

Unification of Multisets with Multiple Labelled Multiset Variables ^{*}

Zan Naeem¹ and Giselle Reis¹

Carnegie Mellon University, Doha, Qatar
znaeem@andrew.cmu.edu, giselle@cmu.edu

Abstract

We look into the problem of unifying multisets containing (first-order) terms and *multiple* multiset variables. The variables are labelled, meaning that a unifier that places a term in a multiset variable M_i is different from another that places a term in a multiset variable M_j , for $i \neq j$. We describe a sound, complete, and terminating algorithm for computing the set of all possible unifiers, and analyse its complexity. We also prove an input pre-processing step that avoids the computation of less general unifiers.

1 Introduction

Multiset is an important data-structure that is used to specify various object systems. Our motivation stems mostly from proof theory, where logical entailment is encoded as sequents $\Gamma \vdash \Delta$, where both Γ and Δ are typically considered as multisets. When reasoning about such objects, one might need to use an implementation of sets/multisets based on lists, since most reasoning tools (i.e., logical frameworks, proof assistants, and logic programming languages) do not have built-in support for these data structures [2, 3, 6, 7]. Adding this kind of support requires, among other things, a unification algorithm.

Multiset unification was studied in [1, 4], where the authors propose solutions for the problem of unifying multisets with at most one multiset variable. We extend those results for multisets with multiple multiset variables. In our setting, each multiset variable is *labelled*, meaning that assigning a term to either a multiset M_i or M_j , where $i \neq j$, should be considered different solutions. We describe a terminating algorithm and analyse its complexity. Moreover, we prove that a simple modification of our algorithm avoids the computation of less general unifiers.

The need for labelled multiset variables emerged when reasoning with multiplicative rules in sequent calculi, such as:

$$\frac{\Gamma_1 \vdash A \quad \Gamma_2, B \vdash C}{\Gamma_1, \Gamma_2, A \rightarrow B \vdash C} \rightarrow_l$$

To apply this rule to, e.g., the sequent $\Gamma, D, A \rightarrow B \vdash C$, where A, B, C , and D are formulas, we need to unify its antecedent $\Gamma, D, A \rightarrow B$ with $\Gamma_1, \Gamma_2, A \rightarrow B$. Assigning formula D to Γ_1 or Γ_2 should be considered two different solutions, since they result in two different applications of the rule.

1.1 Preliminaries

A *multiset \mathcal{M} with multiple labelled multiset variables* is denoted by $\{t_1, \dots, t_n | M_1, \dots, M_k\}$. Each t_i is a *term* ranging over a first-order language $\mathcal{L} = \langle \Sigma, \mathcal{V} \rangle$, where Σ is a set of constants

^{*}This publication was made possible by the support of Qatar Foundation through Carnegie Mellon University in Qatar's Seed Research program. The statements made herein are solely the responsibility of the authors.

and function symbols, and \mathcal{V} is a denumerable set of *term variables*. Each M_i is a *multiset variable* ranging over a denumerable set \mathcal{V}_M of multiset variables. When n and k are not relevant, we abbreviate t_1, \dots, t_n as \bar{t} and M_1, \dots, M_k as \bar{M} . When $k = 0$, we write the multiset as $\{\{t_1, \dots, t_n\}\}$. Henceforth, we refer to multisets with multiple labelled multiset variables as *mmsets* for brevity.

Definition 1 (Substitution). *An mmset substitution σ is a finite mapping of term variables to terms, and of multiset variables to mmsets. The application of a substitution σ to an mmset $\{\{t_1, \dots, t_n | M_1, \dots, M_k\}\}$ is defined as:*

$$\{\{t_1, \dots, t_n | M_1, \dots, M_k\}\}\sigma = \{\{t_1\sigma, \dots, t_n\sigma\}\} \uplus M_1\sigma \uplus \dots \uplus M_k\sigma$$

where each $t_i\sigma$ is the usual first-order substitution and \uplus is left-associative and defined as:

$$\{\{t_1, \dots, t_n | M_1, \dots, M_k\}\} \uplus \{\{s_1, \dots, s_m | N_1, \dots, N_l\}\} = \{\{t_1, \dots, t_n, s_1, \dots, s_m | M_1, \dots, M_k, N_1, \dots, N_l\}\}$$

