$ for $|\Pi|$ multiples of j. Hence, $sub_j(t)$ is finite iff $sub_j(k$-$standard(t))$ is finite. So, k-$standard(t)$ is regular if t is regular.

For every inequality $t \leq_i u$ in $[\Gamma]_k$ let $t' \leq_i u'$ be the corresponding ith inequality in k-$standard(S)([\Gamma]_k)$. To verify that k-$standard(S)$ is indeed a solution for $[\Gamma]_k$, we show that the following is true for every inequality of the form $t' \leq_i u'$ in the instance k-$standard(S)([\Gamma]_k)$:

1. For every path $\Pi \in \{f^1, f^2\}^*$ if $\Pi(t') = c^{\Pi_2}$ then $\Pi(u') = c^{\Pi_2}$.

   If $\Pi(t') = c^{\Pi_2}$ then there is a Π_1 such that $\Pi_2\Pi_1 = \Pi$ and $\Pi_1(S(t)) = c$. But since S is a solution for $[\Gamma]_k$ $\Pi_1(S(u)) = c$. Hence, $\Pi(u') = c^{\Pi_2}$.

2. If there are paths $\Pi, \Delta \in \{f^1, f^2\}^*$ such that, $\Pi(t') = \Delta(t') = \alpha^{\Pi_2} \in X$ then $\Pi(u') = \Delta(u')$.

   If $\Pi(t') = \Delta(t') = \alpha^{\Pi_2}$ then $\Pi_1(t) = \Delta_1(t) = \alpha$ where $\Pi = \Pi_2\Pi_1$ and $\Delta = \Delta_1\Pi_2$. Hence $\Pi_1(S(u)) = \Delta_1(S(u))$. From the definition of k-$standard()$ and because S is a solution for $[\Gamma]_k$, we conclude that $\Pi(u') = \Delta(u')$,

3. If there is a path $\Pi \in \{f^1, f^2\}^*$ such that $\Pi(u') = c^{\Pi_2}$, for some constant c then either $\Pi(t') = \Pi(u')$ or $\Pi'(t') = \alpha^{\Delta}$ where $\Pi = \Pi''\Pi'$ and α is a variable.

   If $\Pi(u') = c^{\Pi_2}$ then $\Pi_1(S(u)) = c$ where $\Pi = \Pi_2\Pi_1$. If S is a solution for $[\Gamma]_k$ then either $\Pi_1(S(t)) = c$, in this case $\Pi_2(t') = c^{\Pi_2}$, or $\Pi_1'(S(t)) = \alpha$ where $\Pi_1 = \Pi_1{}''\Pi_1{}'$ and α is a variable. But Π_1' could at most have the same length as Π and in this case $\Pi(t') = \alpha$ otherwise $\Pi'(t') = \alpha^{\Pi_3}$ where $\Pi_3\Pi_1' = \Pi'$.

4. If there is a path $\Pi \in \{f^1, f^2\}^*$ such that $\Pi(u') = \alpha^{\Pi_2}$ for some variable α, then $\Pi'(t') = \alpha^{\Delta}$ where $\Pi = \Pi''\Pi'$ and α is a variable.

   This case is treated similar to case 3. ∎

Theorem 35 *Let Γ be an instance of k-ary semi-unification. Γ has a general (regular,finite) solution iff the corresponding instance of binary semi-unification $[\Gamma]_k$ has a general (resp. regular,finite) solution.*

Proof: Follows directly from Lemmas 33 and 34. ∎

6.2 Instances with two inequalities

Given an instance Γ with $n > 2$ inequalities, we construct an instance Γ' with 2 inequalities such that Γ has a solution iff Γ' has one. This construction is originally due to [18]. We give a slightly modified version of it. We put all the details in because the previous reference is not readily available. The result due to [18] is for a finite solution. We use a similar construction to obtain the result for general and regular solutions also.

