# UNIVERSAL UNIFICATION AND REGULAR EQUATIONAL ACFM THEORIES

J. Siekmann, P. Szab6

Universitst Karlsruhe
Institut fur Informatik I
D-7500 Karlsruhe, Postfach 6380

ABSTRACT. A universal minimal type conformal matching algorithm and a universal minimal unification algorithm based on [FA79], [SL74], [LB79], [HU80] are presented for a restricted class of equational theories (the Regular ACFM Theories), i.e. it is shown that the set of most general unifiers is recursively enumerable for this class. The class of Regular ACFM Theories is wide enough to contain most special cases of unification algorithms that have been investigated so far. This paper is a (very) abbreviated version of [SS81c], all proofs and most of the technical material are omitted for lack of space. For reasons of consistency with the original paper, the numbering of definitions, lemmas and theorems has been retained.

## 1.1 INTRODUCTION

*Unification theory* is concerned with problems of the following kind: Let f and g be function symbols, a and b constants and let x and y be variables and consider two *first order terms* built from these symbols; for example:

$$t_1 = f(x,g(a,b))$$
$$t_2 = f(g(y,b),x).$$

The first question which arises is whether or not there exist terms which can be substituted for the variables x and y such that $t_1$ and $t_2$ become equal: in the example g(a,b) and a are two such terms. We shall write

$\sigma_1 = \{x \leftarrow g(a,b),\ y \leftarrow a\}$ for such a unifying substitution: $\sigma_1$ is a *unifier* of $t_1$ and $t_2$ since $\sigma_1 t_1 = \sigma_1 t_2$.

In addition to the *decision problem* there is also the problem of finding a *unification algorithm* which generates the unifiers for a given pair $t_1$ and $t_2$.

Consider a variation of the above problem, which arises when we assume that f is commutative:

(C) $f(x,y) = f(y,x)$

Now $\sigma_1$ is still a unifying substitution and moreover $\sigma_2 = \{y \leftarrow a\}$ is also a unifier for $t_1$ and $t_2$, since

$\sigma_2 t_1 = f(x,g(a,b)) =_C f(g(a,b),x) = \sigma_2 t_2$.

But $\sigma_2$ is *more general* than $\sigma_1$, since $\sigma_1$ is an instance of $\sigma_2$ obtained as the *composition* $\lambda \circ \sigma_2$ with $\lambda = \{x \leftarrow g(a,b)\}$; hence a unification algorithm only needs to compute $\sigma_2$.

There are pairs of terms which have more than one most general unifier (i.e. they are not an instance of any other unifier) under commutativity, but they always have at most *finitely many*. This is in contrast to the first situation (of free terms), where every pair of terms has at most *one most* general unifying substitution.

The problem becomes entirely different when we assume that the function denoted by f is associative:

(A) $f(x,f(y,z)) = f(f(x,y),z)$

In that case $\sigma_1$ is still a unifying substitution, but

$\sigma_3 = \{x \leftarrow f(g(a,b),\ g(a,b)),\ y \leftarrow a\}$

is also unifier:

$\sigma_3 t_1 = f(f(g(a,b),\ g(a,b)),\ g(a,b))$
$=_A f(g(a,b),\ f(g(a,b),\ g(a,b))) = \sigma_3 t_2$.

But $\sigma_4 = \{x \leftarrow f(g(a,b),\ f(g(a,b),\ g(a,b))),\ y \leftarrow a\}$ is again a unifying substitution and it is not difficult to see that there are *infinitely many* unifiers, all of which are most general.

Finally, if we assume that both axioms (A) and (C) hold for f then the situation changes yet again and for any pair of terms there are at most *finitely many* most general unifiers *under (A) and (C)*.

Many special unification algorithms for common theories have been developed in the past, since they have important applications in computer science [RSS79]. In particular, they are a crucial component of automatic theorem provers and of many pro-

grams developed in artificial intelligence. We shall now briefly review a formalism to express problems of the above kind.

## 1.2 UNIVERSAL UNIFICATION

Let $F_\Omega$ be the initial algebra of free terms whose elements are given a concrete representation by:

(i) X, the set of variables, is in $F_\Omega$

(ii) for $f \in \Omega$, the set of function symbols, and $t_1,\ldots,t_2 \in F_\Omega : f(t_1,\ldots,t_n) \in F_\Omega$ iff $\Omega f = n$.

[SS81c] provides a brief survey of the logical and algebraic background of unification theory. In this abbreviated version it is however sufficient for the reader to have an intuitive notion of first order terms.

