

ON NOMINAL SYNTAX AND PERMUTATION FIXED POINTS

MAURICIO AYALA-RINCÓN, MARIBEL FERNÁNDEZ, AND DANIELE NANTES-SOBRINHO

Departments Mathematics and Computer Science, Universidade de Brasília, Brasília, Brazil
e-mail address: ayala@unb.br

Department of Informatics, King’s College London, London, UK
e-mail address: maribel.fernandez@kcl.ac.uk

Department Mathematics, Universidade de Brasília, Brasília, Brazil
e-mail address: dnantes@unb.br

ABSTRACT. We propose a new axiomatisation of the alpha-equivalence relation for nominal terms, based on a primitive notion of fixed-point constraint. We show that the standard freshness relation between atoms and terms can be derived from the more primitive notion of permutation fixed-point, and use this result to prove the correctness of the new α -equivalence axiomatisation. This gives rise to a new notion of nominal unification, where solutions for unification problems are pairs of a fixed-point context and a substitution. Although it may seem less natural than the standard notion of nominal unifier based on freshness constraints, the notion of unifier based on fixed-point constraints behaves better when equational theories are considered: for example, nominal unification remains finitary in the presence of commutativity, whereas this is not the case when unifiers are expressed using freshness contexts. We provide a definition of α -equivalence modulo equational theories that takes into account A, C and AC theories. Based on this notion of equivalence, we show that C-unification is finitary and we provide a sound and complete C-unification algorithm, as a first step towards the development of nominal unification modulo AC and other equational theories with permutative properties.

1. INTRODUCTION

This paper presents a new approach for the definition of nominal languages, based on the use of permutation fixed points. More precisely, we give a new axiomatisation of the α -equivalence relation for nominal terms using permutation fixed-points, and revisit nominal unification in this setting.

In nominal syntax [UPG04], *atoms* are used to represent object-level variables and *atom permutations* are used to implement renamings, following the nominal-sets approach advocated by Gabbay and Pitts [Gab00, GP02, Pit13]. Atoms can be abstracted over terms; the syntax $[a]s$ represents the abstraction of a in s . To rename an abstracted atom a to b ,

Key words and phrases: alpha-conversion, nominal syntax, unification, equational theories.

Work supported by FAPDF grant 193001369/2016. M. Ayala-Rincón partially funded by CNPq research grant number 307672/2017-4.

a *swapping* permutation $\pi = (ab)$ is applied. Thus, the action of π over $[a]s$, written as $(ab) \cdot [a]s$, produces the nominal term $[b]s'$, where s' is the result of swapping a and b in s . The α -equivalence relation between nominal terms is specified using swappings together with a *freshness relation* between atoms and terms, written $b\#s$, which roughly corresponds to b not occurring free in s .

In this setting, checking α -equivalence requires another first-order specialised calculus to check freshness constraints. For instance, checking whether $[a]s \approx_\alpha [b]t$ reduces to checking whether $s \approx_\alpha (ba) \cdot t$ and $a\#t$.

The action of a permutation propagates down the structure of nominal terms, until a variable is reached: permutations suspend over variables. Thus, $\pi \cdot s$ represents the action of a permutation over a nominal term, but is not itself a nominal term unless s is a variable; for instance, $\pi \cdot X$ is a *suspension* (also called *moderated variable*), which is a nominal term.

The presence of moderated variables and atom-abstractions makes reasoning about equality of nominal terms more involved than in standard first-order syntax. For example, $\pi \cdot X \approx_\alpha \rho \cdot X$ is only true when X ranges over nominal terms, say s , for which all atoms in the difference set of π and ρ (i.e., the set $\{a : \pi(a) \neq \rho(a)\}$) are fresh in s .

If the domain of a permutation π is fresh for X then $\pi \cdot X \approx_\alpha Id \cdot X$. Thus a set of freshness constraints (i.e., a freshness context) can be used to specify that a permutation will have no effect on the instances of X . This is why in *nominal unification* [UPG04], the solution for a problem is a pair consisting of a freshness context and a substitution.

The use of freshness contexts is natural when dealing with “syntactic” nominal unification, but in the presence of equational axioms (i.e., equational nominal unification) it is not straightforward. For example, in the case of C-nominal unification (nominal unification modulo commutativity), to specify that a permutation has no effect on the instances of X modulo C, in other words, to specify that the permutation does not affect a given C-equivalence class, we need something more than a freshness constraint (note that $(ab) \cdot (a + b) = b + a =_C a + b$, so the permutation (ab) fixes the term $a + b$, despite the fact that a and b are not fresh).

In this paper, we propose to axiomatise α -equivalence of nominal terms using permutation fixed-point constraints: we write $\pi \lambda t$ (read “ π fixes t ”) if t is a fixed-point of π . We show how to derive fixed-point constraints $\pi \lambda t$ from primitive constraints of the form $\pi \lambda X$, and show the correctness of this approach by proving that the α -equivalence relation generated in this way coincides with the one axiomatised via freshness constraints. We then show how fixed-point constraints can be used to specify α -equivalence modulo equational theories containing A, C, and AC operators, and provide an algorithm to solve nominal unification problems modulo C, which outputs a finite set of most general solutions.

Related Work. Equational reasoning has been extensively explored since the early development of modern abstract algebra (see, e.g., the E -unification surveys by Siekmann [Sie90] and Baader et al [BS94] and [BSN⁺01]). For AC-equality checking, AC matching and AC unification, refined techniques have been applied. For instance, AC-equality check and linear AC-matching problems can be reduced to searching a perfect matching in a bipartite graph [BKN87], whereas AC unification problems can be reduced to solving a system of Diophantine equations [Sti81, Fag87].

Techniques to deal with α -equivalence modulo the equational theories A, C and AC were proposed in [AdCSFN17, AdCSMFR19, ARdCSFNS17, ARdCSFNS18], using the standard nominal approach via freshness constraints. Solving nominal C-unification problems

requires to deal with fixed-point equations, for which there is no finitary representation of the set of solutions using only freshness constraints and substitutions [ARdCSFNS17, ARdCSFNS18]. A combinatorial algorithm permits to find all the solutions of fixed-point equations [ARdCSFNS18].

Fixed-point constraints arise also in other contexts: Schmidt-Schauß et al [SKLV17] show how nominal unification problems in a language with recursive let operators give rise to freshness constraints and nominal fixed-point equations. The approach to nominal unification via permutation fixed-points proposed in this paper could also be used to reason about equality in this language.

This paper is a revised and extended version of [AFN18], where nominal terms with fixed-point permutation constraints were first presented. In this paper we prove the correctness of the approach, by showing it is equivalent to the standard presentation via freshness constraints. We also provide proofs of soundness and completeness of the unification algorithm. A generalisation of the notion of α -equivalence to take into account equational theories including **A**, **C** and **AC** is also provided (only **C** was considered in [AFN18]), as well as a **C**-unification algorithm.

Overview. The paper is organised as follows. Section 2 provides the necessary background material on nominal syntax and semantics. Section 3 introduces fixed-point constraints and α -equivalence, and shows that these relations behave as expected. In particular, we show that there is a two-way translation between the freshness-based α -equivalence relation and its permutation fixed-point counter-part, which confirms that the fixed-point approach is equivalent to the standard approach via freshness constraints. Section 4 presents a nominal unification algorithm specified by a set of simplification rules, and proves its soundness and completeness. In Section 5 we generalise the approach to take into account equational theories: we define a notion of permutation fixed-point α -equivalence modulo **A**, **C** and **AC** theories, and develop a **C**-unification algorithm. Finally, Section 6 concludes and discusses future work.

2. PRELIMINARIES

Let \mathbb{A} be a fixed and countably infinite set of elements a, b, c, \dots , which will be called *atoms* (atomic names). A permutation on \mathbb{A} is a bijection on \mathbb{A} with finite domain.

Fix a countably infinite set $\mathcal{X} = \{X, Y, Z, \dots\}$ of variables and a countable set $\mathcal{F} = \{f, g, \dots\}$ of function symbols.

Definition 2.1 (Nominal grammar). Nominal terms are generated by the following grammar.

$$s, t := a \mid [a]t \mid (t_1, \dots, t_n) \mid f^E t \mid \pi \cdot X$$

where a is an *atom term*, $[a]t$ denotes the *abstraction* of the atom a over the term t , (t_1, \dots, t_n) is a tuple, function symbols are equipped with an equational theory **E**, hence $f^E t$ denotes the *application of f^E to t* and $\pi \cdot X$ is a *moderated variable* or *suspension*, where π is an atom permutation. We write f^\emptyset , or simply f , to emphasise that no equational theory is assumed for f , that is, f is just a function symbol.

We follow the *permutative convention* [GM08, Convention 2.3] for atoms throughout the paper, i.e., atoms a, b, c range permutatively over \mathbb{A} so that they are always pairwise different, unless stated otherwise.

Atom *permutations* are represented by finite lists of *swappings*, which are pairs of different atoms $(a\ b)$; hence, a permutation π is generated by the following grammar:

$$\pi ::= Id \mid (a\ b)\pi.$$

where Id is the identity permutation, usually omitted from the list of swappings representing a permutation π . Suspensions of the form $Id \cdot X$ will be represented just by X . We write π^{-1} for the *inverse* of π , and use \circ to denote the composition of permutations. For example, if $\pi = (a\ b)(b\ c)$ then $\pi(c) = a$ and $c = \pi^{-1}(a)$.

The *difference set* of two permutations π, π' is $\mathbf{ds}(\pi, \pi') = \{a \mid \pi(a) \neq \pi'(a)\}$.

We write $\mathbf{Var}(t)$ for the set of variables occurring in t . Ground terms are terms without variables, that is, $\mathbf{Var}(t) = \emptyset$. A ground term may still contain atoms, for example a is a ground term and X is not.

Definition 2.2 (Permutation action). The action of a permutation π on a term t is defined by induction on the number of swappings in π :

$Id \cdot t = t$ and $((a\ b)\pi) \cdot t = (a\ b) \cdot (\pi \cdot t)$, where

$$\begin{array}{lll} (a\ b) \cdot a = b & (a\ b) \cdot (\pi \cdot X) = ((a\ b) \circ \pi) \cdot X & (a\ b) \cdot [a]t = [b](a\ b) \cdot t \\ (a\ b) \cdot b = a & (a\ b) \cdot \mathbf{f}\ t = \mathbf{f}\ (a\ b) \cdot t & (a\ b) \cdot [b]t = [a](a\ b) \cdot t \\ (a\ b) \cdot c = c & (a\ b) \cdot (t_1, \dots, t_n) = ((a\ b) \cdot t_1, \dots, (a\ b) \cdot t_n) & (a\ b) \cdot [c]t = [c](a\ b) \cdot t \end{array}$$

Definition 2.3 (Substitution). *Substitutions* are generated by the grammar

$$\sigma ::= id \mid [X \mapsto s]\sigma.$$

Postfix notation is used for substitution application and \circ for composition: $t(\sigma \circ \sigma') = (t\sigma)\sigma'$. Substitutions act on terms elementwise in the natural way: $t\ id = t$, $t[X \mapsto s]\sigma = (t[X \mapsto s])\sigma$, where

$$\begin{array}{ll} a[X \mapsto s] = a & (t_1, \dots, t_n)[X \mapsto s] = (t_1[X \mapsto s], \dots, t_n[X \mapsto s]) \\ (\mathbf{f}\ t)[X \mapsto s] = \mathbf{f}(t[X \mapsto s]) & (\pi \cdot X)[X \mapsto s] = \pi \cdot s \\ ([a]t)[X \mapsto s] = [a](t[X \mapsto s]) & (\pi \cdot Y)[X \mapsto s] = \pi \cdot Y \end{array}$$

The following well-known property of substitution and permutation justifies the notation $\pi \cdot \sigma$ (without brackets), see [UPG04] for more details.

Lemma 2.4. *Substitution and permutation commute: $\pi \cdot (\sigma) = (\pi \cdot s)\sigma$.*

2.1. Alpha-equivalence via freshness constraints. In the standard nominal approach (see, e.g., [Pit03, UPG04, FG07]), the α -equivalence relation $s \approx_\alpha t$ is defined using a freshness relation between atoms and terms, written $a\#t$ – read “ a fresh for t ” – which, intuitively, corresponds to the idea of an atom not occurring free in a term. These relations are axiomatised using the rules in Figures 1 and 2, respectively.

$$\begin{array}{lll} \frac{}{\Delta \vdash a\#b} (\#\mathbf{a}) & \frac{\pi^{-1}(a)\#X \in \Delta}{\Delta \vdash a\#\pi \cdot X} (\#\mathbf{var}) & \frac{\Delta \vdash a\#t}{\Delta \vdash a\#\mathbf{f}\ t} (\#\mathbf{f}) \\ \frac{\Delta \vdash a\#t_1 \quad \dots \quad \Delta \vdash a\#t_n}{\Delta \vdash a\#(t_1, \dots, t_n)} (\#\mathbf{tuple}) & \frac{}{\Delta \vdash a\#[a]t} (\#\mathbf{a}) & \frac{\Delta \vdash a\#t}{\Delta \vdash a\#[b]t} (\#\mathbf{abs}) \end{array}$$

Figure 1: Rules for freshness

We call $s \approx_\alpha t$ and $a \# t$ α -equality and freshness constraints, respectively. Note that to define \approx_α we use the difference set of two permutations in rule (\approx_α **var**); we denote by $\mathbf{ds}(\pi, \pi') \# X$ the following set of freshness constraints:

$$\mathbf{ds}(\pi, \pi') \# X = \{a \# X \mid a \in \mathbf{ds}(\pi, \pi')\}.$$

$\frac{}{\Delta \vdash a \approx_\alpha a} (\approx_\alpha \mathbf{a})$	$\frac{\mathbf{ds}(\pi, \pi') \# X \subseteq \Delta}{\Delta \vdash \pi \cdot X \approx_\alpha \pi' \cdot X} (\approx_\alpha \mathbf{var})$
$\frac{\Delta \vdash t \approx_\alpha t'}{\Delta \vdash f t \approx_\alpha f t'} (\approx_\alpha \mathbf{f})$	$\frac{\Delta \vdash t_1 \approx_\alpha t'_1 \quad \dots \quad \Delta \vdash t_n \approx_\alpha t'_n}{\Delta \vdash (t_1, \dots, t_n) \approx_\alpha (t'_1, \dots, t'_n)} (\approx_\alpha \mathbf{tuple})$
$\frac{\Delta \vdash t \approx_\alpha t'}{\Delta \vdash [a]t \approx_\alpha [a]t'} (\approx_\alpha \mathbf{[a]})$	$\frac{\Delta \vdash s \approx_\alpha (a b) \cdot t \quad \Delta \vdash a \# t}{\Delta \vdash [a]s \approx_\alpha [b]t} (\approx_\alpha \mathbf{ab})$

Figure 2: Rules for α -equality via freshness

The symbols Δ and ∇ denote *freshness contexts*, which are sets of freshness constraints of the form $a \# X$. The domain of a freshness context Δ , denoted by $\mathbf{dom}(\Delta)$, consists of the atoms occurring in Δ ; $\Delta|_X$ consists of the restriction of Δ to the freshness constraints on variable X , that is, the set $\{a \# X \mid a \# X \in \Delta\}$.

2.2. Nominal sets and support. Let S be a set equipped with an action of the group $\mathbf{Perm}(\mathbb{A})$ of finite permutations of \mathbb{A} .

Definition 2.5. A set $A \subset \mathbb{A}$ is a *support* for an element $x \in S$ if for all $\pi \in \mathbf{Perm}(\mathbb{A})$, the following holds

$$((\forall a \in A) \pi(a) = a) \Rightarrow \pi \cdot x = x \quad (2.1)$$

A *nominal set* is a set equipped with an action of the group $\mathbf{Perm}(\mathbb{A})$, that is, a $\mathbf{Perm}(\mathbb{A})$ -set, all of whose elements have finite support.

As in [Pit13], we denote by $\mathbf{supp}_S(x)$ the least finite support of x , that is,

$$\mathbf{supp}_S(x) := \bigcap \{A \in \mathcal{P}(\mathbb{A}) \mid A \text{ is a finite support for } x\}.$$

We write $\mathbf{supp}(x)$ when S is clear from the context. Clearly, each $a \in \mathbb{A}$ is finitely supported by $\{a\}$, therefore $\mathbf{supp}(a) = \{a\}$.

2.3. The “new” quantifier. In Nominal Logic [Pit03], the “new” quantifier \mathbb{N} is used to quantify over new names. Given two elements x, y of a nominal set, write $x \# y$ as an abbreviation of $\mathbf{supp}(x) \cap \mathbf{supp}(y) = \emptyset$; then $\mathbb{N}a.P(a, x)$ (read “for some/any new atom a , $P(a, x)$ ”) abbreviates $\forall a \in \mathbb{A}. a \# x \Rightarrow P(a, x)$ or equivalently $\exists a \in \mathbb{A}. a \# x \wedge P(a, x)$ (see Theorem 3.9 in [Pit13] for more details).

Intuitively, the “new” quantifier is used to indicate that a predicate holds for some (any) new atoms. The formula $\mathbb{N}a.\phi$ means “for a fresh name a , ϕ holds”. Since there is an infinite supply of names and elements of nominal sets have a finite support, it is always possible to pick a fresh name. Moreover, it can be shown that if a property $\phi(a)$ holds for some fresh name a then it holds for all fresh names, so the way in which the fresh atom is chosen is not important.

In proof systems for Nominal Logic, several approaches have been proposed to define rules to introduce and eliminate the \mathbb{N} quantifier (see, e.g., [Pit03, Che05, CU08]). We follow Cheney’s approach [Che05], where the formula $\mathbb{N}a.\phi$ (read “ a is a *newly quantified atom* in ϕ ”) is well formed if ϕ is a well-formed formula and $a \in \mathbb{A}$ is semantically fresh. In other words, the introduction of a \mathbb{N} quantifier for a in a formula ϕ requires a to be a new atom, fresh for all the variables in ϕ (see [Che05], Figure 5, where the well-formedness rules are given). We implicitly assume that a is a newly generated atom when we introduce \mathbb{N} in our formulas in the next section. Operationally, proving a formula $\mathbb{N}a.\phi$ boils down to picking a fresh atom c and proving that ϕ holds when the occurrences of a are replaced by the fresh atom c .

3. FIXED-POINT CONSTRAINTS

The native notion of equality on nominal terms is α -equivalence, written $s \approx_\alpha t$. As mentioned in section 2.1, this relation is usually axiomatised using the *freshness relation* between atoms and terms. The notion of freshness is derived from the notion of support, which in turn is defined using permutation fixed-points (see Definition 2.5). We can define freshness using the quantifier \mathbb{N} combined with a notion of fixed-point, as shown by Pitts [Pit13] (page 53):

$$a \# X \Leftrightarrow \mathbb{N}a'.(a \ a') \cdot X = X.$$

In this paper, instead of defining α -equivalence using freshness, we define it using the more primitive notion of *fixed-point* under the action of permutations. We will denote this relation $\overset{\wedge}{\approx}_\alpha$, and show that it coincides with \approx_α on ground terms, i.e., the relation defined using permutation fixed points corresponds to the relation defined using freshness. For non-ground terms, there is also a correspondence, but under different kinds of assumptions (fixed-point constraints vs. freshness constraints).