Let Γ be an instance of binary semi-unification consisting of the following n inequalities. $\{t^1 \leq u^1, \ldots, t^n \leq u^n\}$. Let $X_\Gamma = \{x^1, \ldots, x^j\}$ be the set of variables occurring in Γ. Define the terms t' and u' where

$$f^1 \underbrace{f^2, \ldots, f^2}_{i-1}(t') = (t^i)_i, \quad \text{for } i = 1, \ldots, n-1$$

$$\underbrace{f^2, \ldots, f^2}_{n-1}(t') = (t^n)_n$$

and

$$f^1 \underbrace{f^2, \ldots, f^2}_{i-1}(u') = (u^i)_i, \quad \text{for } i = 1, \ldots, n-1$$

$$\underbrace{f^2, \ldots, f^2}_{n-1}(u') = (u^n)_n$$

For every variable $\alpha \in X_\Gamma$ define the terms t^α and u^α where

$$f^1 \underbrace{f^2, \ldots, f^2}_{i-1}(t^\alpha) = (\alpha)_i, \quad \text{for } i = 1, \ldots, n-1$$

$$\underbrace{f^2, \ldots, f^2}_{n-1}(t^\alpha) = (\alpha)_n$$

and

$$f^1 \underbrace{f^2, \ldots, f^2}_{i-1}(u^\alpha) = (\alpha)_{i+1}, \quad \text{for } i = 1, \ldots, n-1$$

$$\underbrace{f^2, \ldots, f^2}_{n-1}(u^\alpha) = (\alpha)_1$$

Define the terms t'' and u'' where

$$f^1 \underbrace{f^2, \ldots, f^2}_{i-1}(t'') = t^{x^i}, \quad \text{for } i = 1, \ldots, j-1$$

$$\underbrace{f^2, \ldots, f^2}_{j-1}(t'') = t^{x^j}$$

and

$$f^1 \underbrace{f^2, \ldots, f^2}_{i-1}(u'') = u^{x^i}, \quad \text{for } i = 1, \ldots, j-1$$

and

$$\underbrace{f^2, \ldots, f^2}_{j-1}(u'') = u^{x^j}$$

Let Γ' be the instance consisting of the following two inequalities

$$\{t' \leq u', t'' \leq u''\}$$

Theorem 36 *An instance of binary semi-unification Γ with $n \geq 2$ inequalities has a general (regular, finite) solution iff the corresponding instance Γ' with 2 inequalities has a general (resp. regular, finite) solution.*

Proof: The proof here is not effected by the type of the solution i.e. general, regular or finite so it works for all the cases.

For the "if" part, assuming Γ' has a solution, there are substitutions S, S' and S'' such that

$$S'(S(t')) = S(u') \quad \text{and} \quad S''(S(t'')) = S(u'') \tag{9}$$

By the second part of 9 we can conclude that, for every $\alpha \in X_\Gamma$, $S''(S(t^\alpha)) = S(u^\alpha)$. From the definition of t^α and u^α we have

$$S''(S(\alpha)) = S((\alpha)_1), S''(S((\alpha)_1)) = S((\alpha)_2), \ldots, S''(S((\alpha)_n) = S(\alpha)$$

For every $i \in \{1, \ldots, n\}$, let P^i be the substitution obtained by composing the substitution S'', i times, more formally,

$$P^i = \underbrace{S'', \ldots, S''}_{i}$$

We conclude that, for every $\alpha \in X_\Gamma$ and for every $i \in \{1, \ldots, n\}$, $P^i(S(\alpha)) = S((\alpha)_i)$. Now, Observe that $(t)_i$ is obtained from t by indexing every variable in t by the index i so, we get, for every $i \in \{1, \ldots, n\}$:

$$P^i(S(t^i)) = S((t^i)_i) \tag{10}$$

From 9 we have, for every $i \in \{1, \ldots, n\}$, $S'(S(t^i)_i) = S(u^i)$. So, using 10 we have, for every $i \in \{1, \ldots, n\}$, $S'(P^i(S(t^i))) = S(u^i)$. Hence, S is a solution for Γ.

For the "only if" part, if Γ has a solution then there are substitutions $S, S^1, \ldots, S^n$ such that, for every $i \in \{1, \ldots, n\}$:

$$S^i(S(t^i)) = S(u^i)$$

We are assuming that indexed variables are not used in any term in Γ and so, we could assume that for any substitution P on Γ $P(\alpha)_i$, for every $i \in \{1, \ldots, n\}$, is not defined.