Definition 2 (Equality). *Mmsets $\mathcal{M} = \{\{t_1, \dots, t_n | M_1, \dots, M_k\}\}$ and $\mathcal{N} = \{\{s_1, \dots, s_m | N_1, \dots, N_l\}\}$ are considered equal modulo a constraint theory \mathcal{T} , written $\mathcal{M} =_{\mathcal{T}} \mathcal{N}$ iff: $n = m$ and t_1, \dots, t_n is a permutation of s_1, \dots, s_m ; and $\mathcal{T} \vdash M_1 \cup \dots \cup M_k \equiv N_1 \cup \dots \cup N_l$.*

2 Mmsets Unification

The *mmset unification problem* of mmsets \mathcal{M}_1 and \mathcal{M}_2 consists of finding a substitution σ and constraint theory \mathcal{T}_σ such that $\mathcal{M}_1\sigma =_{\mathcal{T}_\sigma} \mathcal{M}_2\sigma$. The theory \mathcal{T}_σ consists of an equality over unions of multiset variables, and it is computed *a posteriori* for each unifier σ .

2.1 Algorithm

In what follows, we use σ to denote a single substitution, Σ s to denote sets of substitutions, \times for the Cartesian product of two sets (or lists), and \setminus for multiset difference. The pseudo code for all algorithms are listed in Appendix A, and an implementation in SML can be found at <https://github.com/meta-logic/mmset-unif>.

The main function for mmset unification is implemented by Algorithm 1. In the most general case (lines 11 to 17), the unifiers of $\{\{\bar{t} | \bar{M}\}\}$ and $\{\{\bar{s} | \bar{N}\}\}$ are computed by choosing a subset of terms (of the same size) from \bar{t} and \bar{s} to be unified, and distributing the rest among the multiset variables \bar{M} and \bar{N} . The number of terms chosen to be unified can vary from 0 to the minimum length of \bar{t} and \bar{s} . Two other cases are considered separately for efficiency purposes. The first one is when there are no multiset variables (line 1). Here a unification is only possible if $|\bar{t}| = |\bar{s}|$. The second case is when one of the mmsets does not have multiset variables (lines 4 and 7). If \bar{M} is empty, then all terms in \bar{s} must be unified with a term from \bar{t} . The remaining terms in \bar{t} can be allocated in \bar{N} .

Function `unify_c` (Algorithm 2) chooses c terms from the multisets \bar{t} and \bar{s} to be unified, and distributes the rest of the terms among the multiset variables. The function `choose(F, c)` returns a set of tuples (F_c, F_r) , where F_c is the multiset with the chosen c elements (thus $|F_c| = c$), and $F_r = F \setminus F_c$. Unifiers for the chosen terms are computed by `unify_terms` and stored in Σ_t . Substitutions for each context variable, containing the remaining terms, are computed by `unify_distribute` and stored in Σ_M . The final set of unifiers consists of the composition $\sigma_M \sigma_t$ for each $(\sigma_M, \sigma_t) \in \Sigma_M \times \Sigma_t$. Note that, since the image of σ_M may contain terms, the composition needs to be in this order.

Function `unify_terms` (Algorithm 3) finds all unifiers of two multisets of terms (without multiset variables) of equal length. This is done by testing all possible pairings of the two multisets, obtained by pairing some order of the first multiset with all possible permutations of the second one. For each pairing, the function `unify_lists` computes the most general unifier.

Function `unify_distribute` (Algorithm 4) computes unifiers for the mmsets $\{\bar{t}|\overline{M}\}$ and $\{\bar{s}|\overline{N}\}$ considering that all \bar{t} occurs in \overline{N} and all \bar{s} occurs in \overline{M} . Let Σ_N denote the substitutions that distribute \bar{t} into \overline{N} , and Σ_M the substitutions that distribute \bar{s} into \overline{M} . The resulting substitutions are $\sigma_M \cup \sigma_N$ for each $(\sigma_M, \sigma_N) \in \Sigma_M \times \Sigma_N$. A simple union can be used in the case, as the image of σ_M is disjoint from the domain of σ_N (see below).