Let $\bar{\sigma}: X \to F_\Omega$ be a mapping which is equal to the identity mapping almost everywhere. A *substitution* $\sigma: F_\Omega \to F_\Omega$ is the homomorphic extension of $\bar{\sigma}$ and is represented as a finite set of pairs

$$\sigma = \{x_1 \leftarrow t_1,\ldots,x_n \leftarrow t_n\}.$$

$\Sigma$ is the set of substitutions on $F_\Omega$. If t is a term and $\sigma$ a substitution, let V(t) denote the set of variables occurring in t and define:

$$DOM(\sigma) := \{x \in X: \sigma x \neq x\}$$
$$COD(\sigma) := \{\sigma x: x \in DOM(\sigma)\}$$
$$IV(\sigma) := \{x \in V(p): p \in COD(\sigma)\}$$

An equation $s = t$, $s,t \in F_\Omega$ is *unifiable* (is *solvable*) in the algebra $A$ iff there exists a homomorphism $\xi: F_\Omega \to A$ such that $\xi s = \xi t$ is valid in $A$. A set of equations T induces a congruence $=_T$ in $F_\Omega$ and $F_{\Omega/=_T}$ is the quotient algebra modulo $=_T$.

For simplicity of notation we assume we have a box of symbols, GENSYM, at our disposal, out of which we can take an unlimited number of "new" symbols:

$$X = X_0 \cup GENSYM.$$

We shall adopt the computational proviso that whenever GENSYM is referenced by $v \in GENSYM$ it is subsequently 'updated' by GENSYM' = GENSYM - {v} and $X_0' = X_0 \cup \{v\}$. Since $F_\Omega \simeq F_{\Omega'}$ we shall not always keep track of the '-s and just write $F_\Omega$.

A *renaming substitution* $\rho \in REN \subset \Sigma$ is defined by

(i) $COD(\rho) \subset X$ ;

(ii) $\forall x,y \in X$: if $x \neq y$ then $\rho x \neq \rho y$

For $s,t \in F_\Omega$ : $s \sim_\rho t$ if $\exists \rho \in REN$ such that $\rho s = \rho t$.

In order to formalize the accessing of a subterm in a term t let $\Pi(t)$ be the set of *positions in t* (i.e. subterm addresses) and we denote a *subterm of t at* $\pi$, by $t|\pi$. A *subterm replacement* of t *by* s *at* $\pi$ is $\hat{\rho}t$, with $\hat{\rho} = [\pi \leftarrow s]$. For definitions see [HT80]. For example if $t = f(g(a,b),c)$ then
$\Pi(t) = \{\Lambda,1,2,1.1,1.2\}$ and $t|1.2 = b$.
We denote replacements by $\hat{\sigma}$, $\hat{\rho}$, $\hat{\delta}$, etc. and substitutions by $\sigma$, $\rho$, $\delta$, etc.

A relation $\to \subset F_\Omega \times F_\Omega$ is *Noetherian* (terminating) if there are no infinite sequences: $s_1\ s_2\ s_3 \ldots$ . $\overset{+}{\to}$ is the transitive and $\overset{*}{\to}$ the reflexive and transitive closure of $\to$. A relation $\to$ is *confluent* if for every $r,s,t \in F_\Omega$ such that $r \overset{*}{\to} s$ and $r \overset{*}{\to} t$ there exists a $u \in F_\Omega$ such that $s \overset{*}{\to} u$ and $t \overset{*}{\to} u$. A confluent Noetherian relation is *canonical*.

We define two important relations $\to_T$ and $\rightarrowtail_T$ on $F_\Omega \times F_\Omega$ as follows:

A *rewrite system* $R = \{(l_1,r_1),\ldots(l_n,r_n)\}$ is any set of pairs $l_i,r_i \in F_\Omega$, usually writen as $l_i \to r_i$, such that $V(r_i) \subseteq V(l_i)$, $1 \leq i \leq n$. For two terms s and t, we say s *is paramodulated* to t, $s \rightarrowtail_R t$, if there exist $\pi \in \Pi(s)$, $\sigma \in \Sigma$ and $l_i \to r_i \in R$ such that $\sigma(s|\pi) = (\sigma \circ \rho)l_i$ for some $\rho \in REN$ and $COD(\rho) \subset GENSYM$ and $t = \hat{\sigma}s$, where $\hat{\sigma} = [\pi \leftarrow (\sigma \circ \rho)r_i]$.
For example, for $R = \{g(x,0) \to 0\}$ we have

$$s = f(g(a,y),b) \rightarrowtail f(0,b) = t$$

with $\pi = 1$ and $\sigma = \{x \leftarrow a,\ y \leftarrow 0\}$. The renaming substitution $\rho$ is used to avoid any conflict between variables.
For two terms s and t, we say s *is rewritten to* t, $s \to_R t$, if there exist $\pi \in \Pi(s)$, $\sigma \in \Sigma$ and $l_i \to r_i \in R$ such that $s|\pi = (\sigma \circ \rho)l_i$ for some $\rho \in REN$ and $COD(\rho) \in GENSYM$ and $t = \hat{\sigma}s$, where $\hat{\sigma} = [\pi \leftarrow (\sigma \circ \rho)r_i]$.
If s and t are the same as in the above example then: $s \not\to_R t$, since we are *not* allowed to substitute into s. Occasionally we keep track of the information by writing $s \overset{[\pi,i,\sigma]}{\rightarrowtail} t$ and $s \overset{[\pi,i,\sigma]}{\longrightarrow} t$. In addition, we use $[\Lambda,\lambda_0,\varepsilon]$ to express the reflexivity of $\rightarrowtail$ resp. $\to$ (i.e. $\lambda_0$ means: *no* 'rule' is applied). Thus $s \overset{[\Lambda,\lambda_0,\varepsilon]}{\longrightarrow} s$ for all $s \in F_\Omega$.