Let $P, P^1, \ldots, P^n$ be substitutions defined on Γ' where

$$P(\alpha_i) = S((\alpha))_i$$

for every $i \in \{1, \ldots, n\}$, and $P(\alpha) = S(\alpha)$. And, for every $i \in \{1, \ldots, n\}$, P^i is defined as follows:

$$P^i(\alpha_i) = S^i((\alpha))$$

For every variable α and for every $i \in \{1, \ldots, n\}$ we have

$$P^i(P((\alpha)_i)) = S^i(S(\alpha)) \tag{11}$$

For every $i \in \{1, \ldots, n\}$, the term $(t^i)_i$ only contains variables indexed by the index i so, using 11, we have :

$$P^i(P((t^i)_i)) = S^i(S(t^i)) \tag{12}$$

Observe that for every $i \in \{1, \ldots, n\}$, $P(u^i) = S(u^i)$. Hence for every $i \in \{1, \ldots, n\}$,

$$P^i(P((t^i)_i)) = S^i(S(t^i)) = S(u^i) = P(u^i)$$

Let P_1 be a substitution defined as : $P_1(\alpha_i) = P^i(\alpha_i)$. Notice that P^i is defined on variables indexed by the index i and hence is disjoint from any other P^j for $i \neq j$. So, for every $i \in \{1, \ldots, n\}$, we have:

$$P_1(P(t^i)_i) = P(u^i)$$

From this we conclude that $P_1(P(t')) = P(u')$. For the second inequality, let P_2 be the substitution defined as, for every variable $\alpha \in X_\Gamma$,

$$P_2(\alpha) = (\alpha)_1, P_2((\alpha)_1) = (\alpha)_2, \ldots, P_2((\alpha)_n) = \alpha$$

This plus the definition of P implies:

$$P_2(P(\alpha)) = P((\alpha)_1), P_2(P((\alpha)_1)) = P((\alpha)_2), \ldots, P_2(P(\alpha)_n)) = P(\alpha)$$

From this we conclude that $P_2(P(t'')) = P(u'')$. Hence, P is a solution for Γ'. ∎

6.3 Decidable Cases

All 3 problems of semi-unification are decidable when $k = 1$. This is due to the fact that there is only one infinite term over such a signature. This term is a regular term. Hence, Γ has a general solution iff it has regular one. This makes all 3 problems decidable, the reason for that is given later in this report.

As in the case when $k = 1$, all 3 problems of semi-unification are decidable for $n = 1$. The proof in the finite case is given in [10]. In the other cases this is due to the fact that if an instance Γ with one inequality has a general solution then it will have a regular one. The proof of this fact in this case is a bit more involved than the case with $k = 1$. Other decidable cases for finite semi-unification are discussed in [7, 11, 15].

7 Comparing The Problems of Semi-Unification

In this section we present a comparative study between the three cases of semi-unification. We study the properties of the solution set for each case. We also discuss the principality property for each case. In addition, we study the set of solvable instances in each case of semi-unification and compare the three sets. In particular, we give a recursive inseparability result between instance that have a regular solution and instances that do not have a general solution.

As a result of the discussion given in Section 6 (Theorems 36 and 35), without loss of generality, we now consider an instance of semi-unification to be an instance on a binary signature consisting of at most two inequalities. We denote this function symbol by $\rightarrow$ written in infix notation. We sometimes look at terms as labeled binary trees.

7.1 Principality Property and the Solution Set

Let Γ be an instance of semi-unification. Define $\mathcal{S}^{\Gamma}_{gen}$ to be the set of all general solutions for Γ. Similarly, define $\mathcal{S}^{\Gamma}_{reg}$ to be the set of all regular solutions for Γ and define $\mathcal{S}^{y}_{fin}\Gamma$ to be the set of all finite solutions for Γ. We now list some observations about the solutions sets. We give some informal reasoning for some of these observations.

1- $\mathcal{S}^{y}_{fin}\Gamma \subseteq \mathcal{S}^{\Gamma}_{reg} \subseteq \mathcal{S}^{\Gamma}_{gen}$.

2- $\mathcal{S}^{y}_{fin}\Gamma$ and $\mathcal{S}^{\Gamma}_{reg}$ are both recursive, because given a finite (resp. regular) substitution S, we can decide whether S is a solution for Γ. Not every element of $\mathcal{S}^{\Gamma}_{gen}$ has a finite description, and therefore cannot decide whether a given general substitution is a member of $\mathcal{S}^{\Gamma}_{gen}$.