Function `distribute` (Algorithm 5) is used by `unify_distribute` and that is where the aforementioned Σ_N and Σ_M are computed. It computes substitutions for the multiset variables N_1, \dots, N_l such that all terms t_1, \dots, t_n occur in one of the multisets. This is done by calculating all ordered l -partitions of the multiset $\{t_1, \dots, t_n\}$. For example, the ordered 2-partitions of the multiset $\{a, b, c\}$ are:

$$\begin{array}{cccc} \{\{a, b, c\}, \{\}\} & \{\{a, b\}, \{c\}\} & \{\{a\}, \{b, c\}\} & \{\{a, c\}, \{b\}\} \\ \{\{\}, \{a, b, c\}\} & \{\{c\}, \{a, b\}\} & \{\{b, c\}, \{a\}\} & \{\{b\}, \{a, c\}\} \end{array}$$

Each computed substitution corresponds to an l -ordered partition. If there are no terms ($n = 0$), then there is only the trivial partition of l empty multisets. In this case, the algorithm returns a list with one substitution, which maps every multiset variable N_i to the mmset with no terms. If there are no multiset variables ($l = 0$), then there are no partitions, and thus no possible substitutions. The exceptional case is when there are no terms nor multiset variables ($l = n = 0$). In this case, the solution is the set containing only the empty substitution ($\{\}$).

The last parameter of `distribute` indicates whether the multiset variables should contain *exactly* the terms t_1, \dots, t_n . If set to true, then there is no more space for other terms, and N_i is mapped to an mmset with the appropriate terms and no multiset variables. Otherwise it is mapped to an mmset with a set of terms and a *fresh* multiset variable. If there are no terms to place in the multiset variable, it is mapped to itself (to avoid unnecessary renamings). This is needed to compute the constraint theory, after which such identity substitutions can be eliminated.

Constraint theory For a unifier σ , the constraint theory \mathcal{T}_σ is defined as:

$$\bigcup \{M'_i \mid M_i \mapsto \{ts|M'_i\} \in \sigma\} \equiv \bigcup \{N'_i \mid N_i \mapsto \{ts|N'_i\} \in \sigma\}$$

Soundness and completeness of the algorithm are straightforward, since it exhaustively checks all possibilities for unifying multisets.

Theorem 1. *Soundness* If $\text{unify}(\mathcal{M}, \mathcal{N}) \mapsto \{\sigma_1, \dots, \sigma_n\}$ then $\forall \sigma_i. \mathcal{M}\sigma_i =_{\mathcal{T}_{\sigma_i}} \mathcal{N}\sigma_i$

Theorem 2. *Completeness* If $\exists \sigma. \mathcal{M}\sigma =_{\mathcal{T}_\sigma} \mathcal{N}\sigma$ then $\text{unify}(\mathcal{M}, \mathcal{N}) \mapsto \{\sigma_1, \dots, \sigma_n\}$ and $\exists \sigma_i$ such that $\sigma = \sigma_i\sigma'$.

Note that the use of a substitution σ' is needed even if the set of computed unifiers is not the minimal one.

2.2 Complexity

The most expensive part of the unification algorithm is the one between lines 11 and 17 in Algorithm 1, so we concentrate our complexity analysis to that case. For each i from 0 to the minimum number of terms, `unify_c` is called. This function has two nested loops over the sets T_c and S_c (Alg. 2, lines 4, 5), which contain all possible ways of choosing i elements from n and m , respectively. Thus $|T_c| = \binom{n}{i}$ and $|S_c| = \binom{m}{i}$. In the inner part of the loops, `unify_terms` is called, which finds all possible unifiers for two multisets of size i . Since all possible pairings of elements must be tried, and for each order `unify_lists` runs in i^2 , the function (Alg. 3) has complexity $i^2 i!$. The function `unify_distribute` (Alg. 4) computes all possible ways of partitioning $n - i$ elements into l parts, and $m - i$ elements into k parts, and returns the Cartesian product of these sets. Therefore, its complexity is $l^{n-i} k^{m-i}$.

Putting those together, we get to the cost for the unification of mmsets $\{t_1, \dots, t_n | M_1, \dots, M_k\}$ and $\{s_1, \dots, s_m | N_1, \dots, N_l\}$:

$$\sum_{i=0}^{\min(n,m)} \binom{n}{i} \binom{m}{i} i^2 i! l^{n-i} k^{m-i}$$

After some arithmetic manipulation, we can conclude that, on the worst case, `unify` runs in $O(n! m! l^n k^m)$. For the special case where the multisets have only one multiset variable: $l = k = 1$, and the unification algorithm runs in $O(n! m!)$.