3.1. Fixed-points of permutations and term equality. We start by defining a binary relation that describes which elements of a nominal set S are fixed-points of a permutation $\pi \in \text{Perm}(\mathbb{A})$:

Definition 3.1 (Fixed-point relation). Let S be a nominal set. The *fixed-point relation* $\lambda \subseteq \text{Perm}(\mathbb{A}) \times S$ is defined as: $\pi \lambda x \Leftrightarrow \pi \cdot x = x$. Read “ $\pi \lambda x$ ” as “ π fixes x ”.

It is easy to see that if $\text{dom}(\pi) \cap \text{supp}(x) = \emptyset$ then $\pi \lambda x$ holds (this follows directly from the definition of support, see Definition 2.5). However, the other direction does not hold in general. For example, consider expressions built using atoms and a binary commutative operator $+$; the permutation $\pi = (a \ b)$ fixes the equivalence class of the expression $a + b$ despite the fact that its support coincides with $\text{dom}(\pi)$.

Permutation fixed-points will play an important role in the definition of α -equality of nominal terms. Below we define *fixed-point constraints* and *equality constraints* using predicates λ and $\overset{\wedge}{\approx}_\alpha$ and then give deduction rules to derive fixed-point and equality judgements.

Intuitively, for s and t ground nominal terms

- $s \overset{\wedge}{\approx}_\alpha t$ will mean that s and t are α -equivalent, i.e., equivalent modulo renaming of abstracted atoms.

- $\pi \lambda t$ will mean that the permutation π fixes the nominal term t , that is, $\pi \cdot t \stackrel{\lambda}{\approx}_\alpha t$. This means that π has “no effect” on t except for the renaming of bound names, for instance, $(a b) \lambda [a]a$ but not $(a b) \lambda f a$.

In the case of non-ground terms, a fixed-point or α -equality constraint has to be evaluated in a context, which provides information about permutations that fix the variables.

Definition 3.2 (Fixed-point and equality constraints). A *fixed-point constraint* is a pair $\pi \lambda t$ of a permutation π and a term t . An *α -equivalence constraint* is a pair of the form $s \stackrel{\lambda}{\approx}_\alpha t$.

We call a fixed-point constraint of the form $\pi \lambda X$ a *primitive fixed-point constraint* and a finite set of such constraints is called a *fixed-point context*. Υ, Ψ, \dots range over fixed-point contexts.

We write $\pi \lambda \mathbf{Var}(t)$ as an abbreviation for the set of constraints $\{\pi \lambda X \mid X \in \mathbf{Var}(t)\}$.

We now introduce some notation:

The set $\mathbf{Var}(\Upsilon)$ of variables is defined as expected: it contains all the variables mentioned in the fixed-point context Υ .

The set of permutations of a fixed-point context Υ with respect to the variable $X \in \mathbf{Var}(\Upsilon)$, denoted by $\mathbf{perm}(\Upsilon|_X)$, is defined as $\mathbf{perm}(\Upsilon|_X) := \{\pi \mid \pi \lambda X \in \Upsilon\}$.

For a substitution σ and a fixed-point context Υ we define $\Upsilon\sigma := \{\pi \lambda X\sigma \mid \pi \lambda X \in \Upsilon\}$.

To axiomatise the relation λ , we rely on the notion of *conjugation* of permutations. Conjugacy plays an important role in group theory [Hun80]. Recall that in a group G , the element b is a conjugate of a if there exists an element g such that $gag^{-1} = b$. It can be easily shown that conjugacy is an equivalence relation: if b is a conjugate of a then a is a conjugate of b . The conjugacy class of a consists of the elements of the form gag^{-1} , for $g \in G$. In the case of permutation groups, the conjugate of π with respect to ρ , denoted as π^ρ , is the result of the composition: $\rho \circ \pi \circ \rho^{-1}$.

$$\pi^\rho : \begin{array}{ccccccc} A & \xrightarrow{\rho^{-1}} & A & \xrightarrow{\pi} & A & \xrightarrow{\rho} & A \\ a & \mapsto & \rho^{-1}(a) & \mapsto & \pi(\rho^{-1}(a)) & \mapsto & \rho(\pi(\rho^{-1}(a))) \end{array}$$

The notion of *support* of a permutation or nominal term will be necessary for the next results. Considering permutations as elements of a nominal set, it follows from Definition 2.5 that $\mathbf{supp}(\pi) = \mathbf{dom}(\pi)$; similarly, in the case of a ground term t , the support coincides with its set of free names, i.e., $\mathbf{supp}(t) = \mathbf{fn}(t)$. In the case a term t is not ground, it should be considered in a context, say Υ , which gives information about its variables. We define it after introducing fixed-point and α -equality judgements.

Definition 3.3 (Judgements). A *fixed-point judgement* is a tuple $\Upsilon \vdash \mathbb{N}\bar{c}. \pi \lambda t$ of a fixed-point context and a fixed-point constraint where the atoms \bar{c} are newly quantified. Here $\bar{c} = c_1, \dots, c_n$, and $n = 0$ stands for empty quantification; in that case we may omit the \mathbb{N} quantifier and write simply $\Upsilon \vdash \pi \lambda t$.

An *α -equivalence judgement* is a tuple $\Psi \vdash \mathbb{N}\bar{c}. s \stackrel{\lambda}{\approx}_\alpha t$ of a fixed-point context and an equality constraint where the atoms \bar{c} are newly quantified.

The derivable fixed-point and α -equivalence judgements are defined by the rules in Figures 3 and 4. We give an example before explaining the rules in more detail.

Example 3.4. The term $[a]fa$ is a fixed-point for the permutation $(a b)$, since $(a b) \cdot [a]fa = [b]fb$, which is α -equivalent to $[a]fa$. As expected, we can derive $\vdash (a b) \lambda [a]fa$. For this we

use the rule $(\lambda \mathbf{abs})$, which requires to prove $\vdash \mathcal{I}c_1.(a \ b) \wedge (a \ c_1) \cdot fa$. Since $(a \ c_1) \cdot fa = fc_1$, we need to prove $\vdash \mathcal{I}c_1.(a \ b) \wedge c_1$, which holds by $(\lambda \mathbf{a})$. However, fa is not a fixed-point for $(a \ b)$: we cannot derive $\vdash (a \ b) \wedge fa$. This is to be expected, since fa and fb are not α -equivalent; we cannot derive $\vdash (a \ b) \cdot fa \stackrel{\wedge}{\approx}_\alpha fa$.

$\frac{\pi(a) = a}{\Upsilon \vdash \mathcal{I}\bar{c}.\pi \wedge a} (\lambda \mathbf{a})$	$\frac{\text{supp}(\pi^{\pi'^{-1}}) \setminus \{\bar{c}\} \subseteq \text{supp}(\text{perm}(\Upsilon _X))}{\Upsilon \vdash \mathcal{I}\bar{c}.\pi \wedge \pi' \cdot X} (\lambda \mathbf{var})$
$\frac{\Upsilon \vdash \mathcal{I}\bar{c}.\pi \wedge t}{\Upsilon \vdash \mathcal{I}\bar{c}.\pi \wedge f \ t} (\lambda \mathbf{f})$	$\frac{\Upsilon \vdash \mathcal{I}\bar{c}.\pi \wedge t_1 \quad \dots \quad \Upsilon \vdash \mathcal{I}\bar{c}.\pi \wedge t_n}{\Upsilon \vdash \mathcal{I}\bar{c}.\pi \wedge (t_1, \dots, t_n)} (\lambda \mathbf{tuple})$
$\frac{\Upsilon \vdash \mathcal{I}\bar{c}, c_1.\pi \wedge (a \ c_1) \cdot t}{\Upsilon \vdash \mathcal{I}\bar{c}.\pi \wedge [a]t} (\lambda \mathbf{abs})$	

Figure 3: Fixed-point rules.

The premise of rules $(\lambda \mathbf{abs})$ and $(\stackrel{\wedge}{\approx}_\alpha \mathbf{ab})$ introduces a newly quantified atom c_1 . Recall that the introduction of a \mathcal{I} quantifier requires the use of a new atom by well formedness (see Section 2.3). Therefore, in rules $(\lambda \mathbf{abs})$ and $(\stackrel{\wedge}{\approx}_\alpha \mathbf{ab})$ we can assume that the atom c_1 does not occur in the conclusion.

In rule $(\lambda \mathbf{var})$, the condition $\text{supp}(\pi^{\pi'^{-1}}) \setminus \{\bar{c}\} \subseteq \text{supp}(\text{perm}(\Upsilon|_X))$ means that the permutation $\pi^{\pi'^{-1}}$ can only affect atoms in $\text{perm}(\Upsilon|_X)$ and new atoms, hence it fixes X . Rules $(\lambda \mathbf{f})$ and $(\lambda \mathbf{tuple})$ are straightforward. Rule $(\lambda \mathbf{abs})$ is the most interesting one. The intuition behind this rule is the following: $[a]t$ is a fixed-point of π if $\pi \cdot [a]t$ is α -equivalent to $[a]t$, that is, $[\pi(a)]\pi \cdot t$ is α -equivalent to $[a]t$; the latter means that the only free atom in t that could be affected by π is a , hence, if we replace occurrences of a in t with another, new atom c , π should have no effect. Note that if a judgement $\Upsilon \vdash \mathcal{I}\bar{c}.\pi \wedge t$ is derivable, then so is $\Upsilon \vdash \mathcal{I}\bar{c}, \bar{c}'.\pi \wedge t$ for a new set of atoms \bar{c}' (there is an infinite supply of new atoms).

$\frac{}{\Upsilon \vdash \mathcal{I}\bar{c}. a \stackrel{\wedge}{\approx}_\alpha a} (\stackrel{\wedge}{\approx}_\alpha \mathbf{a})$	$\frac{\text{supp}((\pi')^{-1} \circ \pi) \setminus \{\bar{c}\} \subseteq \text{supp}(\text{perm}(\Upsilon _X))}{\Upsilon \vdash \mathcal{I}\bar{c}.\pi \cdot X \stackrel{\wedge}{\approx}_\alpha \pi' \cdot X} (\stackrel{\wedge}{\approx}_\alpha \mathbf{var})$
$\frac{\Upsilon \vdash \mathcal{I}\bar{c}. t \stackrel{\wedge}{\approx}_\alpha t'}{\Upsilon \vdash \mathcal{I}\bar{c}. f \ t \stackrel{\wedge}{\approx}_\alpha f \ t'} (\stackrel{\wedge}{\approx}_\alpha \mathbf{f})$	$\frac{\Upsilon \vdash \mathcal{I}\bar{c}. t_1 \stackrel{\wedge}{\approx}_\alpha t'_1 \quad \dots \quad \Upsilon \vdash \mathcal{I}\bar{c}. t_n \stackrel{\wedge}{\approx}_\alpha t'_n}{\Upsilon \vdash \mathcal{I}\bar{c}.(t_1, \dots, t_n) \stackrel{\wedge}{\approx}_\alpha (t'_1, \dots, t'_n)} (\stackrel{\wedge}{\approx}_\alpha \mathbf{tuple})$
$\frac{\Upsilon \vdash \mathcal{I}\bar{c}. t \stackrel{\wedge}{\approx}_\alpha t'}{\Upsilon \vdash \mathcal{I}\bar{c}. [a]t \stackrel{\wedge}{\approx}_\alpha [a]t'} (\stackrel{\wedge}{\approx}_\alpha \mathbf{[a]})$	$\frac{\Upsilon \vdash \mathcal{I}\bar{c}. s \stackrel{\wedge}{\approx}_\alpha (a \ b) \cdot t \quad \Upsilon \vdash \mathcal{I}\bar{c}, c_1.(a \ c_1) \wedge t}{\Upsilon \vdash \mathcal{I}\bar{c}. [a]s \stackrel{\wedge}{\approx}_\alpha [b]t} (\stackrel{\wedge}{\approx}_\alpha \mathbf{ab})$

Figure 4: Rules for α -equality.

The α -equivalence relation is defined using fixed-point judgements (see rule $(\stackrel{\wedge}{\approx}_\alpha \mathbf{ab})$). Rules $(\stackrel{\wedge}{\approx}_\alpha \mathbf{a})$, $(\stackrel{\wedge}{\approx}_\alpha \mathbf{f})$, $(\stackrel{\wedge}{\approx}_\alpha \mathbf{[a]})$ and $(\stackrel{\wedge}{\approx}_\alpha \mathbf{tuple})$ are defined as expected, whereas the intuition behind rule $(\stackrel{\wedge}{\approx}_\alpha \mathbf{var})$ is similar to the corresponding rule in Figure 3. The most interesting rule is $(\stackrel{\wedge}{\approx}_\alpha \mathbf{ab})$. Intuitively, it states that for two abstractions $[a]s$ and $[b]t$ to be equivalent,

we must obtain equivalent terms if we rename in one of them, in our case t , the abstracted atom b to a , so that they both use the same atom. Moreover, the atom a should not occur free in t , which is checked by stating that $(a\ c_1)$ fixes t for some newly generated atom c_1 .

Example 3.5. Notice that $(be) \wedge X \vdash (ab) \wedge [a](fa, (cd) \cdot X)$. This is derived as follows: by rule $(\wedge \mathbf{abs})$ we have to show $(be) \wedge X \vdash \mathcal{V}g.(ab) \wedge (fg, (ag)(cd) \cdot X)$; by rule $(\wedge \mathbf{tuple})$ this requires $(be) \wedge X \vdash \mathcal{V}g.(ab) \wedge fg$ and $(be) \wedge X \vdash \mathcal{V}g.(ab) \wedge (ag)(cd) \cdot X$; the former holds by application of rules $(\wedge \mathbf{f})$ and $(\wedge \mathbf{a})$ since $(ab) \cdot g = g$, and the latter by rule $(\wedge \mathbf{var})$, since $\mathbf{supp}((a\ b)^{(c\ d)(a\ g)}) \setminus \{g\} = \{g, b\} \setminus \{g\} = \{b\}$ and $\mathbf{supp}((be)) = \{b, e\}$.

However, it is not possible to derive $(ae) \wedge X \vdash (ab) \wedge [a](fa, (cd) \cdot X)$ since in this case b is not in the support of the permutations that fix X (only (ae) fixes X according to our context).

We prove below that $\overset{\wedge}{\approx}_\alpha$ is indeed an equivalence relation, for which we need to study the properties of the relations $\overset{\wedge}{\approx}_\alpha$ and \wedge , starting with *inversion* and *equivariance*.

Lemma 3.6 (Inversion). *The inference rules for $\overset{\wedge}{\approx}_\alpha$ are invertible.*

Proof. Straightforward since there is only one rule for each syntactic class of terms. Note that in the case of rule $(\overset{\wedge}{\approx}_\alpha \mathbf{ab})$, the permutative convention ensures that a and b are different atoms. \square

Equivariance is an important notion in nominal logic: equivariant subsets of nominal sets are closed under permutation (see Definition 1.8 in [Pit13]). Below we show that the fixed-point and α -equivalence relations are equivariant.

Lemma 3.7 (Equivariance).

- (1) $\Upsilon \vdash \mathcal{V}\bar{c}.\pi \wedge t$ iff $\Upsilon \vdash \mathcal{V}\bar{c}.\pi^\rho \wedge \rho \cdot t$, for any permutation ρ .
- (2) $\Upsilon \vdash \mathcal{V}\bar{c}.s \overset{\wedge}{\approx}_\alpha t$ iff $\Upsilon \vdash \mathcal{V}\bar{c}.\pi \cdot s \overset{\wedge}{\approx}_\alpha \pi \cdot t$, for any permutation π .

Proof. For both parts the “if” statement is trivial, by taking the identity permutation, Id . We prove the “only if” statements.

- (1) The proof is by rule induction using the rules in Figure 3. We consider cases depending on the last rule applied in the derivation.

- (a) The rule applied is $(\wedge \mathbf{a})$.

In this case, $t = a$. Notice that, $\Upsilon \vdash \pi' \wedge a$ if, and only if, $\pi'(a) = a$. We want to show that $\Upsilon \vdash \mathcal{V}\bar{c}.\pi^\rho \wedge \rho(a)$, for any permutation ρ . Notice that $\pi^\rho(\rho(a)) = (\rho \circ \pi \circ \rho^{-1})(\rho(a)) = (\rho \circ \pi)(a) = \rho(\pi(a)) = \rho(a)$, and the result follows.

- (b) The rule applied is $(\wedge \mathbf{var})$.

In this case one has $t = \pi_1 \cdot X$ and $\Upsilon \vdash \mathcal{V}\bar{c}.\pi \wedge \pi_1 \cdot X$. We need to prove that $\Upsilon \vdash \mathcal{V}\bar{c}.\pi^\rho \wedge \rho \cdot (\pi_1 \cdot X)$, which requires $\mathbf{supp}((\rho \circ \pi_1)^{-1} \circ \pi^\rho \circ (\rho \circ \pi_1)) \setminus \{\bar{c}\} \subseteq \mathbf{supp}(\mathbf{perm}(\Upsilon|_X))$.

Note that, since $\Upsilon \vdash \mathcal{V}\bar{c}.\pi \wedge \pi_1 \cdot X$, it follows that $\mathbf{supp}(\pi_1^{-1} \circ \pi \circ \pi_1) \setminus \{\bar{c}\} \subseteq \mathbf{supp}(\mathbf{perm}(\Upsilon|_X))$. By definition of π^ρ , it follows that

$$\begin{aligned} \mathbf{supp}((\rho \circ \pi_1)^{-1} \circ \pi^\rho \circ (\rho \circ \pi_1)) &= \mathbf{supp}((\rho \circ \pi_1)^{-1} \circ (\rho \circ \pi \circ \rho^{-1}) \circ (\rho \circ \pi_1)) \\ &= \mathbf{supp}((\pi_1^{-1} \circ \rho^{-1})) \circ (\rho \circ \pi \circ \rho^{-1}) \circ (\rho \circ \pi_1) \\ &= \mathbf{supp}(\pi_1^{-1} \circ \pi \circ \pi_1) \end{aligned}$$

Therefore, the result follows.

(c) The rule applied is $(\lambda \mathbf{abs})$.

In this case, $t = [a]t'$ and $\Upsilon \vdash \mathcal{V}\bar{c}. \pi \lambda [a]t'$, therefore, there exists a proof Π :

$$\frac{\frac{\Pi}{\Upsilon \vdash \mathcal{V}\bar{c}. \pi \lambda (a \ c_1) \cdot t'}}{\Upsilon \vdash \mathcal{V}\bar{c}. \pi \lambda [a]t'} (\lambda \mathbf{abs})$$

By induction hypothesis, for an arbitrary permutation ρ , there exists a derivation Π' for :

$$\Upsilon \vdash \mathcal{V}\bar{c}. \pi^\rho \lambda \rho \cdot ((a \ c_1) \cdot t').$$

Since $\rho \cdot ((a \ c_1) \cdot t') = (\rho(a) \ \rho(c_1)) \cdot (\rho \cdot t') = (\rho(a) \ c_1) \cdot (\rho \cdot t')$ because c_1 is a new name, we have

$$\frac{\frac{\Pi'}{\Upsilon \vdash \mathcal{V}\bar{c}. \pi^\rho \lambda (\rho(a) \ c_1) \cdot (\rho \cdot t')}}{\Upsilon \vdash \mathcal{V}\bar{c}. \pi^\rho \lambda [\rho(a)]\rho \cdot t'}$$

and the result follows.