We now redefine the ordering $\sqsubseteq_{\Gamma}$ on all substitutions from $\overline{\mathcal{T}}^{*}$ to $\overline{\mathcal{T}}^{*}$. The definition of $\sqsubseteq_{\Gamma}$ here differs from Definition 16 in a technical way only.

Definition 37 *$S' \sqsubseteq_{\Gamma} S''$ iff there is a substitution $P : \overline{X} \rightarrow \overline{\mathcal{T}}^{*}$ such that $P(S'(\xi)) = S''(\xi)$ for every variable ξ occurring in Γ.*

A general (regular, finite) solution S is a principal solution for Γ iff for every solution $S' \in \mathcal{S}^{\Gamma}_{gen}$ (resp. $\in \mathcal{S}^{\Gamma}_{reg}$, $\in \mathcal{S}^{y}_{fin}\Gamma$), $S \sqsubseteq_{\Gamma} S'$. Notice that the principal solution is unique up to variable renaming in Γ. For any of the three cases of semi-unification, the principality property holds iff for every instance Γ, if Γ has a solution (in that case of semi-unification) then it has a principal solution.

3- The principality property holds in finite and general semi-unification. In the finite case the proof is given in [6, 13]. For the general case the proof is given in this report (Corollary 17). In the regular case, it is still open whether the principality property holds.

4- In the finite case, given an instance Γ and a finite solution S, we can decide whether S is a principal solution for Γ. The reason is that we can run the redex procedure on Γ until it stops because we know that Γ has a finite solution. When the redex procedure stops, we compare the solution S and the solution generated by the redex procedure S_0. Now, we can decide whether S is principal by simply checking whether the S is equivalent to S_0 up to variable renaming.

   In the regular case, given an instance Γ and a regular solution S, we cannot decide in general whether S is a principal solution for Γ. The proof is a bit more involved and we give it in Lemma 38. This means that, even if the principality property holds in the regular case, we cannot construct such a principal solution (provided that there is a principal solution). This is a partial answer for the question about the principality property in the regular case.

5- Observe that, by contrast to first order unification, for general (regular, finite) semi-unification: it is *not* the case that, for an instance Γ, if $S \in \mathcal{S}^{\Gamma}_{gen}$ (resp. $S \in \mathcal{S}^{\Gamma}_{reg}$, $S \in \mathcal{S}^{y}_{fin}\Gamma$) then for any general (resp. regular, finite) substitution S' such that $S \sqsubseteq_{\Gamma} S'$, $S' \in \mathcal{S}^{\Gamma}_{gen}$ (resp. $S' \in \mathcal{S}^{\Gamma}_{reg}$, $S' \in \mathcal{S}^{y}_{fin}\Gamma$). For an example, let us consider an instance Γ consisting of one inequality $\{t \leq u\}$, where $t = (x \to y)$ and $u = (c \to x)$. Consider the substitution S_I that maps every variable in Γ to itself (the identity substitution). Clearly, S_I is a finite solution, and hence, a regular and general solution, for Γ, (in this case the matching substitution would map x to c and y to x). Furthermore, S_I is a principal solution for Γ. For any substitution S, $S_I \sqsubseteq_{\Gamma} S$. Let S be the substitution that is the same as S_I except for x, where x is now mapped to the term

$(x \rightarrow x)$. It is readily checked that S is not a finite , and not a regular or general, solution for Γ. That is why the principal solution in semi-unification is different from the most general unifier in first-order unification, where if S is a most general unifier for a certain set of pairs then any substitution instance of the most general unifier is a solution (in the sense of unification) for that set of pairs.

For an instance Γ, Let $S_\Gamma^\top$ be the substitution where every variable in Γ is mapped to the full binary infinite tree. Clearly, $S_\Gamma^\top$ is a regular substitution.