The notation and definitions of term rewrite systems are consistent with [HT80]. Suppose for an equational theory T there is a rewrite system $R_T$ such that for $s,t \in F_\Omega$:

$s =_T t$ iff $\exists p \in F_\Omega$ such that $s \overset{*}{\rightarrow}_{R_T} p$ and $t \overset{*}{\rightarrow}_{R_T} p$.
In that case we say T is *embedded into* $R_T$ and write

$$T \hookrightarrow R_T .$$

For an equational theory T there are techniques to obtain a system $R_T$ such that $T \hookrightarrow R_T$; moreover for many theories of practical interest it is possible to obtain a rewrite system $R_T$ such that $\rightarrow_{R_T}$ is canonical [KB70], [HT80], [PS81]. Canonical relations $\rightarrow$ are an important basis for *computations* in *equational logics*, since they define a unique *normal form* $\|t\|$ for any $t \in F_\Omega$, given by $t \overset{*}{\rightarrow} \|t\|$ and $\nexists\, s \in F_\Omega$ such that $\|t\| \rightarrow s$. Hence

(i) $s =_T t$ iff $\|s\| = \|t\|$ .

An equational theory T is *decidable* iff $s =_T t$ is decidable for any $s,t \in F_\Omega$. Let $\mathcal{T}_\Delta$ denote the family of decidable equational theories.
A *T-unification problem* $\langle s = t \rangle_T$ consists of a pair of terms $s,t \in F_\Omega$ and a theory $T \in \mathcal{T}_\Delta$.
A substitution $\sigma \in \Sigma$ is a *T-unifier* for $\langle s = t \rangle_T$ iff $\sigma s =_T \sigma t$. The subset of $\Sigma$ which unifies $\langle s = t \rangle_T$ is $U\Sigma_T(s,t)$, the *set of unifiers* (for s and t) *under T*.
It is easy to see that $U\Sigma_T$ is recursively enumerable (r.e.) for any s and t: Since $F_\Omega$ is r.e. so is $\Sigma$. Now for any $\delta \in \Sigma$. check if $\delta s =_T \delta t$ (which is decidable since $T \in \mathcal{T}_\Delta$) then $\delta \in U\Sigma_T(s,t)$ otherwise $\delta \notin U\Sigma_T(s,t)$.

The composition of substitutions is defined by the usual composition of mappings: $(\sigma\circ\tau)t = \sigma(\tau t)$. If $W \subseteq X$, then T-equality is extended to substitutions by

$$\sigma =_T \tau\ [W] \quad \text{iff} \quad \forall x \in W \quad \sigma x =_T \tau x \ ,$$

$\sigma$ *and* $\tau$ *are T-equal in W*. We say $\sigma$ is an *instance* of $\tau$ and $\tau$ is *more general* than $\sigma$, in symbols

$$\tau \leq_T \sigma[W] \text{ iff } \exists \lambda \in \Sigma \quad \tau =_T \lambda\circ\sigma[W] \text{ for some } W \subset X.$$

If $\sigma \leq_T \tau[W]$ and $\tau \leq_T \sigma[W]$ then $\sigma \approx_T \tau[W]$, $\sigma$ and $\tau$ are *T-equivalent in W*.

For $\Sigma_1, \Sigma_2 \subseteq \Sigma$ we define
$\Sigma_1\circ\Sigma_2 = \{\sigma_1\circ\sigma_2 : \sigma_1 \in \Sigma_1,\ \sigma_2 \in \Sigma\}$. $\Sigma_1 \subseteq_T \Sigma_2$ [W] iff $\forall \sigma_1 \in \Sigma_1\ \exists\, \sigma_2 \in \Sigma_2$ such that $\sigma_1 =_T \sigma_2$ [W],
$\Sigma_1 =_T \Sigma_2$ [W] iff $\Sigma_1 \subseteq_T \Sigma_2$ [W] and $\Sigma_2 \subseteq_T \Sigma_1$ [W].