(d) The rule applied is $(\lambda \mathbf{f})$ or $(\lambda \mathbf{tuple})$.

These cases follow directly by induction hypothesis.

(2) The proof is by rule induction using the rules in Figure 4.

(a) The rule applied is $(\overset{\wedge}{\approx}_\alpha \mathbf{a})$.

In this case, $s \overset{\wedge}{\approx}_\alpha t$ is an instance of $a \overset{\wedge}{\approx}_\alpha a$. It is straightforward to check that $\Upsilon \vdash \mathcal{V}\bar{c}. \pi \cdot a \overset{\wedge}{\approx}_\alpha \pi \cdot a$, via an application of the rule $(\overset{\wedge}{\approx}_\alpha \mathbf{a})$.

(b) The rule applied is $(\overset{\wedge}{\approx}_\alpha \mathbf{var})$.

In this case, $\Upsilon \vdash \mathcal{V}\bar{c}. \pi_1 \cdot X \overset{\wedge}{\approx}_\alpha \pi_2 \cdot X$ and $\mathbf{supp}(\pi_2^{-1} \circ \pi_1) \setminus \{\bar{c}\} \subseteq \mathbf{supp}(\mathbf{perm}(\Upsilon|_X))$. Notice that $\mathbf{supp}((\pi \circ \pi_2)^{-1} \circ \pi \circ \pi_1) = \mathbf{supp}(\pi_2^{-1} \circ \pi_1)$. Therefore, by rule $(\overset{\wedge}{\approx}_\alpha \mathbf{var})$, $\Upsilon \vdash \mathcal{V}\bar{c}. \pi \cdot (\pi_1 \cdot X) \overset{\wedge}{\approx}_\alpha \pi \cdot (\pi_2 \cdot X)$.

(c) The cases corresponding to rule $(\overset{\wedge}{\approx}_\alpha \mathbf{f})$ and rule $(\overset{\wedge}{\approx}_\alpha [\mathbf{a}])$ follow directly by induction hypothesis.

(d) The rule applied is $(\overset{\wedge}{\approx}_\alpha \mathbf{ab})$.

In this case, there exist derivations Π_1 and Π_2 such that

$$\frac{\frac{\Pi_1}{\Upsilon \vdash \mathcal{V}\bar{c}. t \overset{\wedge}{\approx}_\alpha (a \ b) \cdot s} \quad \frac{\Pi_2}{\Upsilon \vdash \mathcal{V}\bar{c}. (a \ c_1) \lambda s}}{\Upsilon \vdash \mathcal{V}\bar{c}. [a]t \overset{\wedge}{\approx}_\alpha [b]s} (\overset{\wedge}{\approx}_\alpha \mathbf{ab})$$

By induction hypothesis, and part (1) of this lemma:

(i) there exists a proof Π'_1 of $\Upsilon \vdash \mathcal{V}\bar{c}. \pi \cdot t \overset{\wedge}{\approx}_\alpha \pi \cdot ((a \ b) \cdot s)$.

(ii) there exists a proof Π'_2 of $\Upsilon \vdash \mathcal{V}\bar{c}. (a \ c_1)^\pi \lambda \pi \cdot s$.

It is easy to show that $\pi \cdot ((a \ b) \cdot s) = (\pi(a) \ \pi(b)) \cdot (\pi \cdot s)$, and since c_1 is new, $\pi(c_1) = c_1$, therefore, the result follows.

$$\frac{\frac{\Pi'_1}{\Upsilon \vdash \mathcal{V}\bar{c}. \pi \cdot t \overset{\wedge}{\approx}_\alpha (\pi(a) \ \pi(b)) \cdot \pi \cdot s} \quad \frac{\Pi'_2}{\Upsilon \vdash \mathcal{V}\bar{c}. (\pi(a) \ c_1) \lambda \pi \cdot s}}{\Upsilon \vdash \mathcal{V}\bar{c}. [\pi(a)]\pi \cdot t \overset{\wedge}{\approx}_\alpha [\pi(b)]\pi \cdot s} (\overset{\wedge}{\approx}_\alpha \mathbf{ab})$$

□

Example 3.8. Notice that $(a\ c) \wedge X \vdash (a\ b) \wedge (b\ c) \cdot X$, for $(a\ c) \wedge X \vdash (a\ c) \wedge X$ and $(a\ b)^{(b\ c)} = (a\ c)$. Using the Equivariance Lemma, with $\rho = (b\ c)$, we deduce:

$$(a\ c) \wedge X \vdash (a\ b)^{(b\ c)} \wedge X \Leftrightarrow (a\ c) \wedge X \vdash (a\ b) \wedge (b\ c) \cdot X \quad (3.1)$$

Proposition 3.9 (Strengthening for \wedge).

If $\Upsilon, \pi \wedge X \vdash \mathcal{V}\bar{c}. \pi' \wedge s$ and $\text{supp}(\pi) \subseteq \text{supp}(\text{perm}(\Upsilon|_X))$ or $X \notin \text{Var}(s)$ then $\Upsilon \vdash \mathcal{V}\bar{c}. \pi' \wedge s$.

Proof. By induction on the structure of s .

- $s = a$

In this case, $\Upsilon, \pi \wedge X \vdash \mathcal{V}\bar{c}. \pi' \wedge a$. By rule $(\wedge \mathbf{a})$,

$$\Upsilon, \pi \wedge X \vdash \mathcal{V}\bar{c}. \pi' \wedge a \Leftrightarrow \pi'(a) = a \Leftrightarrow \Upsilon \vdash \mathcal{V}\bar{c}. \pi' \wedge a$$

- $s = \pi_1 \cdot X$ (the case $s = \pi_1 \cdot Y$ is trivial).

In this case, $\Upsilon, \pi \wedge X \vdash \mathcal{V}\bar{c}. \pi' \wedge \pi_1 \cdot X$.

$$\begin{aligned} \Upsilon, \pi \wedge X \vdash \mathcal{V}\bar{c}. \pi' \wedge \pi_1 \cdot X &\Leftrightarrow \text{supp}((\pi')^{\pi_1^{-1}}) \setminus \{\bar{c}\} \subseteq \text{supp}(\text{perm}(\Upsilon|_X) \cup \{\pi\}) \quad (\text{rule } (\wedge \mathbf{var})) \\ &\Leftrightarrow \text{supp}((\pi')^{\pi_1^{-1}}) \setminus \{\bar{c}\} \subseteq \text{supp}(\text{perm}(\Upsilon|_X)) \quad (\text{by assumption}) \\ &\Leftrightarrow \Upsilon \vdash \mathcal{V}\bar{c}. \pi' \wedge \pi_1 \cdot X \quad (\text{by rule } (\wedge \mathbf{var})) \end{aligned}$$

- $s = [a]s'$

In this case, $\Upsilon, \pi \wedge X \vdash \mathcal{V}\bar{c}. \pi' \wedge [a]s'$, $\text{supp}(\pi) \subseteq \text{supp}(\text{perm}(\Upsilon|_X))$ or $X \notin \text{Var}(s')$.

Then there exists a derivation

$$\frac{\mathcal{D}}{\frac{\Upsilon, \pi \wedge X \vdash \mathcal{V}\bar{c}, c_1. \pi' \wedge (a\ c_1) \cdot s'}{\Upsilon, \pi \wedge X \vdash \mathcal{V}\bar{c}. \pi' \wedge [a]s'} (\wedge \mathbf{abs})}$$

By induction hypothesis, there exists a derivation \mathcal{D}' such that

$$\frac{\mathcal{D}'}{\frac{\Upsilon \vdash \mathcal{V}\bar{c}, c_1. \pi' \wedge (a\ c_1) \cdot s'}{\Upsilon \vdash \pi' \wedge [a]s'} (\wedge \mathbf{abs})}$$

- $s = fs'$ or $s = (s_1, \dots, s_n)$

These cases follow easily by induction hypothesis. \square

Proposition 3.10 (Strengthening for $\overset{\wedge}{\approx}_\alpha$).

If $\Upsilon, \pi \wedge X \vdash \mathcal{V}\bar{c}. s \overset{\wedge}{\approx}_\alpha t$ and $\text{supp}(\pi) \subseteq \text{supp}(\text{perm}(\Upsilon|_X))$ or $X \notin \text{Var}(s, t)$, then $\Upsilon \vdash \mathcal{V}\bar{c}. s \overset{\wedge}{\approx}_\alpha t$.

Proof. The proof is similar to that of Proposition 3.9 and omitted. \square

The following auxiliary lemma permits to deduce a fixed-point constraint for a permutation π and term t from constraints involving t . Intuitively, it states that we can deduce that π fixes t if the support of π is contained in permutations that fix t .

Lemma 3.11. Let I be a set of indices. If $\Upsilon \vdash \mathcal{V}\bar{c}_i. \pi_i \wedge t$ for all $i \in I$, and π is a permutation such that $\text{supp}(\pi) \setminus \{\bar{c}\} \subseteq \bigcup_{i \in I} \text{supp}(\pi_i) \setminus \bigcup_{i \in I} \bar{c}_i$ then $\Upsilon \vdash \mathcal{V}\bar{c}. \pi \wedge t$.

Proof. By induction on t .

- If $t = a$ then by assumption $\pi_i(a) = a$ for all $i \in I$ since $\Upsilon \vdash \mathcal{V}\bar{c}_i.\pi_i \lambda a$ (rule $(\lambda \mathbf{a})$). Then, $\pi(a) = a$ by the assumption on the support of π . The result follows by rule $(\lambda \mathbf{a})$.
- If $t = \pi' \cdot X$ then $\text{supp}(\pi_i^{\pi'^{-1}}) \setminus \{\bar{c}_i\} \subseteq \text{supp}(\text{perm}(\Upsilon|_X))$ because $\Upsilon \vdash \mathcal{V}\bar{c}_i.\pi_i \lambda t$ for all $i \in I$ (rule $(\lambda \mathbf{var})$).
 Since $\text{supp}(\pi) \setminus \{\bar{c}\} \subseteq \bigcup_{i \in I} \text{supp}(\pi_i) \setminus \bigcup_{i \in I} \bar{c}_i$, also $\text{supp}(\pi^{\pi'^{-1}}) \setminus \{\bar{c}\} \subseteq \bigcup_{i \in I} \text{supp}(\pi_i^{\pi'^{-1}}) \setminus \bigcup_{i \in I} \bar{c}_i$. The result follows by rule $(\lambda \mathbf{var})$.
- The other cases follow directly by induction. \square

The following correctness property states that λ is indeed the fixed-point relation.

Theorem 3.12 (Correctness). *Let Υ , π and t be a fixed-point context, a permutation and a nominal term, respectively. $\Upsilon \vdash \mathcal{V}\bar{c}.\pi \lambda t$ iff $\Upsilon \vdash \mathcal{V}\bar{c}.\pi \cdot t \stackrel{\lambda}{\approx}_\alpha t$.*

Proof. In both directions the proof is by induction on the structure of the term t , distinguishing cases according to the last rule applied in the derivation.

(\Rightarrow) The interesting cases correspond to variables and abstractions; the other cases follow directly by the induction hypothesis.

(1) The last rule is $(\lambda \mathbf{var})$.

In this case, $t = \pi' \cdot X$ and $\text{supp}((\pi')^{-1} \circ \pi \circ \pi') \setminus \{\bar{c}\} \subseteq \text{supp}(\text{perm}(\Upsilon|_X))$, and therefore $\Upsilon \vdash \mathcal{V}\bar{c}.\pi \cdot (\pi' \cdot X) \stackrel{\lambda}{\approx}_\alpha \pi' \cdot X$, via rule $(\stackrel{\lambda}{\approx}_\alpha \mathbf{var})$.

(2) The last rule is $(\lambda \mathbf{abs})$. In this case, $t = [a]t'$ and by assumption there is a derivation of the form:

$$\frac{\Pi}{\frac{\Upsilon \vdash \mathcal{V}\bar{c}, c_1. \pi \lambda (a \ c_1) \cdot t'}{\Upsilon \vdash \mathcal{V}\bar{c}.\pi \lambda [a]t'}}$$

We need to prove that $\Upsilon \vdash \mathcal{V}\bar{c}. [\pi(a)]\pi \cdot t' \stackrel{\lambda}{\approx}_\alpha [a]t'$, that is, $\Upsilon \vdash \mathcal{V}\bar{c}.\pi \cdot t' \stackrel{\lambda}{\approx}_\alpha (\pi(a) \ a) \cdot t'$ and also $\Upsilon \vdash \mathcal{V}\bar{c}, c_1. (\pi(a) \ c_1) \lambda t'$.

By IH, there exists a proof Π' for $\Upsilon \vdash \mathcal{V}\bar{c}, c_1. \pi \cdot ((a \ c_1) \cdot t') \stackrel{\lambda}{\approx}_\alpha (a \ c_1) \cdot t'$. The following equivalence holds, since c_1 is newly quantified:

$$\Upsilon \vdash \mathcal{V}\bar{c}, c_1. \pi \cdot ((a \ c_1) \cdot t') \stackrel{\lambda}{\approx}_\alpha (a \ c_1) \cdot t' \iff \Upsilon \vdash \mathcal{V}\bar{c}, c_1. (\pi(a) \ c_1) \cdot (\pi \cdot t') \stackrel{\lambda}{\approx}_\alpha (a \ c_1) \cdot t' \quad (3.2)$$

Also, $\Upsilon \vdash \mathcal{V}\bar{c}, c_1. (\pi \cdot t') \stackrel{\lambda}{\approx}_\alpha (\pi(a) \ c_1) \cdot ((a \ c_1) \cdot t')$ by Equivariance. Note that $(\pi(a) \ c_1) \cdot ((a \ c_1) \cdot t') = (\pi(a) \ a) \cdot t'$, hence $\Upsilon \vdash \mathcal{V}\bar{c}.\pi \cdot t' \stackrel{\lambda}{\approx}_\alpha (\pi(a) \ a) \cdot t'$ as required.

Since $\Upsilon \vdash \mathcal{V}\bar{c}, c_1. \pi \lambda (a \ c_1) \cdot t'$, we also have $\Upsilon \vdash \mathcal{V}\bar{c}, c_1. \pi^{(a \ c_1)} \lambda t'$ by equivariance and we deduce $\Upsilon \vdash \mathcal{V}\bar{c}, c_1. (\pi(a) \ c_1) \lambda t'$ as required.

(\Leftarrow) Similarly, we proceed by induction on the derivation of $\Upsilon \vdash \mathcal{V}\bar{c}.\pi \cdot t \stackrel{\lambda}{\approx}_\alpha t$. We show the interesting cases, the others follow directly by induction.

(1) The rule is $(\stackrel{\lambda}{\approx}_\alpha \mathbf{var})$.

In this case $t = \pi_1 \cdot X$ and there exists a proof of $\Upsilon \vdash \mathcal{V}\bar{c}.\pi \cdot (\pi_1 \cdot X) \stackrel{\lambda}{\approx}_\alpha \pi_1 \cdot X$. Therefore, $\text{supp}(\pi_1^{-1} \circ \pi \circ \pi_1) \setminus \{\bar{c}\} \subseteq \text{supp}(\text{perm}(\Upsilon|_X))$, and one can conclude that $\Upsilon \vdash \mathcal{V}\bar{c}.\pi \lambda \pi_1 \cdot X$, via application of rule $(\lambda \mathbf{var})$.

(2) The rule is $(\stackrel{\lambda}{\approx}_\alpha \mathbf{ab})$.

In this case, $t = [a]t'$ and since $\Upsilon \vdash \mathcal{V}\bar{c}. [\pi(a)]\pi \cdot t' \stackrel{\lambda}{\approx}_\alpha [a]t'$, we know (Inversion):

$$\Upsilon \vdash \mathcal{V}\bar{c}.\pi \cdot t' \stackrel{\lambda}{\approx}_\alpha (\pi(a) \ a) \cdot t' \text{ and } \Upsilon \vdash \mathcal{V}\bar{c}, c_1. (\pi(a) \ c_1) \lambda t'.$$

By Equivariance $\Upsilon \vdash \mathcal{W}\bar{c}.((\pi(a) a) \circ \pi) \cdot t' \stackrel{\lambda}{\approx}_\alpha t'$ and by induction $\Upsilon \vdash \mathcal{W}\bar{c}.((\pi(a) a) \circ \pi) \lambda t'$.

From $\Upsilon \vdash \mathcal{W}\bar{c}, c_1.(\pi(a) c_1) \lambda t'$ and $\Upsilon \vdash \mathcal{W}\bar{c}.((\pi(a) a) \circ \pi) \lambda t'$ we deduce $\Upsilon \vdash \mathcal{W}\bar{c}, c_1.\pi^{(a c_1)} \lambda t'$ by Lemma 3.11, since $\text{supp}(\pi^{(a c_1)}) \setminus \{\bar{c}, c_1\} \subseteq (\text{supp}((\pi(a) a) \circ \pi) \setminus \{\bar{c}\}) \cup (\text{supp}((\pi(a) c_1)) \setminus \{c_1\})$. By Equivariance, $\Upsilon \vdash \mathcal{W}\bar{c}, c_1.\pi \lambda (a c_1) \cdot t'$, and therefore $\Upsilon \vdash \mathcal{W}\bar{c}. \pi \lambda [a]t'$ by $(\lambda \mathbf{ab})$, as required. \square

Recall that $\Upsilon\sigma$ denotes the set of fixed-point constraints obtained by applying the substitution σ to the constraints in Υ (see Section 3.1). Below we abbreviate $\Upsilon \vdash \pi_1 \lambda t_1, \dots, \Upsilon \vdash \pi_n \lambda t_n$ as $\Upsilon \vdash \pi_1 \lambda t_1, \dots, \pi_n \lambda t_n$. Thus, $\Upsilon \vdash \Upsilon'\sigma$ means that each of the constraints in $\Upsilon'\sigma$ is derivable from Υ .

Proposition 3.13 (Preservation under Substitution). *Suppose that $\Upsilon \vdash \Upsilon'\sigma$. Then,*

- (1) $\Upsilon' \vdash \mathcal{W}\bar{c}. \pi \lambda s \implies \Upsilon \vdash \mathcal{W}\bar{c}. \pi \lambda s\sigma$.
- (2) $\Upsilon' \vdash \mathcal{W}\bar{c}. s \stackrel{\lambda}{\approx}_\alpha t \implies \Upsilon \vdash \mathcal{W}\bar{c}. s\sigma \stackrel{\lambda}{\approx}_\alpha t\sigma$.

Proof. By induction on the rules in Figures 3 and 4.

- (1) We distinguish cases depending on the last rule applied in the derivation of $\Upsilon' \vdash \mathcal{W}\bar{c}.\pi \lambda s$.

- (a) The rule is $(\lambda \mathbf{a})$.

In this case, $s = a$ and

$$\frac{\pi(a) = a}{\Upsilon' \vdash \mathcal{W}\bar{c}.\pi \lambda a} (\lambda \mathbf{a})$$

The result follows trivially, since $a\sigma = a$ for any substitution σ .