2.3 Removing Less General Unifiers

The algorithm described in Section 2.1 does not compute the set of minimal unifiers. For example, given multiset $\{a, a | M\}$ and $\{a | N\}$, `unify` computes three unifiers with constraint: $\{N \mapsto \{a | N'\}\}$, where $M \equiv N'$, twice (once for each occurrence of a), and $\{M \mapsto \{a | M'\} ; N \mapsto \{a, a | N'\}\}$, where $M' \equiv N'$. These are the same unifiers obtained by the non-deterministic algorithm from [5].

In order to reduce the number of less general unifiers, we can remove every pair of equal terms t_i and s_j from the mmsets (i.e., t_i and s_j unify with the empty substitution). The rationale behind this is that, every other unifier that is obtained by unifying these terms with something else, or placing them in a multiset variable, can be recovered from the set of unifiers obtained when this pair is not in the mmset.

We start by showing that it is safe to eliminate pairs of equal terms from the problem of unifying multisets without multiset variables.

Theorem 3. *Let \bar{t} and \bar{s} be two multisets of terms such that $t_a \in \bar{t}$ and $s_b \in \bar{s}$ are equal, for some a and b . Let $\Sigma_{all} = \text{unify_terms}(\bar{t}, \bar{s})$, and $\Sigma = \text{unify_terms}(\bar{t} \setminus \{t_a\}, \bar{s} \setminus \{s_b\})$. Then for every $\sigma \in \Sigma_{all}$, there exists $\mu \in \Sigma$ s.t. $\sigma = \mu\sigma'$ for some substitution σ' .*

The proof for this theorem can be found in Appendix B. The overall idea is as follows. σ was obtained by some pairing of terms in \bar{t} and \bar{s} . We choose μ as the unifier that used a pairing that is as close as possible as the one used for σ . Those pairings differ only for the terms involving t_a and s_b . Suppose t_a is paired with s_x and s_b is paired with t_y . Using the most general unifiers of t_y and s_x , we can conclude the existence of σ' such that $\sigma = \mu\sigma'$.

Theorem 4. *Let $\{\bar{t} | \bar{M}\}$ and $\{\bar{s} | \bar{N}\}$ be two mmsets such that $t_a \in \bar{t}$ is equal to $s_b \in \bar{s}$ for some a and b . Moreover, let $\Sigma_{all} = \text{unify}(\{\bar{t} | \bar{M}\}, \{\bar{s} | \bar{N}\})$, and $\Sigma = \text{unify}(\{\bar{t} \setminus \{t_a\} | \bar{M}\}, \{\bar{s} \setminus \{s_b\} | \bar{N}\})$. Then for every $\sigma \in \Sigma_{all}$ there exists $\mu \in \Sigma$ such that $\sigma = \mu\sigma'$ for some σ' .*

The proof for this theorem can be found in Appendix B. We proceed by a case analysis on whether t_a and s_b were chosen to be unified, or to be placed in a multiset variable. There are four cases. The case in which both are chosen to be unified is solved using Theorem 3. For the case in which both are placed in a multiset variable, we can construct σ' . The other two (dual) cases are the more involved ones. They use a combination of the two strategies of the first cases.

This modification is implemented in the algorithm available online, and extensive testing has shown that all less general unifiers are eliminated. In particular, only the unifier $\{N \mapsto \{a|N'\}\}$, where $M \equiv N'$ is computed for mmsets $\{a, a|M\}$ and $\{a|N\}$.

3 Conclusion

We have developed a sound and complete algorithm for finding unifiers of multisets with multiple multiset variables. The algorithm is deterministic and terminating. It is implemented in SML, and we also provide the pseudo code for reproducibility. The same algorithm can be used for the particular case where there is only one multiset variable.

The complexity of the unification procedure is analysed, and its cost is high. This is inherent to the problem, since it is of combinatorial nature. It may be possible to improve this result by using the right data-structures and heuristics, but for our purposes, since the numbers are quite small, it runs fast enough.

We have also tried to eliminate all sources of redundancy, so that the set of computed unifiers is as close as possible to the minimal one. In particular, we have shown that a simple pre-processing of the input problem will produce fewer unifiers, and all those that are no longer produced can be recovered. We conjecture that this optimization leads to the computation of the *minimal* set of unifiers, but we leave this investigation as future work.