Universal unification is concerned with two fundamental problems:

PROBLEM ONE (*Decidability Problem*)
*For a given equational theory* $T \in \mathcal{T}_\Delta$, *is it decidable for any s and t whether s and t are unifiable under T?*

We are interested in classes of theories such that "s and t are unifiable under T" is decidable for every T in that class.
A unifier $\sigma$ for $\langle s = t \rangle_T$ is called *a most general unifier* (mgu) if for any unifier $\delta \in U\Sigma_T(s,t)$: $\delta \leq \sigma$ [W], where $V(s,t) \subseteq W$. Since in general a single most general unifier does not exist for $\langle s = t \rangle_T$, we define $\mu U\Sigma_T(s,t)$, the *set of most general unifiers*, as:

(i) $\mu U\Sigma \subseteq U\Sigma$ (correctness)
(ii) $U\Sigma =_T \Sigma\circ\mu U\Sigma$ [W] (completeness)
(iii) $\sigma_i \not\leq_T \sigma_k$ for $i \neq k$; $\sigma_i, \sigma_k \in \mu U\Sigma$ (minimality).

From condition (ii) it follows in particular that $U\Sigma = \Sigma\circ U\Sigma$, i.e. $U\Sigma$ is a *left ideal* in the semigroup $(\Sigma,\circ)$ and $U\Sigma$ is *generated by* $\mu U\Sigma$. For practical applications these conditions are sometime too general and there are additional technical requirements on DOM($\sigma$), COD($\sigma$) and IV($\sigma$) for $\sigma \in \mu U\Sigma$, which we shall state when the need arises.

PROBLEM TWO (*Enumeration Problem*)
*For a given equational theory* $T \in \mathcal{T}_\Delta$, *is* $\mu U\Sigma_T(s,t)$ *recursively enumerable for any* $s,t \in F_\Omega$?

That is, we are interested in algorithms which generate all mgus for a given problem $\langle s = t \rangle_T$.
TABLE 1 summarizes the major results that have been obtained for special theories, which consist of combinations of the following equations:

| | | |
|---|---|---|
| A | (associativity) | $f(f(x,y),z) = f(x,f(y,z))$ |
| C | (commutativity) | $f(x,y) = f(y,x)$ |
| D | (distributivity) | $f(x,g(y,z)) = g(f(x,y),f(x,z))$<br>$f(g(x,y),z) = g(f(x,z),f(y,z))$ |
| H,E | (homomorphism, endomorphism) | $\varphi(x\circ y) = \varphi(x)*\varphi(y)$ |
| I | (idempotence) | $f(x,x) = x$ |

Moreover we use the following abbreviations:

QG: Quasi-Groups
AG: Abelian-Groups
H10: Hilbert's 10th Problem
Sot: Second order terms
Hot: Higher order terms (i.e. $\geq 3^{rd}$ order)

The column under $\mathbf{A}_T$ indicates whether or not a type conformal algorithm has been presented in the literature. The type of a theory and type conformality are defined below.

| Theory T | Type of T | $U\Sigma_T \neq \emptyset$ decidable | $\mu U\Sigma_T$ recursive | $\mathbf{A}_T$ |
|---|---|---|---|---|
| ∅ | 1 | Yes | Yes | Yes |
| A | ∞ | Yes | Yes | No |
| C | ω | Yes | Yes | Yes |
| I | ω | Yes | Yes | Yes |
| A+C | ω | Yes | Yes | Yes |
| A+I | ω | Yes | Yes | Yes |
| C+I | ω | Yes | Yes | Yes |
| A+C+I | ω | Yes | Yes | Yes |
| D | ∞ | ? | Yes | No |
| D+A | ∞ | No | Yes | No |
| D+C | ∞ | ? | Yes | No |
| D+A+C | ∞ | No | Yes | No |
| D+A+I | ω | Yes | Yes | Yes |
| H,E | 1 | Yes | Yes | Yes |
| H+A | ∞ | Yes | Yes | No |
| H+A+C | ω | Yes | Yes | Yes |
| E+A+C | ∞ | ? | ? | No |
| QG | ω | Yes | Yes | Yes |
| AG | ω | Yes | Yes | Yes |
| H10 | ∞ | No | No | No |
| Sot, T = ∅ | 0 | No | - | - |
| Hot, T = ∅ | 0 | No | - | - |

TABLE 1 (The references to the work presented in this table are collected in the full paper [SS81])

Except for Hilbert's tenth problem, we have not included the classical work on equation solving in 'concrete' structures such as rings and fields, which is well known. The relationship of universal unification to these classical results is similar to that of universal algebra to classical algebra.

The central notion $\mu U\Sigma_T$ induces a hierarchy of classes of equational theories [SS81b].

(i) A theory T is *unitary* if every $\langle s = t\rangle_T$ has at most one mgu. The class of such theories is $\mathcal{J}_1$ (*type one*).

(ii) A theory T is *finitary* if it is not unitary and for every $\langle s = t\rangle_T$ $\mu U\Sigma_T$ always exists and is finite. The class of such theories is $\mathcal{J}_\omega$ (*type* ω).

(iii) A theory T is *infinitary* if $\mu U\Sigma_T$ always exists and there exists $\langle s = t\rangle_T$ such that $\mu U\Sigma_T$ is infinite. The class of such theories is $\mathcal{J}_\infty$ (*type* ∞).