- (b) The rule is $(\lambda \mathbf{f})$

In this case, $s = fs'$ and there exists a proof Π' of

$$\frac{\frac{\Pi'}{\Upsilon' \vdash \mathcal{W}\bar{c}.\pi \lambda s'}}{\Upsilon' \vdash \mathcal{W}\bar{c}.\pi \lambda fs'} (\lambda \mathbf{f})$$

By induction hypothesis, there exists a proof Π'' such that

$$\frac{\frac{\Pi''}{\Upsilon \vdash \mathcal{W}\bar{c}.\pi \lambda s'\sigma}}{\Upsilon \vdash \mathcal{W}\bar{c}.\pi \lambda (fs')\sigma} (\lambda \mathbf{f})$$

- (c) The rule is $(\lambda \mathbf{abs})$

In this case $s = [a]s'$ and there exists a proof Π' of the form

$$\frac{\frac{\Pi'}{\Upsilon' \vdash \mathcal{W}\bar{c}, c_1.\pi \lambda (a c_1) \cdot s'}}{\Upsilon' \vdash \mathcal{W}\bar{c}.\pi \lambda [a]s'} (\lambda \mathbf{abs})$$

By induction hypothesis, there exists a proof Π'' of the form

$$\frac{\frac{\Pi''}{\Upsilon \vdash \mathcal{W}\bar{c}, c_1.\pi \lambda ((a c_1) \cdot s')\sigma}}{\Upsilon \vdash \mathcal{W}\bar{c}, c_1.\pi \lambda (a c_1) \cdot (s'\sigma)} (\lambda \mathbf{abs})$$

Since $((a c_1) \cdot s')\sigma = (a c_1) \cdot (s'\sigma)$ and $[a](s'\sigma) = ([a]s')\sigma$, it follows that

$$\frac{\frac{\Pi''}{\Upsilon \vdash \mathcal{W}\bar{c}, c_1.\pi \lambda ((a c_1) \cdot s')\sigma}}{\Upsilon \vdash \mathcal{W}\bar{c}, c_1.\pi \lambda (a c_1) \cdot (s'\sigma)} (\lambda \mathbf{abs})$$

- (d) The rule is $(\lambda \mathbf{tuple})$

This case is analogous to the case for $(\lambda \mathbf{f})$, and follows directly by IH.

(e) The rule is $(\lambda \mathbf{var})$

In this case, $s = \rho \cdot X$ and $\Upsilon' \vdash \mathcal{V}\bar{c}.\pi \lambda \rho \cdot X$ holds, that is, $\mathbf{supp}(\pi^{\rho^{-1}}) \setminus \{\bar{c}\} \subseteq \mathbf{supp}(\mathbf{perm}(\Upsilon'|_X))$.

From $\Upsilon \vdash \Upsilon'\sigma$, it follows that, $\Upsilon \vdash \pi_1 \lambda X\sigma$, for all $\pi_1 \lambda X \in \Upsilon'$.

Therefore, $\Upsilon \vdash \mathcal{V}\bar{c}.\pi^{\rho^{-1}} \lambda X\sigma$ by Lemma 3.11, and the result follows by Equivariance.

(2) We proceed by analysing the last rule used in the derivation of $\Upsilon' \vdash \mathcal{V}\bar{c}.s \stackrel{\lambda}{\approx}_{\alpha} t$.

(a) The last rule is $(\stackrel{\lambda}{\approx}_{\alpha} \mathbf{a})$.

This case is trivial.

(b) The last rule is $(\stackrel{\lambda}{\approx}_{\alpha} \mathbf{var})$.

In this case we have $\Upsilon' \vdash \mathcal{V}\bar{c}.\pi \cdot X \stackrel{\lambda}{\approx}_{\alpha} \pi' \cdot X$ and therefore $\mathbf{supp}((\pi')^{-1} \circ \pi) \setminus \{\bar{c}\} \subseteq \mathbf{supp}(\mathbf{perm}(\Upsilon'|_X))$.

From $\Upsilon \vdash \Upsilon'\sigma$, it follows that $\Upsilon \vdash \pi_1 \lambda X\sigma$, for all $\pi_1 \lambda X \in \Upsilon'$.

Therefore, $\Upsilon \vdash \mathcal{V}\bar{c}.(\pi')^{-1} \circ \pi \lambda X\sigma$ by Lemma 3.11, and the result follows by Theorem 3.12 and Equivariance.

(c) The last rule is $(\stackrel{\lambda}{\approx}_{\alpha} \mathbf{f})$ or $(\stackrel{\lambda}{\approx}_{\alpha} [\mathbf{a}])$. These cases follow directly by induction.

(d) The last rule is $(\stackrel{\lambda}{\approx}_{\alpha} \mathbf{ab})$.

In this case, we know $\Upsilon \vdash \mathcal{V}\bar{c}.[a]s \stackrel{\lambda}{\approx}_{\alpha} [b]t$ and therefore (by Inversion) $\Upsilon \vdash \mathcal{V}\bar{c}.s \stackrel{\lambda}{\approx}_{\alpha} (a b) \cdot t$ and $\Upsilon \vdash \mathcal{V}\bar{c}.c_1.(a c_1) \lambda t$.

By induction hypothesis, $\Upsilon \vdash \mathcal{V}\bar{c}.s\sigma \stackrel{\lambda}{\approx}_{\alpha} (a b) \cdot t\sigma$, and by part (1) of this proposition, $\Upsilon \vdash \mathcal{V}\bar{c}.c_1.(a c_1) \lambda t\sigma$. The result then follows using rule $(\stackrel{\lambda}{\approx}_{\alpha} \mathbf{ab})$. \square

Corollary 3.14 (Weakening). *Assume $\Upsilon' \subseteq \Upsilon$.*

(1) $\Upsilon' \vdash \mathcal{V}\bar{c}.\pi \lambda s \implies \Upsilon \vdash \mathcal{V}\bar{c}.\pi \lambda s$.

(2) $\Upsilon' \vdash \mathcal{V}\bar{c}.s \stackrel{\lambda}{\approx}_{\alpha} t \implies \Upsilon \vdash \mathcal{V}\bar{c}.s \stackrel{\lambda}{\approx}_{\alpha} t$.

3.2. Alternative approaches to new name generation. In the previous section we used the “new” quantifier to deal with new names in constraints: judgements involve newly quantified constraints, and in rules $(\lambda \mathbf{ab})$ and $(\stackrel{\lambda}{\approx}_{\alpha} \mathbf{ab})$, when new fresh names are needed, a newly quantified atom is used.

An alternative approach consists of quantifying judgements instead of constraints. More precisely, we can define judgements of the form $\mathcal{V}\bar{c}.(\Upsilon \vdash \pi \lambda t)$ and $\mathcal{V}\bar{c}.(\Upsilon \vdash s \stackrel{\lambda}{\approx}_{\alpha} t)$. Rules to derive this kind of judgements can be easily obtained by adapting the rules in Figures 3 and 4. We show the rules for fixed-point judgements in Figure 5. For the sake of readability, we omit brackets around judgements in the rules and assume the quantifiers have maximal scope.

Yet another approach consists of using a name generator whenever new names are required (see [CU08, CM17]), as done in [AFN18]. In this case, the new quantifier is not needed, and new names are generated dynamically by using an external generator, under the assumption that the generator outputs a new (unused name) whenever needed. For comparison, we recall in Figure 6 the rules given in [AFN18] to derive fixed-point judgements $\Upsilon \vdash \pi \lambda t$. Note that in rule $(\lambda \mathbf{abs})$ the fixed-point context Υ is augmented with constraints $(c_1 c_2) \lambda \mathbf{Var}(t)$. These constraints serve to store the information about the fact that the generated atoms are “new”. The trick is to generate two atoms, even though only one new

$$\begin{array}{c}
\frac{\pi(a) = a}{\mathcal{N}\bar{c}.\Upsilon \vdash \pi \lambda a} (\lambda \mathbf{a}) \qquad \frac{\text{supp}(\pi^{\pi'^{-1}}) \setminus \{\bar{c}\} \subseteq \text{supp}(\text{perm}(\Upsilon|_X))}{\mathcal{N}\bar{c}.\Upsilon \vdash \pi \lambda \pi' \cdot X} (\lambda \mathbf{var}) \\
\frac{\mathcal{N}\bar{c}.\Upsilon \vdash \pi \lambda t}{\mathcal{N}\bar{c}.\Upsilon \vdash \pi \lambda f t} (\lambda \mathbf{f}) \qquad \frac{\mathcal{N}\bar{c}.\Upsilon \vdash \pi \lambda t_1 \quad \dots \quad \mathcal{N}\bar{c}.\Upsilon \vdash \pi \lambda t_n}{\mathcal{N}\bar{c}.\Upsilon \vdash \pi \lambda (t_1, \dots, t_n)} (\lambda \mathbf{tuple}) \\
\frac{\mathcal{N}\bar{c}, c_1.\Upsilon \vdash \pi \lambda (a c_1) \cdot t}{\mathcal{N}\bar{c}.\Upsilon \vdash \pi \lambda [a]t} (\lambda \mathbf{abs})
\end{array}$$

Figure 5: Alternative fixed-point rules.

atom is needed (c_1 , to replace a), in order to be able to express the fact that c_1 is new. This solution is inspired by the link between freshness, the new quantifier and fixed-point equations, $a \# X \Leftrightarrow \mathcal{N}a'.(a a') \cdot X = X$, mentioned at the beginning of Section 3.

$$\begin{array}{c}
\frac{\pi(a) = a}{\Upsilon \vdash \pi \lambda a} (\lambda \mathbf{a}) \qquad \frac{\text{supp}(\pi^{\pi'^{-1}}) \subseteq \text{supp}(\text{perm}(\Upsilon|_X))}{\Upsilon \vdash \pi \lambda \pi' \cdot X} (\lambda \mathbf{var}) \\
\frac{\Upsilon \vdash \pi \lambda t}{\Upsilon \vdash \pi \lambda f t} (\lambda \mathbf{f}) \qquad \frac{\Upsilon \vdash \pi \lambda t_1 \quad \dots \quad \Upsilon \vdash \pi \lambda t_n}{\Upsilon \vdash \pi \lambda (t_1, \dots, t_n)} (\lambda \mathbf{tuple}) \\
\frac{\Upsilon, (c_1 c_2) \lambda \mathbf{Var}(t) \vdash \pi \lambda (a c_1) \cdot t}{\Upsilon \vdash \pi \lambda [a]t} (\lambda \mathbf{abs}) \text{ where } c_1 \text{ and } c_2 \text{ are new names}
\end{array}$$

Figure 6: Fixed-point rules using a name generator.

The rules to derive α -equivalence judgements following this approach are recalled in Figure 7.

$$\begin{array}{c}
\frac{}{\Upsilon \vdash a \overset{\wedge}{\approx}_\alpha a} (\overset{\wedge}{\approx}_\alpha \mathbf{a}) \qquad \frac{\text{supp}((\pi')^{-1} \circ \pi) \subseteq \text{supp}(\text{perm}(\Upsilon|_X))}{\Upsilon \vdash \pi \cdot X \overset{\wedge}{\approx}_\alpha \pi' \cdot X} (\overset{\wedge}{\approx}_\alpha \mathbf{var}) \\
\frac{\Upsilon \vdash t \overset{\wedge}{\approx}_\alpha t'}{\Upsilon \vdash f t \overset{\wedge}{\approx}_\alpha f t'} (\overset{\wedge}{\approx}_\alpha \mathbf{f}) \qquad \frac{\Upsilon \vdash t_1 \overset{\wedge}{\approx}_\alpha t'_1 \quad \dots \quad \Upsilon \vdash t_n \overset{\wedge}{\approx}_\alpha t'_n}{\Upsilon \vdash (t_1, \dots, t_n) \overset{\wedge}{\approx}_\alpha (t'_1, \dots, t'_n)} (\overset{\wedge}{\approx}_\alpha \mathbf{tuple}) \\
\frac{\Upsilon \vdash t \overset{\wedge}{\approx}_\alpha t'}{\Upsilon \vdash [a]t \overset{\wedge}{\approx}_\alpha [a]t'} (\overset{\wedge}{\approx}_\alpha \mathbf{[a]}) \qquad \frac{\Upsilon \vdash s \overset{\wedge}{\approx}_\alpha (a b) \cdot t \quad \Upsilon, (c_1 c_2) \lambda \mathbf{Var}(t) \vdash (a c_1) \lambda t}{\Upsilon \vdash [a]s \overset{\wedge}{\approx}_\alpha [b]t} (\overset{\wedge}{\approx}_\alpha \mathbf{ab})
\end{array}$$

Figure 7: Equality rules using a name generator. In rule $(\overset{\wedge}{\approx}_\alpha \mathbf{ab})$, c_1 and c_2 are new names.

The latter approach has the advantage of having a simpler syntax for judgements (without quantification) but relies on an external name generator, which can be seen as adding a form of state. In other words, in the latter approach rules produce side-effects. Although not as elegant as the approaches using \mathcal{N} , this approach is convenient from an implementation point of view. In the next sections we use this technique to translate primitive freshness constraints to primitive fixed-point constraints, and to specify a unification algorithm.

3.3. From freshness to fixed-point constraints and back again. In this section we show that the α -equivalence relation defined in terms of *freshness constraints*, denoted as \approx_α , is equivalent to $\overset{\wedge}{\approx}_\alpha$, given that a transformation $[-]^\wedge$ from primitive freshness to primitive fixed-point constraints and a transformation $[-]^\#$ from primitive fixed-point to primitive freshness constraints can be defined.

Below we denote by $\mathfrak{F}_\#$ the family of freshness contexts, and by \mathfrak{F}_\wedge the family of fixed-point contexts. The mapping $[-]_\wedge$ associates each primitive freshness constraint with a fixed-point constraint; it extends to freshness contexts in the natural way.

$$\begin{aligned} [-]_\wedge : \quad \mathfrak{F}_\# &\longrightarrow \mathfrak{F}_\wedge \\ a\#X &\mapsto (a\ c_a) \wedge X \text{ where } c_a \text{ is a new name.} \end{aligned}$$

We denote by $[\Delta]_\wedge$ the image of Δ under $[-]_\wedge$.

The mapping $[-]_\#$ associates each primitive fixed-point constraint with a freshness constraint; it extends to fixed-point contexts in the natural way.

$$\begin{aligned} [-]_\# : \quad \mathfrak{F}_\wedge &\longrightarrow \mathfrak{F}_\# \\ \pi \wedge X &\mapsto \text{supp}(\pi)\#X. \end{aligned}$$

We denote by $[\Upsilon]_\#$ the image of Υ under $[-]_\#$.

Below we abbreviate the set of constraints $\{a_1\#t, \dots, a_n\#t \mid a_1, \dots, a_n \in A\}$ as $\overline{A\#t}$.

Using the above-specified translation of primitive constraints, we translate freshness judgements into fixed-point judgements with newly quantified constraints:

$$\Delta \vdash a\#t \text{ is translated as } [\Delta]_\wedge \vdash \mathcal{V}c.(a\ c) \wedge t$$

and vice-versa:

$$\Upsilon \vdash \mathcal{V}\bar{c}.\pi \wedge t \text{ is translated as } [\Upsilon]_\# \vdash \overline{\text{supp}(\pi) \setminus \{\bar{c}\}\#t}.$$

Theorem 3.15.

- (1) $\Delta \vdash a\#t \Leftrightarrow [\Delta]_\wedge \vdash \mathcal{V}c.(a\ c) \wedge t.$
- (2) $\Upsilon \vdash \mathcal{V}\bar{c}.\pi \wedge t \Leftrightarrow [\Upsilon]_\# \vdash \overline{\text{supp}(\pi) \setminus \{\bar{c}\}\#t}.$

Proof. Part (1):

(\Rightarrow) By induction on the derivation of $\Delta \vdash a\#t$ (see the rules in Figure 1). We distinguish cases based on the last rule used in the derivation.

- If the rule is ($\#a$) then t is an atom b (it cannot be a), and the result follows by rule ($\wedge a$).
- If the rule is ($\#\text{var}$) then $t = \pi \cdot X$ and since $\Delta \vdash a\#\pi \cdot X$ we know $\pi^{-1}(a)\#X \in \Delta$ by Inversion. Hence, $(\pi^{-1}(a)\ c_{\pi^{-1}(a)}) \wedge X \in [\Delta]_\wedge$, and therefore $\text{supp}((a\ c)^{\pi^{-1}}) \setminus \{c\} \subseteq \text{supp}(\text{perm}([\Delta]_\wedge|_X))$ (recall that $\pi^{-1}(c) = c$ since c is a new atom). The result then follows by rule ($\wedge\text{var}$).
- The cases for ($\#\text{tuple}$) and ($\#\text{f}$) follow directly by induction.
- If the rule is ($\#[a]$) then $t = [a]s$ and we need to prove $[\Delta]_\wedge \vdash \mathcal{V}c.(a\ c) \wedge [a]s$. By rule ($\wedge\text{abs}$), it suffices to show $[\Delta]_\wedge \vdash \mathcal{V}c, c_1.(a\ c) \wedge (a\ c_1) \cdot s$. By Equivariance, this is equivalent to $[\Delta]_\wedge \vdash \mathcal{V}c, c_1.(a\ c)^{(a\ c_1)} \wedge s$, or equivalently $[\Delta]_\wedge \vdash \mathcal{V}c, c_1.(c_1\ c) \wedge s$, which holds trivially since c and c_1 are new atoms ($(c_1\ c) \cdot s = s$).
- If the rule is ($\#\text{abs}$) then $t = [b]s$. By assumption, $\Delta \vdash a\#[b]s$, hence $\Delta \vdash a\#s$. By induction hypothesis, $[\Delta]_\wedge \vdash \mathcal{V}c.(a\ c) \wedge s$, and by Equivariance, $[\Delta]_\wedge \vdash \mathcal{V}c, c_1.(a\ c) \wedge (b\ c_1) \cdot s$. The result follows by rule ($\wedge\text{abs}$).

(\Leftarrow) By induction on the derivation of $[\Delta]_\wedge \vdash \mathcal{V}c.(a\ c) \wedge t$. We distinguish cases based on the last rule used in the derivation.

- If the rule is $(\lambda \mathbf{a})$ then t is an atom b (it cannot be a), and the result follows by rule $(\# \mathbf{a})$.
- If the rule is $(\lambda \mathbf{var})$ then $t = \pi \cdot X$ and since $[\Delta]_\lambda \vdash \mathcal{V}c.(a\ c) \lambda \pi \cdot X$ we know $\text{supp}((a\ c)^{\pi^{-1}}) \setminus \{c\} \subseteq \text{supp}(\text{perm}([\Delta]_\lambda | X))$ by Inversion. Hence, $\pi^{-1}(a) \in \text{supp}(\text{perm}([\Delta]_\lambda | X))$, and therefore $\pi^{-1}(a) \# X \in \Delta$ by definition of the mapping $[\cdot]_\lambda$. The result then follows by rule $(\# \mathbf{var})$.
- The cases for $(\lambda \mathbf{tuple})$ and $(\lambda \mathbf{f})$ follow directly by induction.
- If the rule is $(\lambda \mathbf{abs})$ then there are two cases, $t = [a]s$ and $t = [b]s$.
If $t = [a]s$ the result follows directly by rule $(\# \mathbf{a})$.
If $t = [b]s$ then by assumption and Inversion, $[\Delta]_\lambda \vdash \mathcal{V}cc_1.(a\ c) \lambda (b\ c_1) \cdot s$. By Equivariance, $[\Delta]_\lambda \vdash \mathcal{V}cc_1.(a\ c) \lambda s$. By induction hypothesis we deduce $\Delta \vdash a \# s$, and the result follows by rule $(\# \mathbf{abs})$.