Lemma 38 *Let Γ be an instance of semi-unification. Assume that Γ has a regular solution and assume that the principality property holds for regular semi-unification . Given a regular solution S for Γ we cannot decide whether S is a principal solution for Γ.*

Proof: The proof is by contradiction. Assume that we can decide whether S is a principal solution for Γ. The claim is, if we can decide this, then regular semi-unification would be decidable. First, we consider a new instance Γ' where we view the constants of Γ as variables. Γ' has a regular solution. In particular, $S_{\Gamma'}^\top$ is a solution for Γ'. We can recursively enumerate all solution for Γ' (see Section 5.2). Whenever we enumerate a solution we ask whether it is a principal solution for Γ'. This means that, eventually, we will be able to construct the principal regular solution S' for Γ'. Given the substitution S', consider the instance $S'(\Gamma')$ obtained by applying S' to all inequalities in Γ'. We can decide whether the original Γ has a solution by simply looking at the instance $S'(\Gamma')$ and observing whether in this instance a constant c is replaced by a non variable term or whether for there is an inequality in $\S'(\Gamma')$ such that every matching substitution for this inequality maps a constant c to a term not equal to c. More formally, for every inequality of the form $t_i \leq u_i$ in Γ' we need to check for the following:

- there is a path Π where $\Pi(t_i) = c$ and $\Pi(S'(t_i))$ is a non variable term such that
$$\Pi(S'(t_i)) \neq \Pi(t_i)$$
- there is a path Π where $\Pi(u_i) = c$ and $\Pi(S'(u_i))$ is a non variable term such that
$$\Pi(S'(u_i)) \neq \Pi(u_i)$$

- there is a path Π where $\Pi(S'(t_i)) = c$ and $\Pi(S'(u_i))$ is a term such that

$$\Pi(S'(u_i)) \neq \Pi(S'(t_i))$$

If one of the above is satisfied, then we know that Γ has no solution. Otherwise a trivial variation S'' on S' will be a solution for Γ. S'' is the same as S' except for the cases where a constant c is mapped to a variable x , in this case S'' will map x to c. ∎

6- We now focus on instances that have no constants. Recall that instances with no constants always have a general and a regular solution. In particular, for any instance Γ, $S_\Gamma^\top$ is a regular and general solution for Γ. The following holds for any such instance Γ:

  - $\mathcal{S}_{fin}^y\Gamma$ has a bottom $-$ with respect to the ordering $\sqsubseteq_\Gamma$. Where for every $S \in \mathcal{S}_{fin}^y\Gamma$, $- \sqsubseteq_\Gamma S$. Notice that, we can use the principal solution for $\mathcal{S}_{fin}^y\Gamma$ as $-$. In this case $-$ is a member of $\mathcal{S}_{fin}^y\Gamma$. $\mathcal{S}_{fin}^y\Gamma$ has a top $\top$ where for every $S \in \mathcal{S}_{fin}^y\Gamma, \S \sqsubseteq_\Gamma \top$. We can use $S_\Gamma^\top$ as $\top$ for $\mathcal{S}_{fin}^y\Gamma$ but notice that a top for $\mathcal{S}_{fin}^y\Gamma$ cannot be a member of $\mathcal{S}_{fin}^y\Gamma$.
  - $\mathcal{S}_{reg}^\Gamma$ has a bottom $-$ with respect to the ordering $\sqsubseteq_\Gamma$. Where for every $S \in \mathcal{S}_{reg}^\Gamma$, $- \sqsubseteq_\Gamma S$. Notice that, in this case, we were not able to find $-$ that is a member of $\mathcal{S}_{reg}^\Gamma$. However, we can use the principal general solution as a bottom for $\mathcal{S}_{reg}^\Gamma$. $\mathcal{S}_{reg}^\Gamma$ has a top $\top$ where for every $S \in \mathcal{S}_{reg}^\Gamma, S \sqsubseteq_\Gamma \top$. We can use $S_\Gamma^\top$ as a top for $\mathcal{S}_{reg}^\Gamma$. Notice the top of $\mathcal{S}_{reg}^\Gamma$ is a member of $\mathcal{S}_{reg}^\Gamma$.
  - $\mathcal{S}_{gen}^\Gamma$ has a bottom $-$ with respect to the ordering $\sqsubseteq_\Gamma$, where for every $S \in \mathcal{S}_{gen}^\Gamma$, $- \sqsubseteq_\Gamma S$. Notice that we can use the principal general solution as a bottom for $\mathcal{S}_{gen}^\Gamma$. $\mathcal{S}_{gen}^\Gamma$ has a top $\top$ where for every $S \in \mathcal{S}_{gen}^\Gamma$ $S \sqsubseteq_\Gamma \top$. We can use $S_\Gamma^\top$ as a top for $\mathcal{S}_{gen}^\Gamma$. Notice that the top and bottom of $\mathcal{S}_{gen}^\Gamma$ are members of $\mathcal{S}_{gen}^\Gamma$.