(iv) A theory T is of *type zero* if it is not one of the above classes. The class of these theories is $\mathcal{J}_0$.

There are several examples for unitary, finitary and infinitary theories in the above table. An example of a type zero problem is the unification of higher order terms (which, however, usually assumes a syntax different from $F_\Omega$), since there are infinitely many ascending chains of unifiers under $\leq_T$ [G066]. We define a partial order $\leq_T$ on terms by: $s \leq_T t$ iff $\exists\delta\in\Sigma$ satisfying $s =_T \delta t$. A matching problem $\langle s \geq t\rangle_T$ consists of a pair of terms and a theory $T\in\mathcal{J}_\Delta$. A substitution $\nu\in\Sigma$ is a T-*matcher* (or one-way unifier) if $\nu s =_T t$. If s and t are matchable we shall write $s \leq_T t$ (resp. $s \geq_T t$). The minimal set of matchers is $\mu M\Sigma_T$, which induces the corresponding hierarchy $\mathcal{M}_0, \mathcal{M}_\infty, \mathcal{M}_\omega, \mathcal{M}_1$.

A *unification algorithm* $\mathbf{A}_T$ (a *matching algorithm* $\mathbf{M}_T$) *for a theory T* is an algorithm which takes two terms s and t as input and generates a set $\Psi_T \subseteq U\Sigma_T$ ($\subseteq M\Sigma_T$) for $\langle s = t\rangle_T$ (for $\langle s \geq t\rangle_T$). A *minimal* algorithm $\mu\mathbf{A}_T$ ($\mu\mathbf{M}_T$) is an algorithm which generates $\mu U\Sigma_T$ ($\mu M\Sigma_T$).

For many practical applications this requirement is not strong enough, since it does not imply that the algorithm terminates for theories $T\in\mathcal{J}_1\cup\mathcal{J}_\omega$. On the other hand, for $T\in\mathcal{J}_\omega$ it is sometimes too rigid, since an algorithm which generates a finite superset of $\mu U\Sigma_T$ may be far more efficient than the algorithm $\mu\mathbf{A}_T$ and for that reason preferable.

An algorithm $\mathbf{A}_T$ is *type conformal* iff:

(i) $\mathbf{A}_T$ generates a set $\Psi_T$ with $U\Sigma_T \supseteq \Psi_T \supseteq \mu U\Sigma_T$;

(ii) $\mathbf{A}_T$ terminates and $\Psi_T$ is finite if $\mu U\Sigma_T$ is; and

(iii) if $\mu U\Sigma_T$ is infinite then $\Psi_T \approx_T \mu U\Sigma_T$.

Similarly: algorithm $\mathbf{M}_T$ is type conformal iff (i)-(iii) hold with U replaced by M.

A *universal unification algorithm (a universal matching algorithm) for a class of theories* $\mathcal{J}_U \subseteq \mathcal{J}_\Delta$ is an algorithm which takes as input a pair of terms (s,t) *and a theory* $T\in\mathcal{J}_U$ and generates $\Psi_T \supseteq \mu U\Sigma_T$ ($\supseteq \mu M\Sigma_T$) for $\langle s = t\rangle_T$ (for $\langle s \geq t\rangle_T$).

Since $U\Sigma_T$ is trivially r.e. for any $T\in\mathcal{J}_\Delta$, there is the important requirement that a universal unification algorithm is either minimal or at least type conformal. The known universal algorithms [FA79], [LB79], [HO80] are neither minimal nor type conformal; and even some special purpose unification algorithms, although they are minimal, are not type conformal either.

## 2. THE CLASS OF ACFM THEORIES

An equational theory T is *admissible* (A) if the matching problem for T is decidable. This restriction pertains to all unification problems of practical interest, as otherwise $\sigma \leq_T \delta$ would be undecidable.

For the purpose of this paper we further restrict the admissible theories by requiring that T possess a *confluent reduction system*, i.e. a *canonical* (C) system.

The final restriction is the requirement that T constitutes a *finitary matching problem* (FM), i.e. that $T \in \mathcal{M}_\omega$. The last restriction is less rigid than it appears: e.g. all special unification problems in TABLE 1 are in this class. On the other hand, there exist very simple theories which do not fall into this class: Let $T = \{g(f(x)) = g(x)\}$; then $\langle g(y) \geq g(a) \rangle_T$ has an infinite matching set:

$$M\Sigma_T = \{y \leftarrow a,\ y \leftarrow f(a),\ y \leftarrow f(f(a)),\ \ldots\}\ .$$

The class of equational theories T with T admissible and $T \in \mathcal{M}_\omega$ and $T \hookrightarrow R_T$ for a canonical $R_T$ is the *class of ACFM theories*.