Part (2):

(\Rightarrow) By induction on the derivation of $\Upsilon \vdash \mathcal{V}\bar{c}.\pi \lambda t$. Again we distinguish cases based on the last rule used in the derivation. The only interesting cases are rule $(\lambda \mathbf{var})$ and $(\lambda \mathbf{abs})$.

- If the last rule applied is $(\lambda \mathbf{var})$ then $t = \pi' \cdot X$. By Inversion, $\text{supp}(\pi^{\pi'^{-1}}) \setminus \{\bar{c}\} \subseteq \text{supp}(\text{perm}(\Upsilon | X))$. Therefore, for any $a \in \text{supp}(\pi) \setminus \{\bar{c}\}$, $\pi'^{-1}(a) \in \text{supp}(\text{perm}(\Upsilon | X))$. By definition of the mapping $[\cdot]_\#$, $\pi'^{-1}(a) \# X \in [\Upsilon]_\#$, and the result follows by rule $(\# \mathbf{var})$.
- If the last rule applied is $(\lambda \mathbf{abs})$ then $t = [a]s$. In this case, we know $\Upsilon \vdash \mathcal{V}c.c_1.\pi \lambda (a\ c_1) \cdot s$. By induction, $[\Upsilon]_\# \vdash \text{supp}(\pi) \setminus \{c, c_1\} \# (a\ c_1) \cdot s$, hence $[\Upsilon]_\# \vdash \text{supp}(\pi) \setminus \{a\} \# s$, and the result follows by rules $(\# \mathbf{a})$ and $(\# \mathbf{abs})$.

(\Leftarrow) By induction on t .

- If t is an atom a then $a \notin \text{supp}(\pi) \setminus \{\bar{c}\}$ (by Inversion, rule $(\# \mathbf{a})$). The result follows by rule $(\lambda \mathbf{a})$.
- If $t = \pi' \cdot X$, then the rule applied in the derivation is rule $(\# \mathbf{var})$. By Inversion, for any atom a in $\text{supp}(\pi) \setminus \{\bar{c}\}$, $\text{supp}(\pi^{\pi'^{-1}}) \subseteq \text{supp}(\text{perm}(\Upsilon | X))$ and the result follows by rule $(\lambda \mathbf{var})$.
- If $t = [a]s$ then for any atom $b \in \text{supp}(\pi) \setminus \{\bar{c}\}$ (b different from a by the permutative convention), we know $[\Upsilon]_\# \vdash b \# s$, since by assumption $[\Upsilon]_\# \vdash b \# [a]s$ and the last rule applied in the derivation must have been rule $(\# \mathbf{abs})$. In particular, for any atom $b \in \text{supp}(\pi) \setminus \{\bar{c}, a, c_1\}$, $[\Upsilon]_\# \vdash b \# s$. Then, by induction hypothesis, $\Upsilon \vdash \mathcal{V}c.c_1.\pi^{(a\ c_1)} \lambda s$ and the result follows by Equivariance and rule $(\lambda \mathbf{abs})$.
- The other cases follow directly by induction. \square

Theorem 3.16. $\overset{\wedge}{\approx}_\alpha$ coincides with \approx_α on ground terms, that is, $\vdash s \approx_\alpha t \iff \vdash s \overset{\wedge}{\approx}_\alpha t$.
More generally,

- (1) $\Delta \vdash s \approx_\alpha t \Rightarrow [\Delta]_\lambda \vdash s \overset{\wedge}{\approx}_\alpha t$.
- (2) $\Upsilon \vdash \mathcal{V}\bar{c}.s \overset{\wedge}{\approx}_\alpha t \Rightarrow [\Upsilon]_\# \vdash \Delta \vdash s \approx_\alpha t$, where $\Delta \vdash \bar{c} \# \mathbf{Var}(s, t)$.

Proof. (1) The first part is proved by induction on the derivation of $\Delta \vdash s \approx_\alpha t$, distinguishing cases according to the last rule applied (see Figure 2). The interesting cases correspond to $(\approx_\alpha \mathbf{var})$ and $(\approx_\alpha \mathbf{ab})$.

- If the last rule applied is $(\approx_\alpha \mathbf{var})$:

$$\frac{\text{ds}(\pi, \pi_1) \# X \subseteq \Delta}{\Delta \vdash \pi \cdot X \approx_\alpha \pi_1 \cdot X} (\approx_\alpha \mathbf{var})$$

We want to show that $[\Delta]_\lambda \vdash \pi \cdot X \overset{\wedge}{\approx}_\alpha \pi_1 \cdot X$. To use rule $(\overset{\wedge}{\approx}_\alpha \mathbf{var})$, we need to show that $\text{supp}(\pi_1^{-1} \circ \pi) \subseteq \text{supp}(\text{perm}([\Delta]_\lambda | X))$. Let $b \in \text{supp}(\pi_1^{-1} \circ \pi)$ and

suppose $b \notin \mathbf{ds}(\pi, \pi_1)$. Then $\pi(b) = \pi_1(b)$ and $\pi_1^{-1}(\pi(b)) = b$, contradiction. Therefore, $b \in \mathbf{ds}(\pi, \pi_1)$ and $(b\ c_b) \lambda X \in [\Delta]_\lambda$ (for c_b a new name), and the result follows.

- If the last rule applied is $(\approx_\alpha \mathbf{ab})$ the result follows directly by induction and Theorem 3.15.

- (2) The second part is proved by induction on the derivation of $\Upsilon \vdash s \overset{\lambda}{\approx}_\alpha t$, distinguishing cases according to the last rule applied. Again the interesting cases correspond to variables and abstractions. The proof follows the lines of the previous part and is omitted. \square

As a corollary, since \approx_α is an equivalence relation [UPG04], we deduce that $\overset{\lambda}{\approx}_\alpha$ is also an equivalence relation.

Theorem 3.17. $\overset{\lambda}{\approx}_\alpha$ is an equivalence relation.

Lemma 3.18 (λ preservation under $\overset{\lambda}{\approx}_\alpha$). If $\Upsilon \vdash s \overset{\lambda}{\approx}_\alpha t$ and $\Upsilon \vdash \mathcal{V}\bar{c}.\pi \lambda s$ then $\Upsilon \vdash \mathcal{V}\bar{c}.\pi \lambda t$.

Proof. Direct consequence of Theorem 3.12, Equivariance and Transitivity. \square

Having proved that $\overset{\lambda}{\approx}_\alpha$ is an equivalence relation, and that λ is correctly defined with respect to $\overset{\lambda}{\approx}_\alpha$, we can use λ to define the support of a non-ground term. We denote the support of a term in context $\Upsilon \vdash t$ as $\mathbf{supp}_\Upsilon(t)$. As indicated in Definition 2.5, the set of atoms in the support of an element of a nominal set can be characterised by using permutation fixed points. In the particular case of terms, the previous results justify the definition of $\mathbf{supp}_\Upsilon(t)$ using fixed-point judgements as follows.

Definition 3.19. Let $\Upsilon \vdash t$ be a term in context. The *support* of t with respect to Υ , $\mathbf{supp}_\Upsilon(t)$, is the smallest set A of atoms such that for any permutation π ,

$$(\forall a \in A, \pi(a) = a) \Rightarrow \Upsilon \vdash \pi \lambda t.$$

As expected, α -equivalent terms have the same support.

Lemma 3.20.

- (1) If $\Upsilon \vdash s \overset{\lambda}{\approx}_\alpha t$ then $\mathbf{supp}_\Upsilon(s) = \mathbf{supp}_\Upsilon(t)$.
- (2) $\Upsilon \vdash \mathcal{V}\bar{c}.\pi \lambda t$ if and only if $\mathbf{supp}(\pi) \setminus \{\bar{c}\} \cap \mathbf{supp}_\Upsilon(t) = \emptyset$.

Proof. Direct consequence of Definition 3.19 and Lemma 3.18. \square

4. NOMINAL UNIFICATION VIA FIXED-POINT CONSTRAINTS

In this section we address the problem of unifying nominal terms. Solutions for unification problems will be represented using fixed-point constraints and substitutions. After defining unification problems, we present a simplification algorithm that computes the most general unifier for a unification problem, provided the problem has a solution (otherwise the algorithm stops indicating that there is no solution). To specify the nominal unification algorithm, in this section we follow the approach to dealing with new atoms that relies on a name generator.

Definition 4.1. A *unification problem* Pr consists of a finite set of equality and fixed-point constraints of the form $s \overset{\lambda?}{\approx}_\alpha t$ and $\pi \lambda? t$, respectively.

Below we call $\pi \lambda^? X$ a **primitive** constraint.

Definition 4.2 (Solution). A *solution* for a problem Pr is a pair of the form $\langle \Phi, \sigma \rangle$ where the following conditions are satisfied:

- (1) $\Phi \vdash \pi \lambda^? t \sigma$, if $\pi \lambda^? t \in \text{Pr}$;
- (2) $\Phi \vdash s \sigma \stackrel{\wedge}{\approx}_\alpha t \sigma$, if $s \stackrel{\wedge}{\approx}_\alpha t \in \text{Pr}$.
- (3) $X \sigma \stackrel{\wedge}{\approx}_\alpha X \sigma \sigma$ for all $X \in \text{Var}(\text{Pr})$ (the substitution is idempotent).

The *solution set* for a problem Pr is denoted by $\mathcal{U}(\text{Pr})$.

Solutions in $\mathcal{U}(\text{Pr})$ can be compared using the following ordering.

Definition 4.3. Let Φ_1, Φ_2 be fixed-point contexts, and σ_1, σ_2 substitutions. Then $\langle \Phi_1, \sigma_1 \rangle \leq \langle \Phi_2, \sigma_2 \rangle$ if there exists some σ' such that

$$\text{for all } X, \quad \Phi_2 \vdash X \sigma_1 \sigma' \stackrel{\wedge}{\approx}_\alpha X \sigma_2 \quad \text{and} \quad \Phi_2 \vdash \Phi_1 \sigma'.$$

Definition 4.4. A **principal** (or **most general**) solution to a problem Pr is a least element of $\mathcal{U}(\text{Pr})$.

We design a unification algorithm via the simplification rules presented in Table 1. These rules act on unification problems Pr by transforming constraints into simpler ones, or instantiating variables in the case of rules $(\approx_\alpha \text{ inst1})$ and $(\approx_\alpha \text{ inst2})$. We call the latter *instantiating rules*. We abbreviate (t_1, \dots, t_n) as $(\tilde{t})_{1..n}$, and for a set S , $\overline{\pi \lambda^? S} = \{\pi \lambda^? X \mid X \in S\}$.

(λat)	$\text{Pr} \uplus \{\pi \lambda^? a\}$	$\implies \text{Pr}$, if $\pi(a) = a$
(λf)	$\text{Pr} \uplus \{\pi \lambda^? ft\}$	$\implies \text{Pr} \cup \{\pi \lambda^? t\}$
(λt)	$\text{Pr} \uplus \{\pi \lambda^? (\tilde{t})_n\}$	$\implies \text{Pr} \cup \{\pi \lambda^? t_1, \dots, \pi \lambda^? t_n\}$
(λabs)	$\text{Pr} \uplus \{\pi \lambda^? [a]t\}$	$\implies \text{Pr} \cup \{\pi \lambda^? (a \ c_1) \cdot t, (c_1 \ c_2) \lambda^? \text{Var}(t)\}$
(λvar)	$\text{Pr} \uplus \{\pi \lambda^? \pi' \cdot X\}$	$\implies \text{Pr} \cup \{\pi^{(\pi')^{-1}} \lambda^? X\}$, if $\pi' \neq Id$
$(\stackrel{\wedge}{\approx}_\alpha a)$	$\text{Pr} \uplus \{a \stackrel{\wedge}{\approx}_\alpha a\}$	$\implies \text{Pr}$
$(\stackrel{\wedge}{\approx}_\alpha f)$	$\text{Pr} \uplus \{f t \stackrel{\wedge}{\approx}_\alpha f t'\}$	$\implies \text{Pr} \cup \{t \approx_\alpha^? t'\}$
$(\stackrel{\wedge}{\approx}_\alpha t)$	$\text{Pr} \uplus \{(\tilde{t})_n \approx_\alpha^? (\tilde{t}')_n\}$	$\implies \text{Pr} \cup \{t_1 \stackrel{\wedge}{\approx}_\alpha t'_1, \dots, t_n \stackrel{\wedge}{\approx}_\alpha t'_n\}$
$(\stackrel{\wedge}{\approx}_\alpha abs1)$	$\text{Pr} \uplus \{[a]t \stackrel{\wedge}{\approx}_\alpha [a]t'\}$	$\implies \text{Pr} \cup \{t \stackrel{\wedge}{\approx}_\alpha t'\}$
$(\stackrel{\wedge}{\approx}_\alpha abs2)$	$\text{Pr} \uplus \{[a]t \stackrel{\wedge}{\approx}_\alpha [b]s\}$	$\implies \text{Pr} \cup \{t \stackrel{\wedge}{\approx}_\alpha (a \ b) \cdot s, (a \ c_1) \lambda^? s, \overline{(c_1 \ c_2) \lambda^? \text{Var}(s)}\}$
$(\stackrel{\wedge}{\approx}_\alpha var)$	$\text{Pr} \uplus \{\pi \cdot X \stackrel{\wedge}{\approx}_\alpha \pi' \cdot X\}$	$\implies \text{Pr} \cup \{(\pi')^{-1} \circ \pi \lambda^? X\}$
$(\stackrel{\wedge}{\approx}_\alpha \text{ inst1})$	$\text{Pr} \uplus \{\pi \cdot X \stackrel{\wedge}{\approx}_\alpha t\}$	$\xRightarrow{[X \mapsto \pi^{-1}.t]} \text{Pr}\{X \mapsto \pi^{-1}.t\}$, if $X \notin \text{Var}(t)$
$(\stackrel{\wedge}{\approx}_\alpha \text{ inst2})$	$\text{Pr} \uplus \{t \stackrel{\wedge}{\approx}_\alpha \pi \cdot X\}$	$\xRightarrow{[X \mapsto \pi^{-1}.t]} \text{Pr}\{X \mapsto \pi^{-1}.t\}$, if $X \notin \text{Var}(t)$

Table 1: Simplification Rules. In (λabs) and $(\stackrel{\wedge}{\approx}_\alpha abs2)$, c_1 and c_2 are new names.

We write $\text{Pr} \implies \text{Pr}'$ when Pr' is obtained from Pr by applying a simplification rule from Table 1, and we write \implies^* for the reflexive and transitive closure of \implies .

Lemma 4.5 (Termination). *There is no infinite chain of reductions \implies starting from a problem Pr .*

Proof. Termination of the simplification rules follows directly from the fact that the following measure of the size of \mathbf{Pr} is strictly decreasing:

$[\mathbf{Pr}] = (n_1, M)$ where n_1 is the number of different variables used in \mathbf{Pr} , and M is the multiset of sizes of equality constraints and non-primitive fixed-point constraints occurring in \mathbf{Pr} . To compare $[\mathbf{Pr}]$ and $[\mathbf{Pr}']$ we use the lexicographic combination of the usual order on natural numbers, $>$, and its multiset extension. We denote the order by $>_{lex}$. Thus, $[\mathbf{Pr}] >_{lex} [\mathbf{Pr}']$ if \mathbf{Pr}' has less variables than \mathbf{Pr} , or if it has the same number of variables but smaller equality constraints and non-primitive fixed-point constraints.

Each simplification step $\mathbf{Pr} \Longrightarrow \mathbf{Pr}'$ either eliminates one variable (when an instantiation rule is used) and therefore decreases the first component of the interpretation, or leaves the first component unchanged but replaces a constraint with smaller ones and/or primitive ones (when a non-instantiating rule is used). Therefore, $\mathbf{Pr} \Longrightarrow \mathbf{Pr}'$ implies $[\mathbf{Pr}] >_{lex} [\mathbf{Pr}']$. Hence, it is not possible to have an infinite descending chain. \square

If $\mathbf{Pr} \Longrightarrow^* \mathbf{Pr}'$ and \mathbf{Pr}' is irreducible, we say that \mathbf{Pr}' is a normal form. We will show next that if instantiation rules are not used, each problem \mathbf{Pr} has a unique normal form. Indeed unicity of normal forms for non-instantiating rules is a consequence of the following property.

Lemma 4.6 (Confluence). *The relation \Longrightarrow defined by the rules in Table 1 without $(\approx_\alpha \text{inst1})$ and $(\approx_\alpha \text{inst2})$ is confluent.*

Proof. Confluence follows from the fact that the rules have no critical pairs (there are only trivial overlaps) and are terminating (by Newman's lemma). \square

If instantiating rules are used, normal forms are not necessarily unique. For example, the problem $(a \ b) \lambda^? X, X \overset{\lambda^?}{\approx}_\alpha Y$ has two normal forms:

$$(a \ b) \lambda^? X, X \overset{\lambda^?}{\approx}_\alpha Y \Longrightarrow (a \ b) \lambda^? Y$$

$$(a \ b) \lambda^? X, X \overset{\lambda^?}{\approx}_\alpha Y \Longrightarrow (a \ b) \lambda^? X$$

However reductions do always terminate with some normal form, and all the normal forms of a problem are equivalent in a natural and useful sense (see Definition 4.9 and Theorem 4.12). We will use the notation $\langle \mathbf{Pr} \rangle_{\text{nf}}$ to refer to any normal form of \mathbf{Pr} .

We say that an equality constraint $s \overset{\lambda^?}{\approx}_\alpha t$ is *reduced* when one of the following holds:

- (1) $s = a$ and $t = b$ are distinct atoms;
- (2) s and t are headed with different function symbols, that is, $s = f \ s'$ and $t = g \ t'$;
- (3) s and t have different term constructors, that is, $s = [a]s'$ and $t = f \ t'$, for some term former f , or $s = \pi \cdot X$ and $t = a$, etc.

A fixed-point constraint $\pi \lambda^? s$ is *reduced* when it is of the form $\pi \lambda^? a$ and $\pi(a) \neq a$, or $\pi \lambda^? X$; the former is *inconsistent* whereas the latter is *consistent*.

Example 4.7. For $\text{Pr} = [a]f(X, a) \stackrel{\lambda?}{\approx}_\alpha [b]f((b\ c) \cdot W, (a\ c) \cdot Y)$, we obtain the following derivation chain using the rules in Table 1:

$$\begin{aligned}
[a]f(X, a) \stackrel{\lambda?}{\approx}_\alpha [b]f((b\ c) \cdot W, (a\ c) \cdot Y) &\Longrightarrow \left\{ \begin{array}{l} f(X, a) \stackrel{\lambda?}{\approx}_\alpha f((a\ b) \circ (b\ c) \cdot W, (a\ b) \circ (a\ c) \cdot Y), \\ (a\ c_1) \lambda? f((b\ c) \cdot W, (a\ c) \cdot Y), \\ (c_1\ c_2) \lambda? W, (c_1\ c_2) \lambda? Y \end{array} \right\} \\
&\xrightarrow{*} \left\{ \begin{array}{l} X \stackrel{\lambda?}{\approx}_\alpha (a\ b) \circ (b\ c) \cdot W, a \stackrel{\lambda?}{\approx}_\alpha (a\ b) \circ (a\ c) \cdot Y, \\ (a\ c_1) \lambda? (b\ c) \cdot W, (a\ c_1) \lambda? (a\ c) \cdot Y, (c_1\ c_2) \lambda? W, (c_1\ c_2) \lambda? Y \end{array} \right\} \\
&\xrightarrow{[Y \mapsto b]} \left\{ \begin{array}{l} X \stackrel{\lambda?}{\approx}_\alpha (a\ b) \circ (b\ c) \cdot W, (a\ c_1)^{(b\ c)} \lambda? W, (a\ c_1) \lambda? b, (c_1\ c_2) \lambda? W, (c_1\ c_2) \lambda? b \end{array} \right\} \\
&\xrightarrow{*} \left\{ \begin{array}{l} X \stackrel{\lambda?}{\approx}_\alpha (a\ b) \circ (b\ c) \cdot W, (a\ c_1) \lambda? W, (c_1\ c_2) \lambda? W \end{array} \right\} \\
&\xrightarrow{[X \mapsto (a\ b) \circ (b\ c) \cdot W]} \left\{ (a\ c_1) \lambda? W, (c_1\ c_2) \lambda? W \right\} = \langle \text{Pr} \rangle_{\text{nf}}.
\end{aligned}$$

Definition 4.8 (Characterisation of normal forms). Let Pr be a problem such that $\langle \text{Pr} \rangle_{\text{nf}} = \text{Pr}'$. We say that Pr' is *reduced* when it consists of reduced constraints, and *successful* when $\text{Pr}' = \emptyset$ or contains only consistent reduced fixed-point constraints; otherwise, $\langle \text{Pr} \rangle_{\text{nf}}$ *fails*.