7.2 Solvable Instances

In this subsection we study the set of instances that have a solution (solvable instances) in each case of semi-unification. Let Ξ_{gen} (Ξ_{reg}, Ξ_{fin}) be the set of instances that have a general (resp. regular, finite) solution.

- Ξ_{fin} is recursively enumerable. The procedure to enumerate this set is straightforward. We can generate all instances [6] in a certain order. And we run every instance we generate (using dove-tailing) on the redex procedure. If the redex procedure stops, because there are no more redexes, on a certain instance Γ then we list Γ.

- Ξ_{reg} is also recursively enumerable. Recall that we can give a procedure that halts if an instance has a regular solution. Again, we can write a procedure, similar to the finite case, that enumerates all instances that are solvable in the regular case.

- Ξ_{gen} is not recursively enumerable. However, its complement $\overline{\Xi}_{gen}$ is recursively enumerable. Again, we generate all instances and run them on the redex procedure. If the redex procedure stops, because there is an illegal redex (see Theorem 15), on a certain instance Γ then we list Γ.

Notice that $\Xi_{fin} \subset \Xi_{reg} \subset \Xi_{gen}$. This implies that there are instances that have a general solution but not a regular one. Otherwise, general and regular semi-unification will both be decidable, because Ξ_{reg} and $\overline{\Xi}_{gen}$ are both recursively enumerable. In what follows ,we extend the relation between Ξ_{reg} and $\overline{\Xi}_{gen}$.

Two (disjoint) sets A and B are *recursively inseparable* if there is no recursive set containing A and disjoint from B [20]. The following lemma is taken from [5].

Lemma 39 *Let Δ range over alphabets containing the symbols 0 and A_0, and let x_i and y_i range over words in Δ^*. Let ϕ range over formulas of the form $x_1 = y_1 \wedge, \ldots, \wedge x_n = y_n \Longrightarrow A_0 = 0$. The following sets are recursively inseparable:*

1. *$\{\phi \mid \phi$ holds in every Δ generated semigroup $\}$,*

2. *$\{\phi \mid \phi$ fails in some finite Δ generated semigroup $\}$.*

Theorem 40 *$\overline{\Xi}_{gen}$ and Ξ_{reg} are recursively inseparable, i.e. instances that have no general solution are recursively inseparable from instances that have a regular solution.*

[6] Recall that we only consider binary instances with at most two inequalities.

Proof: As mentioned in the outline given in Section 5.1, there is a Γ' such that ϕ fails in some finite semigroup iff Γ' has a regular solution. Also, using the same construction in [4], there is a Γ'' such that ϕ holds in every semigroup iff Γ'' has no general solution. Hence, for every ϕ which is a member of the set mentioned in Part 1 of Theorem 39, there is a corresponding $\Gamma'' \in \overline{\Xi}_{gen}$. Similarly, for every ϕ which is a member of the set mentioned in Part 2 of Theorem 39, there is a corresponding $\Gamma' \in \Xi_{reg}$. Hence, if we can recursively separate Ξ_{reg} from $\overline{\Xi}_{gen}$ then we can recursively separate the sets mentioned in Parts 1 and 2 of Theorem 39. ∎

Proposition 41 *Let $\mathcal{T}'$ be a set of terms where $\mathcal{T}_{reg} \subseteq \mathcal{T}' \subseteq \mathcal{T}^*$. Consider the case of semi-unification obtained by allowing substitutions to replace variables by members of $\mathcal{T}'$, i.e. a substitution S in this case is a map from X to $\mathcal{T}'$.*

> *Every semi-unification problem extending regular semi-unification in this sense is undecidable.*

Proof: Let Ξ' be the set of solvable instances in this new case. Clearly, $\Xi_{reg} \subseteq \Xi'$. Hence, if we can recursively separate Ξ' and $\overline{\Xi}_{gen}$, we can use the same recursive set that separates Ξ' and $\overline{\Xi}_{gen}$ to separate Ξ_{reg} and $\overline{\Xi}_{gen}$. This implies the recursive inseparability of Ξ' and $\overline{\Xi}_{gen}$. Which implies the undecidability of any case of semi-unification obtained by such an extension of a regular substitution. ∎

References

[1] A. K. Chandra and M. Y. Vardi. The implication problem for functional and inclusion dependencies is undecidable. *SIAM Journal of Computing*, 14(3):761 – 677, 1985.