References

- [1] I. Cervesato. Solution Count for Multiset Unification with Trailing Multiset Variables. In C. Ringers, C. Tinelli, F. Trinen, and R. Verma, editors, *Sixteenth International Workshop on Unification — UNIF'02*, pages 64–68, 2002.
- [2] K. Chaudhuri, L. Lima, and G. Reis. Formalized Meta-Theory of Sequent Calculi for Substructural Logics. *Electronic Notes in Theoretical Computer Science*, 332:57 – 73, 2017. LSFA 2016 - 11th Workshop on Logical and Semantic Frameworks with Applications (LSFA).
- [3] J. E. Dawson and R. Goré. Generic Methods for Formalising Sequent Calculi Applied to Provability Logic. In *Logic for Programming, Artificial Intelligence, and Reasoning - 17th International Conference, LPAR-17, 2010. Proceedings*, pages 263–277, 2010.
- [4] A. Dovier, A. Policriti, and G. Rossi. Integrating Lists, Multisets, and Sets in a Logic Programming Framework. In F. Baader and K. U. Schulz, editors, *Frontiers of Combining Systems: First International Workshop, Munich, March 1996*, pages 303–319. Springer Netherlands, 1996.
- [5] A. Dovier, A. Policriti, and G. Rossi. A uniform axiomatic view of lists, multisets, and sets, and the relevant unification algorithms. *Fundam. Inf.*, 36(2-3):201–234, 1998.
- [6] H. Tews. Formalizing Cut Elimination of Coalgebraic Logics in Coq. In *Automated Reasoning with Analytic Tableaux and Related Methods: 22nd International Conference, TABLEAUX 2013, Proceedings*, pages 257–272. Springer, 2013.
- [7] B. Xavier, C. Olarte, G. Reis, and V. Nigam. Mechanizing Linear Logic in Coq. In *Proceedings of the 12th Workshop on Logical and Semantic Frameworks with Applications (LSFA)*, pages 60–77, 2017.

A Algorithms

Algorithm 1 $\text{unify}(\{t_1, \dots, t_n | M_1, \dots, M_k\}, \{s_1, \dots, s_m | N_1, \dots, N_l\})$

```

1: if  $k = 0 \wedge l = 0$  then
2:   if  $n = m$  then return  $\text{unify\_terms}([t_1, \dots, t_n], [s_1, \dots, s_m])$ 
3:   else return  $\square$ 
4: else if  $k = 0$  then
5:   if  $m \leq n$  then return  $\text{unify\_c}(m, \{t_1, \dots, t_n | M_1, \dots, M_k\}, \{s_1, \dots, s_m | N_1, \dots, N_l\})$ 
6:   else return  $\square$ 
7: else if  $l = 0$  then
8:   if  $n \leq m$  then return  $\text{unify\_c}(n, \{t_1, \dots, t_n | M_1, \dots, M_k\}, \{s_1, \dots, s_m | N_1, \dots, N_l\})$ 
9:   else return  $\square$ 
10: else
11:    $c \leftarrow \min(n, m)$ 
12:    $\Sigma \leftarrow \square$ 
13:   for  $i = 0$  to  $c$  do
14:      $\Sigma' \leftarrow \text{unify\_c}(i, \{t_1, \dots, t_n | M_1, \dots, M_k\}, \{s_1, \dots, s_m | N_1, \dots, N_l\})$ 
15:      $\Sigma \leftarrow \Sigma \cup \Sigma'$ 
16:   end for
17:   return  $\Sigma$ 
18: end if

```

Algorithm 2 $\text{unify_c}(c, \{t_1, \dots, t_n | M_1, \dots, M_k\}, \{s_1, \dots, s_m | N_1, \dots, N_l\})$

```

1:  $\Sigma \leftarrow \square$ 
2:  $T_c \leftarrow \text{choose}(\{t_1, \dots, t_n\}, c)$ 
3:  $S_c \leftarrow \text{choose}(\{s_1, \dots, s_m\}, c)$ 
4: for  $(\{t'_1, \dots, t'_c\}, \{t'_{c+1}, \dots, t'_n\}) \in T_c$  do
5:   for  $(\{s'_1, \dots, s'_c\}, \{s'_{c+1}, \dots, s'_m\}) \in S_c$  do
6:      $\Sigma_t \leftarrow \text{unify\_terms}([t'_1, \dots, t'_c], [s'_1, \dots, s'_c])$ 
7:     if  $\Sigma_t \neq \square$  then
8:        $\Sigma_M \leftarrow \text{unify\_distribute}(\{t'_{c+1}, \dots, t'_n\}, \{M_1, \dots, M_k\}, (\{s'_{c+1}, \dots, s'_m\}, \{N_1, \dots, N_l\}))$ 
9:        $\Sigma' \leftarrow \text{map}(\lambda(\sigma_t, \sigma_M). \sigma_M \sigma_t)(\Sigma_t \times \Sigma_M)$ 
10:       $\Sigma \leftarrow \Sigma \cup \Sigma'$ 
11:     end if
12:   end for
13: end for
14: return  $\Sigma$ 