The simplification rules (Table 1) specify a *unification algorithm*: we apply the simplification rules in a problem Pr until we reach a normal form $\langle \text{Pr} \rangle_{\text{nf}}$.

Definition 4.9 (Computed Solutions). If $\langle \text{Pr} \rangle_{\text{nf}}$ *fails* or contains reduced equational constraints, we say that Pr is *unsolvable*; otherwise, $\langle \text{Pr} \rangle_{\text{nf}}$ is *solvable* and its solution, denoted $\langle \text{Pr} \rangle_{\text{sol}}$, consists of the composition σ of substitutions applied through the simplification steps and the fixed-point context $\Phi = \{\pi \lambda? X \mid \pi \lambda? X \in \langle \text{Pr} \rangle_{\text{nf}}\}$.

Example 4.10 (Continuing example 4.7). Notice that $\langle \Psi, \sigma \rangle$, where $\Psi = \{(a\ c_1) \lambda? W, (c_1\ c_2) \lambda? W\}$ and $\sigma = \{Y \mapsto b, X \mapsto (a\ b) \circ (b\ c) \cdot W\}$, is a solution for Pr .

We will now show that every solvable unification problem has a principal solution, computed by the unification algorithm, such that any other solution can be obtained as an instance of the principal one.

We start by proving that the non-instantiating rules preserve the set of solutions.

Lemma 4.11 (Correctness of Non-Instantiating rules). *Let Pr be a unification problem such that $\text{Pr} \Longrightarrow^* \text{Pr}'$ without using instantiating rules ($\stackrel{\lambda?}{\approx}_\alpha \text{inst1}$) and ($\stackrel{\lambda?}{\approx}_\alpha \text{inst2}$) then*

- (1) $\mathcal{U}(\text{Pr}) = \mathcal{U}(\text{Pr}')$, and
- (2) if Pr' contains equational or inconsistent reduced fixed-point constraints then $\mathcal{U}(\text{Pr}) = \emptyset$.

Proof. The proof is by induction on the length of the derivation $\text{Pr} \xrightarrow{n} \text{Pr}'$.

Base Case. $n = 0$. Then $\text{Pr} = \text{Pr}'$ and the result is trivial.

Induction Step. Suppose, $n > 0$ and consider the reduction chain

$$\text{Pr} = \text{Pr}_1 \Longrightarrow \dots \Longrightarrow \text{Pr}_{n-1} \Longrightarrow \text{Pr}_n = \text{Pr}'.$$

The proof follows by case analysis on the last rule applied in Pr_{n-1} .

- (1) The rule is (λat).

In this case, $\text{Pr}_{n-1} = \text{Pr}'_{n-1} \uplus \{\pi \lambda? a\} \Longrightarrow \text{Pr}'_{n-1} = \text{Pr}_n$, and $\pi(a) = a$.

Let $\langle \Psi, \sigma \rangle \in \mathcal{U}(\text{Pr}_{n-1})$, then

- (a) $\Psi \vdash \pi' \lambda? t\sigma$, for all $\pi' \lambda? t \in \text{Pr}'_{n-1}$

- (b) $\Psi \vdash t\sigma \overset{\lambda}{\approx}_\alpha s\sigma$, for all $t \overset{\lambda}{\approx}_\alpha s \in \mathbf{Pr}'_{n-1}$;
- (c) $X\sigma = X\sigma\sigma$, for all $X \in \mathbf{Var}(\mathbf{Pr}'_{n-1})$.

Therefore, $\langle \Psi, \sigma \rangle \in \mathcal{U}(\mathbf{Pr}_n)$ and $\mathcal{U}(\mathbf{Pr}_{n-1}) \subseteq \mathcal{U}(\mathbf{Pr}_n)$. The other inclusion is trivial.

(2) The rule is (λvar) .

In this case, $\mathbf{Pr}_{n-1} = \mathbf{Pr}'_{n-1} \uplus \{\pi \lambda^? \pi' \cdot X\} \implies \mathbf{Pr}'_{n-1} \cup \{\pi(\pi')^{-1} \lambda^? X\} = \mathbf{Pr}_n$, and $\pi' \neq Id$.

Let $\langle \Psi, \sigma \rangle \in \mathcal{U}(\mathbf{Pr}_{n-1})$, then

- (a) $\Psi \vdash \pi' \lambda t\sigma$, for all $\pi' \lambda^? t \in \mathbf{Pr}'_{n-1}$, and $\Psi \vdash \pi \lambda \pi' \cdot X\sigma$.
- (b) $\Psi \vdash t\sigma \overset{\lambda}{\approx}_\alpha s\sigma$, for all $t \overset{\lambda}{\approx}_\alpha s \in \mathbf{Pr}'_{n-1}$;
- (c) $X\sigma = X\sigma\sigma$, for all $X \in \mathbf{Var}(\mathbf{Pr}'_{n-1})$.

Notice that

$$\begin{aligned} \Psi \vdash \pi \lambda \pi' \cdot X\sigma &\implies \Psi \vdash \pi \cdot (\pi' \cdot X\sigma) \overset{\lambda}{\approx}_\alpha (\pi' \cdot X\sigma), \text{ hence} \\ &\Psi \vdash (\pi')^{-1} \circ \pi \circ \pi' \cdot (X\sigma) \overset{\lambda}{\approx}_\alpha X\sigma \text{ by Equivariance (Lemma 3.7)} \\ &\implies \Psi \vdash \pi(\pi')^{-1} \lambda X\sigma. \end{aligned}$$

Therefore, $\langle \Psi, \sigma \rangle \in \mathcal{U}(\mathbf{Pr}_n)$ and $\mathcal{U}(\mathbf{Pr}_{n-1}) \subseteq \mathcal{U}(\mathbf{Pr}_n)$. The other inclusion is similar.

(3) The rule is λabs . Then

$$\mathbf{Pr}_{n-1} = \mathbf{Pr}' \uplus \{\pi \lambda^? [a]s\} \implies \mathbf{Pr}' \cup \{\pi \lambda^? (a \ c_1).s, \overline{(c_1 \ c_2) \lambda \mathbf{Var}(s)}\} = \mathbf{Pr}_n.$$

where c_1 and c_2 are new names not occurring anywhere in the problem.

Let $\langle \Psi, \sigma \rangle \in \mathcal{U}(\mathbf{Pr}_{n-1})$ be a solution for \mathbf{Pr}_{n-1} :

- (a) $\Psi \vdash \pi' \lambda t\sigma$, for all $\pi' \lambda^? t \in \mathbf{Pr}'$, and $\Psi \vdash \pi \lambda ([a]s)\sigma$.
- (b) $\Psi \vdash t\sigma \overset{\lambda}{\approx}_\alpha s\sigma$, for all $t \overset{\lambda}{\approx}_\alpha s \in \mathbf{Pr}'$.

Since $\Psi \vdash \pi \lambda ([a]s)\sigma$ and $([a]s)\sigma = [a]s\sigma$, it follows that $\Psi \vdash \pi \lambda [a](s\sigma)$. From inversion and rule $(\lambda \mathbf{abs})$ (see Figure 6), this implies that there exists a proof for $\Psi, \overline{(c_1 \ c_2) \lambda \mathbf{Var}(s\sigma)} \vdash \pi \lambda (a \ c_1).s\sigma$.

Notice that we can always choose c_1 and c_2 such that $\mathbf{supp}((c_1 \ c_2)) \cap \mathbf{supp}(s\sigma) = \emptyset$, from Lemma 3.20, it follows that $\Psi \vdash (c_1 \ c_2) \lambda s\sigma$. Since $\Psi, \overline{(c_1 \ c_2) \lambda \mathbf{Var}(s\sigma)} \vdash \pi \lambda (a \ c_1).s\sigma$, it follows that $\Psi \vdash \pi \lambda (a \ c_1).s\sigma$, by Proposition 3.13. The other inclusion is similar.

The cases corresponding to the other non-instantiating simplification rules are similar to the above and omitted. \square

We now show that the result of the algorithm is a principal solution.

Theorem 4.12. *Let \mathbf{Pr} be a unification problem, and suppose $\langle \mathbf{Pr} \rangle_{sol} = \langle \Phi, \sigma \rangle$. Then:*

- (1) $\langle \Phi, \sigma \rangle \in \mathcal{U}(\mathbf{Pr})$, and
- (2) $\langle \Phi, \sigma \rangle \leq \langle \Phi', \sigma' \rangle$ for all other $\langle \Phi', \sigma' \rangle \in \mathcal{U}(\mathbf{Pr})$. That is, the solution is also a least or principal solution.

Proof. We work by induction on the length of the reduction $\mathbf{Pr} \implies^* \langle \mathbf{Pr} \rangle_{sol}$.

- Suppose \mathbf{Pr} is in normal form. Then the result is trivial since:

- (1) $\sigma = Id$, $\Phi \vdash \mathbf{Pr}Id$ and Id is idempotent;
- (2) For any other $\langle \Phi', \sigma' \rangle \in \mathcal{U}(\mathbf{Pr})$, σ' is such that $\Phi' \vdash \Phi\sigma$ and $\Phi' \vdash XId\sigma \overset{\lambda}{\approx}_\alpha X\sigma$ for all X .

- Suppose $\mathbf{Pr} \implies \mathbf{Pr}'$ by some non-instantiating simplification. Then using Lemma 4.11, we know that $\mathcal{U}(\mathbf{Pr}) = \mathcal{U}(\mathbf{Pr}')$. Both parts of the result follow by induction.

- Suppose $\text{Pr} \Longrightarrow^\theta \text{Pr}'\theta$ by an instantiating rule. Assume $\text{Pr} = \text{Pr}' \cup \{\pi \cdot X \stackrel{\wedge}{\approx}_\alpha t\}$ where $\theta = [X \mapsto \pi^{-1} \cdot t]$ and $X \notin \text{Var}(t)$ (the case for the other instantiating rule is similar). Suppose $\langle \text{Pr}'\theta \rangle_{\text{sol}} = \langle \Phi, \sigma \rangle$, so that by construction $\langle \text{Pr} \rangle_{\text{sol}} = \langle \Phi, \theta \circ \sigma \rangle$.
 - (1) It is easy to see that $\theta \circ \sigma$ is idempotent and by the first part of the inductive hypothesis $\Phi \vdash \text{Pr}'\theta\sigma$, that is, $\langle \Phi, \theta \circ \sigma \rangle \in \mathcal{U}(\text{Pr})$.
 - (2) Suppose $\langle \Phi', \sigma' \rangle \in \mathcal{U}(\text{Pr})$. Then $\Phi' \vdash X\sigma' \stackrel{\wedge}{\approx}_\alpha \pi^{-1} \cdot t\sigma'$ by Equivariance. Hence, $\langle \Phi', \theta \circ \sigma'' \rangle \in \mathcal{U}(\text{Pr})$ where σ'' acts just like σ' only it maps X to X , $\theta = [X \mapsto \pi^{-1} \cdot t]$, and $\sigma' = \theta \circ \sigma''$. Note that $\langle \Phi', \sigma'' \rangle \in \mathcal{U}(\text{Pr}'\theta)$ and by inductive hypothesis $\langle \Phi, \sigma \rangle \leq \langle \Phi', \sigma'' \rangle$. It follows that $\langle \Phi, \theta \circ \sigma \rangle \leq \langle \Phi', \theta \circ \sigma'' \rangle$. \square

Matching and unification are closely related notions. While unification is the basis of logic programming languages, matching is at the heart of rewriting and functional programming. A matching algorithm can be easily derived from a unification algorithm. First we recall the definition of matching.

Definition 4.13. A *matching problem* is a particular kind of unification problem Pr in which the variables in right-hand sides of equality constraints are disjoint from the variables in left-hand sides. A solution $\langle \Phi, \sigma \rangle$ for a matching problem Pr satisfies $\Phi \vdash s \stackrel{\wedge}{\approx}_\alpha t\sigma$ and $\Phi \vdash \pi \wedge t\sigma$, for each $s \stackrel{\wedge}{\approx}_\alpha t, \pi \wedge t \in \text{Pr}$ (i.e., in a matching problem, only the variables in right-hand sides of terms can be instantiated).

A nominal matching algorithm can be obtained from the unification algorithm specified in Table 1 by restricting the instantiation rules, so that only variables which were in left-hand sides of equality constraints in the initial problem can be instantiated.

Observation 4.14. Theorem 3.12 guarantees the equivalence between \approx_α and $\stackrel{\wedge}{\approx}_\alpha$, therefore, we can associate the unification algorithm proposed, with the standard nominal unification algorithm proposed in [UPG04]. The problem Pr introduced in Example 4.7, is equivalent to the nominal unification problem $\mathcal{P} = \{[a]f(X, a) \approx_\alpha [b]f((b c) \cdot W, (a c) \cdot Y)\}$, and using the standard simplification rules [UPG04]:

$$\begin{aligned} \mathcal{P} &\xrightarrow{*} \xrightarrow{[Y \mapsto b]} \xrightarrow{*} \mathcal{P}' = \{X \approx_\alpha (a b) \cdot ((b c) \cdot W), a \# W\} \\ &\xrightarrow{[X \mapsto (a b) \circ (b c) \cdot W]} \{a \# W\} = \mathcal{P}' \end{aligned} \quad (4.1)$$

The pair $\langle \mathcal{P} \rangle_{\text{sol1}} = \langle \{a \# W\}, \delta \rangle$, where $\delta = \{Y \mapsto b, X \mapsto (a b) \circ (b c) \cdot W\}$ is a solution for \mathcal{P} . Using the translation $[-]_\lambda$, we obtain $[\langle \mathcal{P} \rangle_{\text{sol1}}]_\lambda = \langle \{[a \# W]_\lambda\}, \delta \rangle = \langle (a c_a) \wedge W, \delta \rangle$, where c_a is a new name, which is equivalent to $\langle (a c_a) \wedge W, (c_a c_1) \wedge W, \delta \rangle$, for c_a and c_1 not occurring anywhere in \mathcal{P} . Therefore, $[\langle \mathcal{P}_{\text{sol1}} \rangle]_\lambda$ is a solution for $\text{Pr} = \{[a]f(X, a) \stackrel{\wedge}{\approx}_\alpha [b]f((b c) \cdot W, (a c) \cdot Y)\}$. Similarly, from the solution $\langle \Psi, \sigma \rangle$ proposed in Example 4.10, we obtain $\langle [\Psi]_\#, \sigma \rangle = \langle a \# W, c_1 \# W, c_2 \# W, \sigma \rangle$, which is a solution for \mathcal{P} .

In the theorem below Pr_λ denotes a unification problem w.r.t. $\stackrel{\wedge}{\approx}_\alpha$ and \wedge , and $\mathcal{P}_\#$ denotes a unification problem w.r.t. \approx_α and $\#$.

Theorem 4.15. Let Pr_λ and $\mathcal{P}_\#$ be unification problems such that $[\text{Pr}_\lambda]_\# = \mathcal{P}_\#$ and $\langle \Psi, \sigma \rangle \in \mathcal{U}(\text{Pr}_\lambda)$ and $\langle \Delta, \delta \rangle \in \mathcal{U}(\mathcal{P}_\#)$ be solutions for Pr_λ and $\mathcal{P}_\#$, respectively. Then

- (1) $\langle [\Psi]_\#, \sigma \rangle \in \mathcal{U}(\mathcal{P}_\#)$.
- (2) $\langle [\Delta]_\lambda, \delta \rangle \in \mathcal{U}(\text{Pr}_\lambda)$.

Proof. Consequence of Theorems 3.15 and 3.16 and the fact that substitution preserves $\#$, \approx_α , λ and $\overset{\lambda}{\approx}_\alpha$ judgements (see Proposition 3.13). \square

It is easy to see that the algorithm in Table 1 is exponential (as the one given in [UPG04]): even if we restrict the problems to first-order unification problems (without atoms), the simplification of the following problem requires a number of steps which is exponential with respect to the size of the problem.

$$h(f(X_0, X_0), \dots, f(X_{n-1}, X_{n-1})) \overset{\lambda}{\approx}_\alpha^? h(X_1, \dots, X_n)$$

Comparing this algorithm with the one based on freshness constraints, one notices that there is a correspondence between the simplification rules in both approaches (the term syntax is the same in both cases). Moreover, the simplification rules for fixed-point constraints work exactly like the simplification rules for freshness constraints (fixed-point rules have linear complexity, same as the simplification rules for freshness constraints in the freshness-based approach). The techniques designed to improve the efficiency of the freshness-based nominal unification algorithm (see, e.g., [Cal10, LV10]) can equally apply to the fixed-point based algorithm. To avoid the exponential complexity, terms should be represented via graphs, and permutations should be applied in a lazy way (see [Cal10, LV10] for details). We leave the development of efficient implementations to future work and in the rest of the paper we focus on unification modulo equational theories.

5. NOMINAL ALPHA-EQUIVALENCE MODULO EQUATIONAL THEORIES

In this section an extension of α -equality to take into account equational theories will be proposed. It is well-known that first-order unification modulo an arbitrary equational theory \mathbf{E} is undecidable, therefore, it is expected that nominal unification modulo \mathbf{E} (nominal \mathbf{E} -unification, for short) inherits this undecidability property. In this work, we propose an approach to deal with particular classes of equational theories, such as, associativity (\mathbf{A}), commutativity (\mathbf{C}) and associativity-commutativity (\mathbf{AC}), using fixed-point constraints.

In the case in which $\mathbf{E} = \mathbf{A}$, \mathbf{C} , \mathbf{AC} , or a combination of these theories, an algorithm to check \mathbf{E} - α -equality via freshness constraints was proposed in [AdCSFN17, AdCSMFR19], where correctness results were formally verified using the Coq proof assistant, and an implementation in Ocaml was given. Unification was considered only for \mathbf{C} theories, for which it was shown that in general there is no finitary representation of the set of solutions if solutions are represented by freshness contexts and substitutions.

We argue that the approach of fixed-point constraints is convenient when dealing with unification problems in equational theories that involve some notion of permutation of elements (such as commutativity), as it was shown in the previous version of this paper [AFN18].