An equational *theory T is regular* iff for every $s = t \in T$ $V(s) = V(t)$; similarly, a *rewrite system* R *is regular* iff for every $l \to r \in R$ $V(l) = V(r)$. For a regular theory $T = \{l_i = r_i : 1 \leq i \leq n\}$ the corresponding regular rewrite system can immediately be obtained as $R_T = \{l_i \to r_i,\ r_i \to l_i : 1 \leq i \leq n\}$, and $T \hookrightarrow R_T$ is obvious.

For regular theories we have an important property concerning the *length of terms*, in symbols $|t|$, defined as the number of symbols in t.

To refer to the (least possible) *length of rewrites*, let $|s \xrightarrow{*} t| \in \mathbb{N} \cup \{0\}$ denote the *smallest* n such that $s = s_0 \to s_1 \to s_2 \ldots \to s_n = t$

*2. Lemma 1.* Let R be a regular rewrite system and $m_R = \max\{|l|, |r| : l \to r \in R\}$. Then for any $s,t \in F_\Omega$ such that $s \xrightarrow{*}_R t$ and $|s \xrightarrow{*}_R t| = n \geq 1$:

$$|s| \leq m_R^n \cdot |t|\ .$$

## 3. LOCALIZING THE TEST FOR $T \in \mathcal{M}_\omega$

In order to test an equational theory T for the ACFM property let us assume it is admissible. We also assume that we have proved $T \hookrightarrow R_T$, where $R_T$ is canonical, e.g. by using the techniques of [HT80]. Then there remains the test for $T \in \mathcal{M}_\omega$ which can be difficult, since it has to be shown that

$\forall s,t \in F_\Omega$: $|\mu M\Sigma(s,t)| \in \mathbb{N}$ .

The test can be localized in the following sense:

*3. Theorem 6:* Let T be an admissible theory with a regular canonical reduction system $R_T$ such that $T \hookrightarrow R_T$. If $T \in \mathcal{M}_\omega$ then there exists a ground substitution $\tau$, $(l_i, r_i)$ and $(l_k, r_k) \in R_T$ such that

$$\langle l_i \geq \tau r_k \rangle_T$$

has an infinite matching set.

Hence, if there does not exist a pair $\langle l \geq \tau r \rangle_T$ with an infinite matching set, then $T \in \mathcal{M}_\omega$. It is an open problem to find a local (and finite) test set for theories which are *not* regular.

## 4. UNIVERSAL UNIFICATION ALGORITHMS IN REGULAR ACFM THEORIES

A universal matching algorithm and a universal unification algorithm are presented in this section. Both algorithms take a regular ACFM theory T and a pair of terms as input and output the sets $\mu M\Sigma_T$ and $\mu U\Sigma_T$ respectively, hence the algorithms are universal in the sense that a universal Turing machine is 'universal'.

Throughout this section we slightly change the notion of $\mu U\Sigma_T(s,t)$ by adding the following technical requirement:

(iv) $\forall \sigma \in \mu U\Sigma_T(s,t)$ $\quad IV(\sigma) \cap W = \emptyset$, where $V(s,t) \subseteq W \subset X$.

In case (iv) is satisfied we shall write $\mu U\Sigma_T \vdash W$ and say '$\mu U\Sigma_T$ *away from* W' [PL72], [HO80].

Let R be a canonical system. A substitution $\sigma \in \Sigma$ is called *normalized* if for all $t \in COD(\sigma)$ t is in normal form (i.e. $t = \|t\|_R$). Let $\bar{\Sigma}_R \subset \Sigma$ be the set of normalized substitutions and analogously $U\bar{\Sigma}_R$ ($M\bar{\Sigma}_R$) the set of normalized unifiers (matchers). We denote the normalization of substitutions as for terms, i.e. if $\sigma \in \Sigma$ then $\|\sigma\| \in \bar{\Sigma}$.

### 4.1 *The Paramodulation Tree*

Let $\rightarrowtail$ be the paramodulation relation on $F_\Omega \times F_\Omega$ as defined in section 1.2 with the additional proviso that we never paramodulate into variables, i.e. if $s \overset{[\pi,i,\sigma]}{\rightarrowtail} t$ then $s|_\pi \notin X$. The set $(t \rightarrowtail) = \{[s]_{\sim_\rho} : t \overset{[\pi,i,\sigma]}{\rightarrowtail} s\}$ is called the *follow set* of $\rightarrowtail$, factored by the equivalence relation $\sim_\rho$. Note that $(t \rightarrowtail)$ is finite. For ease of notation we just write s for some representative of $[s]_{\sim_\rho}$. For a given term t we define

the *labeled paramodulation tree* $P_t$ as:

(i) t (the *root*) is a node in $P_t$,

(ii) if $\tilde{s}$ is a node in $P_t$ and $s \in (\tilde{s} \rightarrowtail)$ then s (the *successor*) is a node in $P_t$;

(iii) the edge $(\tilde{s},s)$, where $\tilde{s} \xrightarrow[[\pi,i,\sigma]]{} s$, is labeled with the triple $[\pi,i,\sigma]$.