```

Algorithm 3 unify_terms($([t_1, \dots, t_n], [s_1, \dots, s_n])$)

```

1:  $\Sigma \leftarrow []$ 
2:  $P_s \leftarrow \text{permutations}([s_1, \dots, s_n])$ 
3:  $P \leftarrow [[t_1, \dots, t_n]] \times P_s$ 
4: for  $(T, S) \in P$  do
5:    $\sigma \leftarrow \text{unify\_lists}(T, S)$ 
6:   if  $\sigma \neq \text{None}$  then  $\Sigma \leftarrow \{\sigma\} \cup \Sigma$ 
7: end for

```

Algorithm 4 unify_distribute($(\{\{t_1, \dots, t_n\}, \{M_1, \dots, M_k\}\}, (\{s_1, \dots, s_m\}, \{N_1, \dots, N_l\}))$)

```

1:  $\Sigma_N \leftarrow \text{distribute}(\{\{t_1, \dots, t_n\}, \{N_1, \dots, N_l\}, k = 0\}$  {List of substitutions for  $N_i$ }
2:  $\Sigma_M \leftarrow \text{distribute}(\{s_1, \dots, s_m\}, \{M_1, \dots, M_k\}, l = 0\}$  {List of substitutions for  $M_i$ }
3:  $\Sigma \leftarrow \text{map}(\lambda(\sigma_M, \sigma_N). \sigma_M \cup \sigma_N)(\Sigma_M \times \Sigma_N)$ 
4: return  $\Sigma$ 

```

Algorithm 5 distribute($(\{\{t_1, \dots, t_n\}, \{M_1, \dots, M_k\}, \text{exact})$)

```

1: if  $n = 0 \wedge k = 0$  then
2:   return  $\{\{\}\}$ 
3: end if
4:  $\Sigma \leftarrow []$ 
5:  $P_t \leftarrow \text{ordered\_partitions}(\{\{t_1, \dots, t_n\}, k)$ 
6: for  $p \in P_t$  do
7:    $\sigma \leftarrow \{\}$ 
8:   for  $i = 1$  to  $k$  do
9:      $ts \leftarrow p[i]$ 
10:    if exact then  $\sigma \leftarrow \sigma\{M_i \mapsto \{ts | \cdot\}\}$ 
11:    if  $\neg \text{exact} \wedge ts = \emptyset$  then  $\sigma \leftarrow \sigma\{M_i \mapsto \{\cdot | M_i\}\}$ 
12:    if  $\neg \text{exact} \wedge ts \neq \emptyset$  then  $\sigma \leftarrow \sigma\{M_i \mapsto \{ts | M'_i\}\}$ 
13:   end for
14:    $\Sigma \leftarrow \{\sigma\} \cup \Sigma$ 
15: end for
16: return  $\Sigma$ 

```

B Proofs

Proof for Theorem 3. We know that Σ_{all} contains at most $n!$ unifiers, one for each way of pairing elements of \bar{t} with elements of \bar{s} . Analogously, Σ contains at most $(n-1)!$ unifiers. Let: $\Sigma_{\text{all}} = \{\sigma_1, \dots, \sigma_{n!}\}$ and $\Sigma = \{\mu_1, \dots, \mu_{(n-1)!}\}$. Then each σ_i can be obtained from some μ_j .

If σ_i is a unifier resulting from pairing t_a with s_b , then there exists $\mu_j = \sigma_i$ and we are done.