5.1. Alpha-equivalence modulo \mathbf{E} via permutation fixed points. In this section the relations $\overset{\lambda}{\approx}_\alpha$ and λ will be extended to $\overset{\lambda}{\approx}_{\alpha, \mathbf{E}}$ and $\lambda_{\mathbf{E}}$, where \mathbf{E} is some equational theory. Inference rules will be parameterised by \mathbf{E} so they can be reused when other theories are exploited.

In this work we have dedicated rules for \mathbf{A} , \mathbf{C} and \mathbf{AC} , and their application depends on whether the signature Σ has function symbols satisfying these theories. Whenever we want to restrict the results to a particular theory, we will explicitly indicate the \mathbf{E} .

Similarly to the case of pure/syntactic α -equality we define notions of \mathbf{E} -fixed-point and \mathbf{E} - α -equality constraints, as well as \mathbf{E} -fixed-point contexts and \mathbf{E} -judgments as expected.

Definition 5.1 (\mathbf{E} -constraints and \mathbf{E} -fixed-point contexts). • An \mathbf{E} -fixed-point constraint is a pair of the form $\pi \lambda_{\mathbf{E}} t$, of a permutation π and a term t . An \mathbf{E} - α -equality constraint (for short, \mathbf{E} -equality constraint or just equality constraint) is a pair of the form $s \overset{\wedge}{\approx}_{\alpha, \mathbf{E}} t$, for nominal terms s and t .

- We call a fixed-point constraint of the form $\pi \lambda_{\mathbf{E}} X$ a *primitive \mathbf{E} -fixed-point constraint* and a finite set of such constraints is called an *\mathbf{E} -fixed-point context*. Υ, Ψ, \dots range over contexts.

Intuitively, $s \overset{\wedge}{\approx}_{\alpha, \mathbf{E}} t$ will mean that s and t are α -equivalent modulo the equational theory \mathbf{E} , and $\pi \lambda_{\mathbf{E}} t$ will mean that the permutation π has no effect on the equivalence class of the term t modulo \mathbf{E} . For instance, $(a c) \lambda_{\mathbf{C}} +(a, c)$, assuming $+$ is commutative, but not $(a c) \lambda_{\mathbf{C}} f(a, c)$, if f is not a commutative symbol.

Below we assume that commutative symbols are always applied to pairs (although the grammar of nominal terms permits application of function symbols to tuples, we assume a syntactic check is carried out in the case of \mathbf{C} -symbols).

Definition 5.2. An \mathbf{E} -fixed-point judgement is a tuple $\Upsilon \vdash \mathcal{V}\bar{c}.\pi \lambda_{\mathbf{E}} t$ of a fixed-point context and a fixed-point constraint, possibly with some newly quantified atoms, whereas an \mathbf{E} - α -equality judgement is a tuple $\Upsilon \vdash \mathcal{V}\bar{c}.s \overset{\wedge}{\approx}_{\alpha, \mathbf{E}} t$ of a fixed-point context and an \mathbf{E} -equality constraint, also with some newly quantified atoms \bar{c} .

\mathbf{E} -judgements are derived using the rules in Figures 8 and 9.

Notice that unlike the syntactical case, where the deduction rules for λ do not use $\overset{\wedge}{\approx}_{\alpha}$, here the rules for $\lambda_{\mathbf{E}}$ and $\overset{\wedge}{\approx}_{\alpha, \mathbf{E}}$ are mutually recursive when the theory \mathbf{E} involves function symbols with commutative and associative-commutativity properties (see rules $(\lambda_{\mathbf{E}}f^{\mathbf{C}})$ and $(\lambda_{\mathbf{E}}f^{\mathbf{AC}})$). We assume that the terms are flattened w.r.t. associative and associative-commutative function symbols.

Despite the fact that the rules are mutually recursive, the relations are well-defined since the recursion is well-founded. To see this, one can use a measure that interprets each equality judgement by the pair consisting of the maximum of the sizes of terms in the equality constraint and the symbol $\overset{\wedge}{\approx}_{\alpha, \mathbf{E}}$, and each fixed-point judgement by the pair consisting of the size of the term in the fixed-point constraint and the fixed-point symbol. For example, $\Upsilon \vdash \mathcal{V}\bar{c}.s \overset{\wedge}{\approx}_{\alpha, \mathbf{E}} t$ is interpreted by $\langle \max(|s|, |t|), \overset{\wedge}{\approx}_{\alpha, \mathbf{E}} \rangle$ and $\Upsilon \vdash \mathcal{V}\bar{c}.(a b) \lambda_{\mathbf{E}} t$ by $\langle |t|, \lambda_{\mathbf{E}} \rangle$. We compare these pairs lexicographically, using the ordering on natural numbers to compare the sizes of terms, and the ordering $\lambda_{\mathbf{E}} > \overset{\wedge}{\approx}_{\alpha, \mathbf{E}}$ for the second component of pairs. We can then check that in all the rules in Figures 8 and 9 the premises are strictly smaller than the conclusion according to this ordering, therefore the rules provide an inductive definition of the set of derivable judgements.

Rules $(\lambda_{\mathbf{E}}\mathbf{a})$, $(\lambda_{\mathbf{E}}\mathbf{var})$, $(\lambda_{\mathbf{E}}\mathbf{abs})$ and $(\lambda_{\mathbf{E}}\mathbf{tuple})$ behave exactly as the corresponding rules in Figure 3 (where $\mathbf{E} = \emptyset$), i.e., the theory \mathbf{E} has no effect on the fixed-point constraint.

Rule $(\lambda_{\mathbf{E}}f)$ is used for associative and uninterpreted function symbols. In the case of commutative or associative-commutative function symbols the rules $(\lambda_{\mathbf{E}}f^{\mathbf{C}})$ and $(\lambda_{\mathbf{E}}f^{\mathbf{AC}})$ are used. The goal is to ensure the analogous of Theorem 3.12 for the relations $\lambda_{\mathbf{E}}$ and $\overset{\wedge}{\approx}_{\alpha, \mathbf{E}}$

that we wish to obtain: $\Upsilon \vdash \mathcal{V}\bar{c}.\pi \lambda_E t$ iff $\Upsilon \vdash \bar{c}.\pi \cdot t \stackrel{\wedge}{\approx}_{\alpha, E} t$ (Theorem 5.6). This is illustrated in the example below.

$$\frac{\pi(a) = a}{\Upsilon \vdash \mathcal{V}\bar{c}.\pi \lambda_E a} (\lambda_E \mathbf{a}) \qquad \frac{\text{supp}(\pi^{\pi'^{-1}}) \setminus \{\bar{c}\} \subseteq \text{supp}(\text{perm}(\Upsilon|_X))}{\Upsilon \vdash \mathcal{V}\bar{c}.\pi \lambda_E \pi' \cdot X} (\lambda_E \mathbf{var})$$

$$\frac{\Upsilon \vdash \mathcal{V}\bar{c}.\pi \lambda_E t}{\Upsilon \vdash \mathcal{V}\bar{c}.\pi \lambda_E f^E t} (\lambda_E \mathbf{f}), \text{ if } f^E = f, f^A \qquad \frac{\Upsilon \vdash \mathcal{V}\bar{c}, c_1.\pi \lambda_E (a \ c_1) \cdot t}{\Upsilon \vdash \mathcal{V}\bar{c}.\pi \lambda_E [a]t} (\lambda_E \mathbf{abs})$$

$$\frac{\Upsilon \vdash \mathcal{V}\bar{c}.\pi \lambda_E t_1 \quad \dots \quad \Upsilon \vdash \mathcal{V}\bar{c}.\pi \lambda_C t_n}{\Upsilon \vdash \mathcal{V}\bar{c}.\pi \lambda_E (t_1, \dots, t_n)} (\lambda_E \mathbf{tuple})$$

$$\frac{\Upsilon \vdash \mathcal{V}\bar{c}.\pi \cdot (f^C(t_0, t_1)) \stackrel{\wedge}{\approx}_{\alpha, E} f^C(t_0, t_1)}{\Upsilon \vdash \mathcal{V}\bar{c}.\pi \lambda_E f^C(t_0, t_1)} (\lambda_E f^C)$$

$$\frac{\Upsilon \vdash \mathcal{V}\bar{c}.\pi \cdot (f^{AC}(t_0, t_1, \dots, t_n)) \stackrel{\wedge}{\approx}_{\alpha, E} f^{AC}(t_0, t_1, \dots, t_n)}{\Upsilon \vdash \mathcal{V}\bar{c}.\pi \lambda_E f^{AC}(t_0, t_1, \dots, t_n)} (\lambda_E f^{AC})$$

Figure 8: Fixed-point rules modulo A, C, AC.

$$\frac{}{\Upsilon \vdash \mathcal{V}\bar{c}.a \stackrel{\wedge}{\approx}_{\alpha, E} a} (\stackrel{\wedge}{\approx}_{\alpha, E} \mathbf{a}) \qquad \frac{\text{supp}((\pi')^{-1} \circ \pi) \setminus \{\bar{c}\} \subseteq \text{supp}(\text{perm}(\Upsilon|_X))}{\Upsilon \vdash \mathcal{V}\bar{c}.\pi \cdot X \stackrel{\wedge}{\approx}_{\alpha, E} \pi' \cdot X} (\stackrel{\wedge}{\approx}_{\alpha, E} \mathbf{var})$$

$$\frac{\Upsilon \vdash \mathcal{V}\bar{c}.s \stackrel{\wedge}{\approx}_{\alpha, E} (a \ b) \cdot t \quad \Upsilon \vdash \mathcal{V}\bar{c}, c_1.(a \ c_1) \lambda_E t}{\Upsilon \vdash \mathcal{V}\bar{c}.[a]s \stackrel{\wedge}{\approx}_{\alpha, E} [b]t} (\stackrel{\wedge}{\approx}_{\alpha, E} \mathbf{ab}) \qquad \frac{\Upsilon \vdash \mathcal{V}\bar{c}.t \stackrel{\wedge}{\approx}_{\alpha, E} t'}{\Upsilon \vdash \mathcal{V}\bar{c}.[a]t \stackrel{\wedge}{\approx}_{\alpha, E} [a]t'} (\stackrel{\wedge}{\approx}_{\alpha, E} \mathbf{[a]})$$

$$\frac{\Upsilon \vdash \mathcal{V}\bar{c}.t \stackrel{\wedge}{\approx}_{\alpha, E} t'}{\Upsilon \vdash \mathcal{V}\bar{c}.f^E t \stackrel{\wedge}{\approx}_{\alpha, E} f^E t'} (\stackrel{\wedge}{\approx}_{\alpha, E} \mathbf{f}), \text{ if } f^E = f, f^A \text{ or } t, t' \text{ are not pairs}$$

$$\frac{\Upsilon \vdash \mathcal{V}\bar{c}.t_1 \stackrel{\wedge}{\approx}_{\alpha, E} t'_1 \quad \dots \quad \Upsilon \vdash \mathcal{V}\bar{c}.t_n \stackrel{\wedge}{\approx}_{\alpha, E} t'_n}{\Upsilon \vdash \mathcal{V}\bar{c}.(\tilde{t})_{1..n} \stackrel{\wedge}{\approx}_{\alpha, E} (\tilde{t}')_{1..n}} (\stackrel{\wedge}{\approx}_{\alpha, E} \mathbf{t})$$

$$\frac{\Upsilon \vdash \mathcal{V}\bar{c}.s_0 \stackrel{\wedge}{\approx}_{\alpha, E} t_i \quad \Upsilon \vdash \mathcal{V}\bar{c}.s_1 \stackrel{\wedge}{\approx}_{\alpha, E} t_{(i+1) \bmod 2}}{\Upsilon \vdash \mathcal{V}\bar{c}.f^C(s_0, s_1) \stackrel{\wedge}{\approx}_{\alpha, E} f^C(t_0, t_1)} (\stackrel{\wedge}{\approx}_{\alpha, E} f^C) \text{ where } i \in \{0, 1\}$$

$$\frac{\Upsilon \vdash \mathcal{V}\bar{c}.s_0 \stackrel{\wedge}{\approx}_{\alpha, E} t_i \quad \Upsilon \vdash \mathcal{V}\bar{c}.f^{AC}(\tilde{s})_{1..n} \stackrel{\wedge}{\approx}_{\alpha, E} f^{AC}(\tilde{t})_{0..n}^{-i}}{\Upsilon \vdash \mathcal{V}\bar{c}.f^{AC}(\tilde{s})_{0..n} \stackrel{\wedge}{\approx}_{\alpha, E} f^{AC}(\tilde{t})_{0..n}} (\stackrel{\wedge}{\approx}_{\alpha, E} f^{AC}) \text{ where } i \in \{0, \dots, n\},$$

$$(\tilde{t})_n^{-i} = (t_0, \dots, t_{i-1}, t_{i+1}, \dots, t_n)$$

Figure 9: Rules for equality modulo A, C, AC

Example 5.3. Let $+$ be a commutative function symbol and suppose that we want to decide whether $\emptyset \vdash (a \ b) \lambda_C ((a+b) + c)$, which corresponds to $\emptyset \vdash (a \ b) \cdot ((a+b) + c) \stackrel{\wedge}{\approx}_{\alpha, C} (a+b) + c$ since $(a \ b) \cdot ((a+b) + c) = (b+a) + c \stackrel{\wedge}{\approx}_{\alpha, C} (a+b) + c$. In general, $\Upsilon \vdash \pi \lambda_C t_0 + t_1$

means that the permutation π fixes $t_0 + t_1$ modulo C (given the information in Υ), that is, $\Upsilon \vdash \pi \cdot (t_0 + t_1) \overset{\wedge}{\approx}_{\alpha, C} t_0 + t_1$. By definition, the permutation π distributes homomorphically over the operator $+$, therefore, we have $\Upsilon \vdash \pi \cdot t_0 + \pi \cdot t_1 \overset{\wedge}{\approx}_{\alpha, C} t_0 + t_1 \overset{\wedge}{\approx}_{\alpha, C} t_1 + t_0$. Thus, two cases can be distinguished: $\Upsilon \vdash \pi \cdot t_i \overset{\wedge}{\approx}_{\alpha, C} t_i$ or $\Upsilon \vdash \pi \cdot t_i \overset{\wedge}{\approx}_{\alpha, C} t_{(i+1) \bmod 2}$, for $i = 0, 1$ (see rule $(\overset{\wedge}{\approx}_{\alpha, C} \mathbf{f}^C)$ in Figure 9).

Similarly, α -equality rules (see Figure 4) have their equational counterpart in Figure 9: rules for atoms, tuples and abstractions are not affected by the theory E , whereas rules involving function symbols \mathbf{f}^E have to be analysed separately.

- Associative symbols: rule $(\overset{\wedge}{\approx}_{\alpha, E} \mathbf{f}^A)$ assumes that terms are flattened with respect to nested occurrences of \mathbf{f}^A .
- Commutative symbols: rule $(\overset{\wedge}{\approx}_{\alpha, E} \mathbf{f}^C)$ is used.
- Associative-commutative symbols: terms are assumed to be flattened with respect to the \mathbf{f}^{AC} function symbol, and rule $(\overset{\wedge}{\approx}_{\alpha, E} \mathbf{f}^{AC})$ is used.

Example 5.4. Consider the signature $\Sigma_A = \{\mathbf{f}^A\} \cup \Sigma^\emptyset$, where Σ^\emptyset is a set of uninterpreted function symbols. Rules $(\lambda_E \mathbf{f}^{AC})$, $(\lambda_E \mathbf{f}^C)$, $(\overset{\wedge}{\approx}_{\alpha, E} \mathbf{f}^C)$ and $(\overset{\wedge}{\approx}_{\alpha, E} \mathbf{f}^{AC})$ will not be used in judgements involving Σ_A -terms, therefore, we can replace E for A and obtain rules for $\overset{\wedge}{\approx}_{\alpha, A}$ and λ_A . For instance, consider $\mathbf{f}^A([a]a, \mathbf{f}^A(b, X))$, which is represented in flattened form as $\mathbf{f}^A([a]a, b, X)$.

$$\frac{\frac{\frac{\Upsilon \vdash \mathcal{M}_{C1} \cdot (a \ d) \ \lambda_A \ c_1}{\Upsilon \vdash (a \ d) \ \lambda_A \ [a]a} \ (\lambda_{Aa}) \quad (\lambda_{Aabs}) \frac{\Upsilon \vdash (a \ d) \ \lambda_A \ b \ (\lambda_{Aa}) \quad \frac{\{a, d\} \subseteq \text{supp}(\text{perm}(\Upsilon|_X))}{\Upsilon \vdash (a \ d) \ \lambda_A \ X} \ (\lambda_{Aa}) \ ?}{\Upsilon \vdash (a \ d) \ \lambda_A \ X} \ (\lambda_{Atuple})}{\Upsilon \vdash (a \ d) \ \lambda_A \ ([a]a, b, X)} \ (\lambda_{Atuple})}{\Upsilon \vdash (a \ d) \ \lambda_A \ \mathbf{f}^A([a]a, b, X)} \ (\lambda_{A\mathbf{f}^A})$$

The conclusion of this derivation depends on the support of the permutations in Υ , this is illustrated with the question mark ‘?’ in the rightmost leaf in the derivation. For example, if, on one hand, $\Upsilon = \emptyset$, then this derivation fails and we cannot conclude that $(a \ d)$ fixes $\mathbf{f}^A([a]a, b, X)$; if, on the other hand, $\Upsilon = \{(a \ b) \wedge X, (d \ e) \wedge X\} \cup \Upsilon'$, then we could conclude the opposite.

Example 5.5. Consider the signature $\Sigma_{\{C, AC\}} = \{\oplus^C, \text{or}^{AC}\} \cup \{\forall^\emptyset, \mathbf{g}^\emptyset\} \cup \Sigma^\emptyset$, where \forall and \mathbf{g} are unary uninterpreted function symbols. To improve readability, we will omit the superscripts of the function symbols in the rest of this example. Since we only have commutative and associative-commutative symbols, we will replace E in the rules by $\{C, AC\}$, therefore, obtaining rules for $\overset{\wedge}{\approx}_{\alpha, \{C, AC\}}$ and $\lambda_{\{C, AC\}}$. By applying the rules one can verify that

- $\emptyset \not\vdash (a \ b)(b \ c) \lambda_{\{C, AC\}} (\mathbf{g}(a) \oplus \mathbf{g}(b)) \oplus \mathbf{g}(c)$, this is due to the fact that \oplus is commutative but not associative, and the permutation $(a \ b)(b \ c)$ swaps the atom a which is an argument of the inner \oplus with the atom c which is an argument of the outer \oplus ;
- $\emptyset \vdash (a \ b)(b \ c) \lambda_{\{C, AC\}} \text{or}(\text{or}(\mathbf{g}(a), \mathbf{g}(b)), \mathbf{g}(c))$;
- $\emptyset \vdash (a \ b) \lambda_{\{C, AC\}} \forall[a] \text{or}(\text{or}((a \oplus b), (b \oplus c)), (a \oplus c))$;
- $(a \ b) \lambda_{\{C, AC\}} X \vdash \forall[a] \text{or}(\text{or}((a \oplus X), \mathbf{g}(c)), \mathbf{g}(a)) \overset{\wedge}{\approx}_{\alpha, \{C, AC\}} \forall[b] \text{or}(\text{or}(\mathbf{g}(c), ((a \ b) \cdot X \oplus b)), \mathbf{g}(b))$.