The composition '$\circ$' of labels is defined as $[\pi_0,i_0,\sigma_0]\circ[\pi_1,i_1,\sigma_1] = [(\pi_0,\pi_1),(i_0,i_1),\sigma_0\circ\sigma_1]$.
In analogy to the notion of a derivation word in formal language theory we define the *paramodulation word* $L^*(P_t,s)$ for $s \in P_t$ by:

(i) $L^*(t,t) = (\lambda_0,\Lambda,\varepsilon)$

(ii) $L^*(t,s) = L^*(t,\tilde{s}) \circ (\pi,i,\sigma)$ where $\tilde{s} \xrightarrow[[\pi,i,\sigma]]{} s$.

It is convenient to have selectors $\lambda,\mathbb{I},\theta$ on $L^*$ defined by: $L^*(t,s) = (\mathbb{I}(t,s), \lambda(t,s), \theta(t,s))$.

The relationship between T-unification and the paramodulation tree is shown up by the following construction: for a problem $\langle s \doteq t\rangle_T$ define a term $H = h(s,t)$, where h is a 'new' symbol with $\Omega h = 2$, and consider $P_H$:

*Theorem:* (i) For every $h(\tilde{s},\tilde{t}) \in P_H$: if $\sigma$ is mgu of $\tilde{s}$ and $\tilde{t}$ then $\sigma\circ\theta \in U\Sigma_T(s,t)$ where $\theta = \theta(H,h(\tilde{s},\tilde{t}))$

(ii) For every $\delta \in U\Sigma_T(s,t)$: there exists $h(\tilde{s},\tilde{t}) \in P_H$ and $\sigma$ mgu of $\tilde{s}$ and $\tilde{t}$ such that $\delta \leq_T \sigma\circ\theta$ where $\theta = \theta(H,h(\tilde{s},\tilde{t}))$

The first part of the theorem states that every substitution which is obtained from a Robinson unifier $\sigma$ for $\tilde{s}$ and $\tilde{t}$ and the labelword component $\sigma\circ\theta$ is a *correct* unifier for $\langle s \doteq t\rangle_T$. The second part of the theorem ensures *completeness:* for every unifier $\delta \in U\Sigma(s,t)$ there is a node in $P_H$ such that $\delta$ is an instance of the substitution obtained at that node.

The proof of the theorem is immediate from Birkhoff's theorem. It is also an immediate consequence of the correctness and completeness of paramodulation [WR73].
In the following we are concerned with refinements of the paramodulation tree incorporating those of [HU80], [FA79], [LB79], [SL74].

Let s be a term and $P_s$ the corresponding paramodulation tree. Then we define a refinement of $P_s$, the *normalized paramodulation tree* $\bar{P}_s$, by: The generation of nodes in $P_s$ is stopped in case one of the following two criteria is true: A node $t \in P_s$ is a *leaf node* iff

C1: $(t \rightarrowtail) = \emptyset$

C2: $\theta(s,t) \notin \bar{\Sigma}$

The tree thus obtained is $\bar{P}_s$.

### 4.2 *Universal Matching Algorithms in Regular ACFM Theories*

For a given matching problem $\langle s \geq t\rangle_T$ assume that T is a regular ACFM theory and let $\bar{P}_s$ be the normalized paramodulation tree for s. We extend the notion of leaf nodes by the following criteria: $\tilde{s} \in P_s$ is also a *leaf node* if:

C3: $M\Sigma(\tilde{s},t) = \emptyset$, i.e. $t \not\leq_T s$, which is decidable since T is admissible.

We shall write $\mu\bar{P}_s(t)$, the *minimal paramodulation tree for a relative to t*, if C3 is applied to $\bar{P}_s$ and $\langle s \geq t\rangle_T$ is the given matching problem.

The set of matchers obtained from $\mu\bar{P}_s(t)$ is the most general set of matchers for s and t up to renaming:

*4. Theorem 8:* Let T be a regular ACFM theory and $\langle s \geq t\rangle_T$. For $M = \{(\sigma_i\circ\theta_i)|_W : W = V(s,t), \theta_i = \theta(s,\tilde{s}), \tilde{s} \in \mu\bar{P}_s(t)$ and $\{\sigma_i\} = \mu U\Sigma_\emptyset(\tilde{s},t)\}$
$M = \mu M\Sigma_T(s,t) \vdash W$.

Since T is an ACFM theory $\mu U\Sigma$ is finite, and hence $\mu\bar{P}_s(t)$ is finite except for the fact that there may be an infinite path in $\mu\bar{P}_s(t)$ which is not terminated by C3 (i.e. there are infinitely many T-equivalent matchers for $\langle s \geq t\rangle_T$ along the path).

*4. Theorem 9:* Let T be a regular ACFM theory. For any $s,t \in F_\Omega$
$\mu\bar{P}_s(t)$ is finite.