Let σ_i be a unifier resulting from pairing t_a with some s_x , s_b with some t_y , and some permutation P_t of $\bar{t} \setminus \{t_a, t_y\}$ with some permutation P_s of $\bar{s} \setminus \{s_b, s_x\}$. There exists a unifier $\mu_j \in \Sigma$ that is the result of unifying the same permutations P_t and P_s , and t_y with s_x . We show how σ_i can be reconstructed from μ_j . Since the order in which terms are unified does not matter, we assume that σ_i and μ_j are obtained as follows:

1. Computation of σ_i :

- (a) mgu σ_{ax} of t_a and s_x
- (b) mgu σ_{by} of $t_y\sigma_{ax}$ and $s_b\sigma_{ax}$
- (c) mgu σ_{P_i} of $P_t\sigma_{ax}\sigma_{by}$ and $P_t\sigma_{ax}\sigma_{by}$

2. Computation of μ_j :

- (a) mgu σ_{xy} of t_y and s_x
- (b) mgu σ_{P_j} of $P_t\sigma_{xy}$ and $P_s\sigma_{xy}$

Therefore $\sigma_i = \sigma_{ax}\sigma_{by}\sigma_{P_i}$ and $\mu_j = \sigma_{xy}\sigma_{P_j}$. We know that:

$$t_a\sigma_{ax} = s_x\sigma_{ax} \quad \text{from 1a} \quad (1)$$

$$t_y\sigma_{ax}\sigma_{by} = s_b\sigma_{ax}\sigma_{by} \quad \text{from 1b} \quad (2)$$

$$t_y\sigma_{ax}\sigma_{by} = t_a\sigma_{ax}\sigma_{by} \quad \text{because } t_a = s_b \quad (3)$$

$$t_y\sigma_{ax}\sigma_{by} = s_x\sigma_{ax}\sigma_{by} \quad \text{from 3 and 1} \quad (4)$$

Thus, $\sigma_{ax}\sigma_{by}$ is a unifier of t_y and s_x . But from 2a we have that σ_{xy} is the most general unifiers of these terms, which means that there exists a σ' such that

$$\sigma_{ax}\sigma_{by} = \sigma_{xy}\sigma' \quad (5)$$

From 5 and 1c, we know that $P_t\sigma_{xy}\sigma'\sigma_{P_i} = P_s\sigma_{xy}\sigma'\sigma_{P_i}$, meaning that $\sigma'\sigma_{P_i}$ is a unifier of $P_t\sigma_{xy}$ and $P_s\sigma_{xy}$. But from 2b, σ_{P_j} is the most general unifier of these two lists, therefore, there exists σ'' such that

$$\sigma'\sigma_{P_i} = \sigma_{P_j}\sigma'' \quad (6)$$

Using 5, 6, and associativity of substitution composition:

$$\sigma_i = \sigma_{ax}\sigma_{by}\sigma_{P_i} = \sigma_{xy}\sigma_{P_j}\sigma'' = \mu_j\sigma''$$

□

Proof for Theorem 4. We case on how t_a and s_b were used to compute σ . Let \vec{t}_c and \vec{s}_c denote the terms and order chosen from \vec{t} and \vec{s} , respectively, to be unified. Let \vec{t}_r and \vec{s}_r denote the rest of the terms in \vec{t} and \vec{s} that will be distributed to \vec{N} and \vec{M} , respectively.

We know that $\sigma = \sigma_M\sigma_t$ (Alg. 2, line 9), where σ_t is the unifier of \vec{t}_c and \vec{s}_c , and σ_M is obtained from partitions denoted as $\pi_t(\vec{t}_r)$ and $\pi_s(\vec{s}_r)$.

1. $t_a \in \vec{t}_c$ and $s_b \in \vec{s}_c$

Take $\mu = \mu_M\mu_t$ such that $\mu_M = \sigma_M$ is obtained from the same partitions $\pi_t(\vec{t}_r)$ and $\pi_s(\vec{s}_r)$, and μ_t is such that $\sigma_t = \mu_t\sigma'$ for some σ' . The existence of such μ_t is guaranteed by Theorem 3. Therefore, $\sigma = \sigma_M\sigma_t = \mu_M\mu_t\sigma' = \mu\sigma'$.

2. $t_a \in \vec{t}_r$ and $s_b \in \vec{s}_r$

Let \vec{t}_{p_a}, t_a be the part from $\pi_t(\vec{t}_r)$ with t_a , and \vec{s}_{p_b}, s_b the part from $\pi_s(\vec{s}_r)$ with s_b .