[2] S.S. Cosmadakis and P.C. Kanellakis. Equational theories and database constraints. In *Proceedings of the 17th Annual ACM Symposium on Theory of Computing*, pages 273 – 284, 1985.

[3] S.S. Cosmadakis and P.C. Kanellakis. Functional and inclusion dependencies: A graph theoretic approach. In P.C. Kanellakis and F. Preparata, editors, *Advances in Computing Research*, pages 164 – 185. JAI Press, 1986.

[4] J. Dörre and W. Rounds. On subsumption and semi-unification in feature algebras. In *Proceedings of IEEE 5th Annual Symposium on Logic in Computer Science*, pages 300 – 310, 1990.

[5] Yuri Gurevich and Harry R. Lewis. The word problem for cancellation semigroups with zero. *Association for Symbolic Logic*, 49:184 – 191, March 1984.

[6] F. Henglein. Type inference and semi-unification. In *ACM Conference on Lisp and Functional Programming*, pages 184 – 197, July 1988.

[7] F. Henglein. Fast left-linear semi-unification. In *Proceedings of International Conference on Computing and Information*, pages 32 – 36, Toronto, May 1990. Canadian Scholar's Press Inc.

[8] F. Henglein. Type inference with polymorphic recursion. *ACM Transactions on Programming Languages and Systems*, 15(2):254–290, 1993.

[9] P.K. Hooper. The undecidability of the Turing machine immortality problem. *Journal of Symbolic Logic*, 31(2):219 – 234, 1966.

[10] D. Kapur, D. Musser, P. Narendran, and J. Stillman. Semi-unification. In Nori and Kumar, editors, *Proceedings of 8th Conference on Foundations of Software Technology and Theoretical Computer Science, LNCS 338*, pages 435 – 454. Springer Verlag, 1988.

[11] A.J. Kfoury, J. Tiuryn, and P. Urzyczyn. Computational consequences and partial solutions of a generalized unification problem. In *Proceedings of IEEE 4th Annual Symposium on Logic in Computer Science*, pages 98 – 105, 1989.

[12] A.J. Kfoury, J. Tiuryn, and P. Urzyczyn. An analysis of ML typability. In Arnold, editor, *15th Colloquium on Trees in Algebra and Programming, CAAP 90, LNCS 431*, pages 206 – 220. Springer Verlag, 1990.

[13] A.J. Kfoury, J. Tiuryn, and P. Urzyczyn. The undecidability of the semi-unification problem. In *Proceedings of the 22nd Annual ACM Symposium on Theory of Computing, Baltimore*, pages 468 – 476, 1990.

[14] A.J. Kfoury, J. Tiuryn, and P. Urzyczyn. Type-reconstruction in the presence of polymorphic recursion. *ACM Transactions on Programming Languages and Systems*, 15(2):290–311, 1993.

[15] H. Leiss. Decidability of semi-unification in two variables. Research Report INF 2-ASE-9-89, Siemens, München, 1989.

[16] H. Leiß and F. Henglein. A decidable case of the semi-unification problem. In *Proc. 16th Int'l Symp. on Mathematical Foundations of Computer Science (MFCS), Poland*. Springer, Sept. 1991. Lecture Notes in Computer Science, Vol. 520.

[17] J.C. Mitchell. The implication problem for functional and inclusion dependencies. *Information and Control*, 56(3):154 – 173, 1983.

[18] P. Pudlák. On a unification problem related to Kreisel's conjecture. *Commentationes Mathematicae Universitatis Carolinae, Prague, Czechoslovakia*, 29(3):551 – 556, 1988.

[19] P.W. Purdom. Detecting looping simplifications. In Lescanne, editor, *Proceedings of 2nd Conference on Rewriting Techniques and Applications, LNCS 250*, pages 54 – 62. Springer Verlag, 1987.

[20] Hartley Rogers. *Theory of Recursive Functions and Effective Computability*. McGraw-Hill, 1967.