We have limited ourselves to regular ACFM theories, since they are applicable in most cases of interest: An equational theory T is often separated into subtheories $T = E \cup R$ such that R allows for a canonical reduction and for each $s = t \in E$ $Var(s) = Var(t)$ holds, and interest is in (finite) unification algorithms for E; see also [HT80], [PS81]. All theories in TABLE 1 of section 1 are regular.

For nonregular ACFM theories the situation is not as simple and we suspect that a *terminating* universal matching algorithm does not exist. For many special cases a terminating matching algorithm is of course possible even in the nonregular case.

### 4.3 *Universal Unification Algorithms in Regular ACFM Theories*

The universal unification algorithm of this section

is based on the matching algorithm of 4.2 and on the solvability of equations in the $(\Sigma_\Omega,T)$-algebra, where $\Omega = \{(\circ,2)\}$. For given substitutions $\delta,\tau\in\Sigma$ we should like to know if there is a third substitution $\sigma\in\Sigma$ such that $\delta =_T \sigma\circ\tau$. That is, we are interested in the solvability of equations in $(\Sigma_\Omega,T)$:

(i) $\delta =_T \tilde{V}\circ\tau$ , where $\delta,\tau\in\Sigma$ and $\tilde{V}$ is a variable ranging over substitution.

We say (i) has a most general set of solutions $\mu M\Sigma_T(\delta,\tau)$ iff for every $\gamma$ such that $\delta =_T \gamma\circ\tau$ there exists $\sigma\in\mu M\Sigma_T(\delta,\tau)$ such that $\sigma \geq_T \gamma$.

*4. Lemma 10:* Let T be admissible and $T\in\mathcal{T}_\omega$. For $\delta =_T \tilde{V}\circ\tau$, (i) $\mu M\Sigma_T(\delta,\tau)$ is finite; (ii) $\mu M\Sigma_T(\delta,\tau) = \emptyset$ is decidable.

4. Lemma 10. is wellknown, but with an unnecessarily complicated proof, in the case $T = \emptyset$ [VA75], since it is used to advantage in the connection graph proof procedure [KO75]. We are now ready to state the main result of this paper.

The enumeration of $\mu U\Sigma_T(s,t)$ is based on the normalized paramodulation tree $\hat{P}$ as defined in section 4.1. The check for minimality is based on the (one-sided) $\mu\hat{P}$-tree.

*4. Theorem 12:* Let T be a regular ACFM theory. For any $s,t\in F_\Omega$ and $W = V(s,t)$, $\mu U\Sigma_T(s,t) \restriction W$ is recursively enumerable.

The proof of 4 Theorem 12 is constructive and hence specifies a universal unification algorithm. The proof is based on lemmas and technical material, which we had to omit for lack of space, see [SS81c]. Note that the universal unification algorithm based on 4. theorem 12 may not terminate.

If a decision procedure is known for T-unifiability the paramodulation tree can be pruned considerably by an application of this procedure to each node: in case of nonunifiability the node is a leaf node. However even under this additional termination criterion the paramodulation tree may grow indefinitely even if $T\in\mathcal{T}_\omega$, which is to be expected in general, since we do not assume $\mu U\Sigma$ to be recursive. For certain classes of theories termination is of course possible.

## 5. CONCLUSION

We have presented theorems for a universal matching algorithm and a universal unification algorithm for regular ACFM theories. Both algorithms based on these theorems generate a complete minimal set of unifiers (of matchers) and termination of the matching algorithm is shown.

For nonregular ACFM theories the matching algorithm may not terminate and we suspect that uniform termination is unattainable.

Since most known special purpose unification algorithms are based on regular ACFM theories and since this class allows for a comparatively simple universal unification algorithm, we believe this class may come to play a prominent role (similar to certain classes, like SLR(I), in formal language theory).

On the other hand, the 'abstract' universal algorithm does not represent a practical solution for a given theory T, because of its gross inefficiency.

This contribution is of *theoretical* relevance in that it exhibits a class of theories which possesses an r.e. set $uUI_T$. This result can be applied in *practice* for the design of an actual algorithm: So far the design of a special purpose unification algorithm was more of an art than a science since for a given theory there was no indication whatsoever of how the algorithm might work. In fact the algorithms in TABLE 1 of section 1 all operate on entirely different principles.

*The next 700 Unification Algorithms*

On the basis of the "abstract" universal unification algorithm it is possible to find a concrete algorithm much more easily by first isolating the crucial parts in the abstract algorithm and then designing a practical and efficient solution. The universal algorithm has been successfully applied to a special case [RS78], yielding a minimal algorithm [SS81a], which in addition is much simpler than the one previously known. A collection of canonical theories [HU80a] is a valuable source for this purpose and has already been used to find the first unification algorithm for Abelian group theory and quasi group theory [LA79], [HU80].

## REFERENCES

We apologize to the reader that for space limitations we had to append the references to the paper by K. Blasius et. al. "The Markgraf Karl Refutation Procedure", this volume.