Take $\mu = \mu_M\mu_t$ such that $\mu_t = \sigma_t$ is the unifier of \vec{t}_c and \vec{s}_c , and μ_M is obtained from partitions $\pi_t(\vec{t}_r)$ where \vec{t}_{p_a}, t_a is replaced by \vec{t}_{p_a} and analogously for $\pi_s(\vec{s}_r)$. Thus the mappings in σ_M and μ_M are the same, except for two multiset variables N_a and M_b :

$$\begin{aligned} \{M_b \mapsto \vec{s}_{p_b}, s_b, M'_b\} \ ; \ N_a \mapsto \vec{t}_{p_a}, t_a, N'_a\} &\subset \sigma_M \\ \{M_b \mapsto \vec{s}_{p_b}, M'_b\} \ ; \ N_a \mapsto \vec{t}_{p_a}, N'_a\} &\subset \mu_M \end{aligned}$$

Since M'_b and N'_a are fresh multiset variable names:

$$\sigma_M = \mu_M\{M'_b \mapsto s_b, M'_b; N'_a \mapsto t_a, N'_a\}$$

And since $\sigma_t = \mu_t$:

$$\sigma_M\sigma_t = \mu_M\{M'_b \mapsto s_b, M'_b; N'_a \mapsto t_a, N'_a\}\mu_t$$

The image of μ_t does not contain M'_b nor N'_a , therefore:

$$\sigma_M\sigma_t = \mu_M\mu_t\{M'_b \mapsto s_b, M'_b; N'_a \mapsto t_a, N'_a\}\mu_t$$

Thus:

$$\sigma = \mu\{M'_b \mapsto s_b, M'_b; N'_a \mapsto t_a, N'_a\}\mu_t$$

3. $\boxed{t_a \in \vec{t}_c \text{ and } s_b \in \overline{s_r}}$

Let s_k be the term from \vec{s}_c that is paired with t_a and $\overline{s_{p_b}}, s_b$ be the part from $\pi_s(\overline{s_r})$ containing s_b . Assume that t_a is unified with s_k with mgu σ_{ak} , and that $(\vec{t}_c \setminus \{t_a\})\sigma_{ak}$ unifies with $(\vec{s}_c \setminus \{s_k\})\sigma_{ak}$ with mgu σ_c . Thus $\sigma_t = \sigma_{ak}\sigma_c$. Let M_b be the variable to which partition $\overline{s_{p_b}}, s_b$ is assigned. Thus:

$$\{M_b \mapsto \overline{s_{p_b}}, s_b, M'_b\} \subset \sigma_M \quad \text{By definition} \quad (1)$$

$$\{M_b \mapsto \overline{s_{p_b}}\sigma_{ak}, s_b\sigma_{ak}, M'_b\} \subset \sigma_M\sigma_{ak} \quad \text{Composition with } \sigma_{ak} \quad (2)$$

$$\{M_b \mapsto \overline{s_{p_b}}\sigma_{ak}, t_a\sigma_{ak}, M'_b\} \subset \sigma_M\sigma_{ak} \quad t_a = s_b \quad (3)$$

$$\{M_b \mapsto \overline{s_{p_b}}\sigma_{ak}, s_k\sigma_{ak}, M'_b\} \subset \sigma_M\sigma_{ak} \quad t_a\sigma_{ak} = s_k\sigma_{ak} \quad (4)$$

Take $\mu = \mu_M\mu_t$ such that the terms $\vec{t}_c \setminus \{t_a\}$ and $\vec{s}_c \setminus \{s_k\}$ are unified with mgu μ_t , and μ_M is computed using partition $\pi_t(\overline{s_r})$ and $\pi_s(\overline{s_r})$ where part $\overline{s_{p_b}}, s_b$ is replaced by $\overline{s_{p_b}}, s_k$. Therefore,

$$\{M_b \mapsto \overline{s_{p_b}}, s_k, M'_b\} \subset \mu_M \quad (5)$$

Because μ_t is the mgu of the two lists, we have that: $\sigma_t = \sigma_{ak}\sigma_c = \mu_t\sigma'_t$ for some σ'_t . And from 4 and 5 we can also conclude: $\sigma_M\sigma_{ak} = \mu_M\sigma_{ak}$. Using these equalities: $\sigma = \sigma_M\sigma_t = \sigma_M\sigma_{ak}\sigma_c = \mu_M\sigma_{ak}\sigma_c = \mu_M\mu_t\sigma'_t = \mu\sigma'_t$.

4. $\boxed{t_a \in \overline{t_r} \text{ and } s_b \in \vec{s}_c}$ Analogous to the previous case.

□