The theorem below extends Theorem 3.12, relating fixed-point constraints to fixed-point equalities, for the case in which equational theories A, C and AC are involved.

Theorem 5.6. *Let Υ, π and t be an E -fixed-point context, a permutation and a nominal term, respectively. $\Upsilon \vdash \mathcal{V}\bar{c}.\pi \lambda_E t$ iff $\Upsilon \vdash \mathcal{V}\bar{c}.\pi \cdot t \stackrel{\wedge}{\approx}_{\alpha, E} t$.*

Proof. The proof is by induction on the derivation, similar to the proof of Theorem 3.12, except for the use of rules dealing with function symbols modulo E . The rule for associative symbols is the same as for syntactic symbols; we consider the cases of commutative and associative-commutative symbols.

- Suppose that $\Upsilon \vdash \mathcal{V}\bar{c}.\pi \lambda_E f^C(t_0, t_1)$, therefore, rule $(\lambda_E f^C)$ was applied, and one gets $\Upsilon \vdash \mathcal{V}\bar{c}.\pi \cdot (f^C(t_0, t_1)) \stackrel{\wedge}{\approx}_{\alpha, E} f^C(t_0, t_1)$, and the result follows trivially. The other direction is also trivial.
- Suppose that $\Upsilon \vdash \mathcal{V}\bar{c}.\pi \lambda_E f^{AC}(t_0, \dots, t_n)$, therefore, rule $(\lambda_E f^{AC})$ was applied, and one gets $\Upsilon \vdash \mathcal{V}\bar{c}.\pi \cdot (f^{AC}(t_0, \dots, t_n)) \stackrel{\wedge}{\approx}_{\alpha, E} f^{AC}(t_0, \dots, t_n)$, and the result follows trivially. The other direction is also trivial. \square

5.2. From freshness to E -fixed-point constraints and back again. In [ARdCSFNS18] relations $\approx_{\{\alpha, A\}}$, $\approx_{\{\alpha, C\}}$, $\approx_{\{\alpha, AC\}}$ and their combination $\approx_{\{A, C, AC\}}$ were defined as extensions of \approx_α using the standard approach via freshness constraints (see the rules in Figures 1 and 2) but using specific rules for associative, commutative and associative-commutative symbols. We recall those rules in Figure 10.

$$\frac{\nabla \vdash s \approx_{\{A, C, AC\}} t}{\nabla \vdash f^E s \approx_{\{A, C, AC\}} f^E t} (\approx_{\{A, C, AC\}} \mathbf{app}), \text{ if } E \notin \{A, C, AC\} \text{ or both } s \text{ and } t \text{ are not pairs}$$

$$\frac{\nabla \vdash s_0 \approx_{\{A, C, AC\}} t_0 \quad f^A(\tilde{s})_{1..n} \approx_{\{A, C, AC\}} f^A(\tilde{t})_{1..n}}{\nabla \vdash f^A(\tilde{s})_{0..n} \approx_{\{A, C, AC\}} f^A(\tilde{t})_{0..n}} (\approx_{\{A, C, AC\}} A)$$

$$\frac{\nabla \vdash s_0 \approx_{\{A, C, AC\}} t_i \quad \nabla \vdash s_1 \approx_{\{A, C, AC\}} t_{(i+1) \bmod 2} \quad i = 0, 1}{\nabla \vdash f^C(s_0, s_1) \approx_{\{A, C, AC\}} f^C(t_0, t_1)} (\approx_{\{A, C, AC\}} C)$$

$$\frac{\nabla \vdash s_0 \approx_{\{A, C, AC\}} t_i \quad \nabla \vdash f^{AC}(\tilde{s})_{1..n} \approx_{\{A, C, AC\}} f^{AC}(\tilde{t})_{0..n}^{-i}}{\nabla \vdash f^{AC}(\tilde{s})_{0..n} \approx_{\{A, C, AC\}} f^{AC}(\tilde{t})_{0..n}} (\approx_{\{A, C, AC\}} AC)$$

Figure 10: Additional rules for equational α -equivalence via freshness constraints

Using a generalisation of the functions $[-]_\lambda$ and $[-]_\#$ defined in Section 3.3, we can obtain results that extend Theorem 3.15 and Theorem 3.16 to the equational case. The functions $[-]_\lambda^E$ and $[-]_\#^E$ defined below are the natural extension of the previous translation functions.

The mapping $[-]_\lambda^E$ associates each primitive freshness constraint in Δ with a primitive fixed-point constraint:

$$[a\#X]_\lambda^E = (a \ c_a) \lambda_E X \text{ where } c_a \text{ is a new name}$$

This mapping extends directly to contexts. We denote by $[\Delta]_\lambda^E$ the image of Δ under $[-]_\lambda^E$.

The mapping $[-]_{\#}^E$ associates each primitive fixed-point constraint in Υ with a primitive freshness context:

$$[\pi \lambda_E X]_{\#}^E = \text{supp}(\pi) \# X.$$

We denote by $[\Upsilon]_{\#}^E$ the union of the freshness contexts obtained by translating each constraint in Υ using $[-]_{\#}^E$.

Theorem 5.7. (1) $\Delta \vdash a \# t \Leftrightarrow [\Delta]_{\#}^E \vdash \mathcal{V}c.(a \ c) \lambda_E t.$
(2) $\Upsilon \vdash \mathcal{V}\bar{c}.\pi \lambda_E t \Leftrightarrow [\Upsilon]_{\#}^E \vdash \overline{\text{supp}(\pi) \setminus \{\bar{c}\}} \# t.$

Proof. The proof follows the same lines of the proof of Theorem 3.15. We discuss only the proof of part (1). The proof is by induction on rules of Figure 1 used in the derivation of $\Delta \vdash a \# t$. We show only the cases corresponding to function symbols.

(\Rightarrow)

• Rule ($\#f$)

In this case $t = f^E t'$, for some theory E . The analysis is based on the specific theory:

(1) $f^E = f^A$

This case is analogous to the case in which $E = \emptyset$.

(2) $f^E = f^C$

There is a proof Π of the form

$$\frac{\frac{\Pi}{\Delta \vdash a \# (t_1, t_2)}}{\Delta \vdash a \# f^C(t_1, t_2)} (\#f)$$

By induction hypothesis, there exists a proof Π' of $[\Delta]_{\lambda}^E \vdash \mathcal{V}c_a.(a \ c_a) \lambda_E (t_1, t_2)$. Therefore, there is a proof of the form

$$\frac{\frac{\Pi'_1}{[\Delta]_{\lambda}^E \vdash \mathcal{V}c_a.(a \ c_a) \lambda_E t_1} \quad \frac{\Pi'_2}{[\Delta]_{\lambda}^E \vdash \mathcal{V}c_a.(a \ c_a) \lambda_E t_2}}{[\Delta]_{\lambda}^E \vdash \mathcal{V}c_a.(a \ c_a) \lambda_E (t_1, t_2)} (\lambda_E \text{tuple})$$

From $[\Delta]_{\lambda}^E \vdash \mathcal{V}c_a.(a \ c_a) \lambda_E t_i$ one can derive that $[\Delta]_{\lambda}^E \vdash \mathcal{V}c_a.(a \ c_a) \cdot t_i \overset{\lambda}{\approx}_{\alpha, E} t_i$, by Theorem 5.6, for $(i = 1, 2)$. By applying rule ($\overset{\lambda}{\approx}_{\alpha, E} f^C$), it follows that

$$[\Delta]_{\lambda}^E \vdash \mathcal{V}c_a.(a \ c_a) \cdot (f^C(t_1, t_2)) \overset{\lambda}{\approx}_{\alpha, E} f^C(t_1, t_2),$$

and the result follows from Theorem 5.6.

(3) $f^E = f^A$

The proof is analogous to the case above.

(\Leftarrow) The interesting case is again for rule $(\lambda_E f^C)$.

Suppose that $[\Delta]_{\lambda}^E \vdash \mathcal{V}c.(a \ c) \lambda_E t_1 \oplus t_2$. We want to prove that $\Delta \vdash a \# t_1 \oplus t_2$. From rule $(\lambda_E f^C)$ we can conclude that

- either there exist proofs of $[\Delta]_{\lambda}^E \vdash \mathcal{V}c.(a \ c) \cdot t_1 \overset{\lambda}{\approx}_{\alpha, E} t_1$ and $[\Delta]_{\lambda}^E \vdash \mathcal{V}c.(a \ c) \cdot t_2 \overset{\lambda}{\approx}_{\alpha, E} t_2$, and by Theorem 5.6 it follows that there exist proofs of $[\Delta]_{\lambda}^E \vdash \mathcal{V}c.(a \ c) \lambda_E t_1$ and $[\Delta]_{\lambda}^E \vdash \mathcal{V}c.(a \ c) \lambda_E t_2$, and the result follows by induction hypothesis.
- or there exist proofs of $[\Delta]_{\lambda}^E \vdash \mathcal{V}c.(a \ c) \cdot t_1 \overset{\lambda}{\approx}_{\alpha, E} t_2$ and $[\Delta]_{\lambda}^E \vdash \mathcal{V}c.(a \ c) \cdot t_2 \overset{\lambda}{\approx}_{\alpha, E} t_1$.

Since these equalities hold for any new name c (not occurring in t_1 or t_2), then a cannot be in the support of t_1 and t_2 and therefore a is fresh in t_1 and t_2 . \square

We can now relate α -equivalence modulo E via freshness constrains ($\approx_{\{\alpha, E\}}$) with its version via fixed-point constraints ($\overset{\wedge}{\approx}_{\alpha, E}$).

Theorem 5.8. (1) $\Upsilon \vdash \mathcal{N}\bar{c}. s \overset{\wedge}{\approx}_{\alpha, E} t \Rightarrow [\Upsilon]_{\#}^E \cup \Delta \vdash s \approx_{\{\alpha, E\}} t$, where $\Delta \vdash \bar{c} \# \text{Var}(s, t)$.
 (2) $\Delta \vdash s \approx_{\{\alpha, E\}} t \Rightarrow [\Delta]_{\lambda}^E \vdash s \overset{\wedge}{\approx}_{\alpha, E} t$.

Proof. The proof is very similar to the proof of Theorem 3.16, except for the case of rules involving function symbols f^E , with $E \neq \emptyset$, where the reasoning is similar to the one in the proof of Theorem 5.7. \square

5.3. Solving nominal C-unification problems via fixed-point constraints. In this section we propose an approach to nominal unification modulo commutativity via the notion of fixed-point constraints.

For example, assuming $+$ is commutative, i.e., $X + Y = Y + X$, a problem of the form

$$+((a\ b) \cdot X, a) \overset{\wedge}{\approx}_{\alpha}^? + (Y, X) \quad (5.1)$$

can be solved by unifying $(a\ b) \cdot X$ with Y and a with X , or $(a\ b) \cdot X$ with X and a with Y .

In [ARdCSFNS18], a simplification algorithm for solving nominal C-unification was proposed. This algorithm was based on the standard nominal unification algorithm [UPG04] where α -equivalence is defined w.r.t. the notion of freshness. Upon the input of a unification problem Pr , the algorithm outputs a finite family of triples of the form $\langle \nabla, \sigma, P \rangle$, where ∇ is a freshness context, σ a substitution and P is a set of fixed-point equations, which are solved using a separate procedure.

In [ARdCSFNS17] it is proved that even a simple unification problem such as $(a\ b) \cdot X \approx_{\alpha} X$ (i.e., a problem consisting of just one fixed-point equation) could produce an infinite and independent set of solutions, whenever the signature contains commutative function symbols. For example, if f is commutative, the following substitutions solve this equation: $\{X \mapsto a + b, X \mapsto f(a + b), X \mapsto [e](a + b, b + a), \dots\}$. Therefore, it is not possible to obtain a finite and complete set of solutions for every solvable unification problem if solutions are expressed using freshness constraints and substitutions. However, we remark that the problem $+((a\ b) \cdot X, a) \overset{\wedge}{\approx}_{\alpha}^? (Y, X)$ mentioned above has in fact a finite number of most general solutions (indeed, two) if we solve it using fixed-point constraints. The most general unifiers are $\{X \mapsto a, Y \mapsto b\}$ and $\{Y \mapsto a, (a\ b) \lambda X\}$. This observation led us to use fixed-point constraints instead of freshness constraints to express solutions.

Similarly to Section 4, below we define the notion of nominal C-unification in terms of C-fixed-point constraints, and provide a nominal C-unification unification algorithm specified by means of simplification rules.

In this section, as in Section 4, we assume a generator of new names exists, and remove the new quantifier from the syntax of unification problems.

Definition 5.9. A C-unification problem Pr is a pair $\langle \Phi, P \rangle$ where P is a finite set of C-equality constraints $s \overset{\wedge}{\approx}_{\mathcal{C}}^? t$ and Φ is a finite set of C-fixed-point constraints $\pi \lambda_{\mathcal{C}}^? t$. To ease the notation, we will denote $s \overset{\wedge}{\approx}_{\mathcal{C}}^? t$ by $s \approx_{\mathcal{C}}^? t$.

Definition 5.10 (Solutions of C-unification problems). A solution for a C-unification problem $\text{Pr} = \langle \Phi, P \rangle$ is a pair $\langle \Upsilon, \sigma \rangle$, where the following conditions are satisfied

- (1) $\Upsilon \vdash \pi \lambda_C t \sigma$, if $\pi \lambda_C^? t \in \Phi$;
- (2) $\Upsilon \vdash s \sigma \stackrel{\wedge}{\approx}_{\alpha, C} t \sigma$, if $s \approx^? t \in P$.
- (3) $\Upsilon \vdash X \sigma \sigma \stackrel{\wedge}{\approx}_{\alpha, C} X \sigma$.

The set of solutions for a C-unification problem Pr is denoted as $\mathcal{U}_C(\text{Pr})$.

Definition 5.11 (Most general solution and complete set of solutions).

- For $\langle \Upsilon, \sigma \rangle$ and $\langle \Psi, \delta \rangle$ in $\mathcal{U}_C(\text{Pr})$, we say that $\langle \Upsilon, \sigma \rangle$ is *more general than* $\langle \Psi, \delta \rangle$, denoted $\langle \Upsilon, \sigma \rangle \preceq \langle \Psi, \delta \rangle$, if there exists a substitution ρ satisfying $\Psi \vdash \sigma \rho \stackrel{\wedge}{\approx}_{\alpha, C} \delta$ and $\Psi \vdash \Upsilon \rho$.
- A subset \mathcal{C} of $\mathcal{U}_C(\text{Pr})$ is a *complete set of solutions* of Pr if for all $\langle \Psi, \sigma \rangle \in \mathcal{U}_C(\text{Pr})$, there exists a $\langle \Upsilon, \delta \rangle \in \mathcal{C}$ such that $\langle \Upsilon, \delta \rangle \preceq \langle \Psi, \sigma \rangle$. We denote a complete set of solutions of the C-unification problem Pr as $\mathcal{C}(\text{Pr})$.

Table 2 presents the simplification rules for C-unification problems. They are derived from the deduction rules for judgements, as done for the syntactic case. The main difference is that now there are two rules for the simplification of fixed-point constraints involving commutative symbols (rules $(\lambda_C \text{f}^{\mathbf{C}1})$ and $(\lambda_C \text{f}^{\mathbf{C}2})$) and two rules to deal with equality of terms rooted by a commutative symbol (rules $(\stackrel{\wedge}{\approx}_{\alpha, C} \text{f}^{\mathbf{C}1})$ and $(\stackrel{\wedge}{\approx}_{\alpha, C} \text{f}^{\mathbf{C}2})$).

$(\lambda_C at)$	$\text{Pr} \uplus \{\pi \lambda_C^? a\}$	\implies	Pr , if $\pi(a) = a$
$(\lambda_C \text{f}^0)$	$\text{Pr} \uplus \{\pi \lambda_C^? ft\}$	\implies	$\text{Pr} \cup \{\pi \lambda_C^? t\}$, f not C
$(\lambda_C \text{f}^{\mathbf{C}1})$	$\text{Pr} \uplus \{\pi \lambda_C^? \text{f}^{\mathbf{C}}(t_0, t_1)\}$	\implies	$\text{Pr} \cup \{\pi \cdot t_0 \approx^? t_0, \pi \cdot t_1 \approx^? t_1\}$
$(\lambda_C \text{f}^{\mathbf{C}2})$	$\text{Pr} \uplus \{\pi \lambda_C^? \text{f}^{\mathbf{C}}(t_0, t_1)\}$	\implies	$\text{Pr} \cup \{\pi \cdot t_0 \approx^? t_1, \pi \cdot t_1 \approx^? t_0\}$
$(\lambda_C \text{tuple})$	$\text{Pr} \uplus \{\pi \lambda_C^? (\tilde{t})_{1..n}\}$	\implies	$\text{Pr} \cup \{\pi \lambda_C^? t_1, \dots, \pi \lambda_C^? t_n\}$
$(\lambda_C \text{abs})$	$\text{Pr} \uplus \{\pi \lambda_C^? [a]t\}$	\implies	$\text{Pr} \cup \{\pi \lambda_C^? (a \ c_1) \cdot t, (c_1 \ c_2) \lambda_C \text{Var}(t)\}$
$(\lambda_C \text{var})$	$\text{Pr} \uplus \{\pi \lambda_C^? \pi' \cdot X\}$	\implies	$\text{Pr} \cup \{\pi^{(\pi')^{-1}} \lambda_C^? X\}$, if $\pi' \neq Id$
$(\stackrel{\wedge}{\approx}_{\alpha, C} a)$	$\text{Pr} \uplus \{a \approx^? a\}$	\implies	Pr
$(\stackrel{\wedge}{\approx}_{\alpha, C} f)$	$\text{Pr} \uplus \{ft \approx^? ft'\}$	\implies	$\text{Pr} \cup \{t \approx^? t'\}$, f not C
$(\stackrel{\wedge}{\approx}_{\alpha, C} \text{f}^{\mathbf{C}1})$	$\text{Pr} \uplus \{\text{f}^{\mathbf{C}}(t_0, t_1) \approx^? \text{f}^{\mathbf{C}}(s_0, s_1)\}$	\implies	$\text{Pr} \cup \{t_0 \approx^? s_0, t_1 \approx^? s_1\}$
$(\stackrel{\wedge}{\approx}_{\alpha, C} \text{f}^{\mathbf{C}2})$	$\text{Pr} \uplus \{\text{f}^{\mathbf{C}}(t_0, t_1) \approx^? \text{f}^{\mathbf{C}}(s_0, s_1)\}$	\implies	$\text{Pr} \cup \{t_0 \approx^? s_1, t_1 \approx^? s_0\}$
$(\stackrel{\wedge}{\approx}_{\alpha, C} \text{tuple})$	$\text{Pr} \uplus \{(\tilde{t})_{1..n} \approx^? (\tilde{t}')_{1..n}\}$	\implies	$\text{Pr} \cup \{t_1 \approx^? t'_1, \dots, t_n \approx^? t'_n\}$
$(\stackrel{\wedge}{\approx}_{\alpha, C} \text{abs1})$	$\text{Pr} \uplus \{[a]t \approx^? [a]t'\}$	\implies	$\text{Pr} \cup \{t \approx^? t'\}$
$(\stackrel{\wedge}{\approx}_{\alpha, C} \text{abs2})$	$\text{Pr} \uplus \{[a]t \approx^? [b]s\}$	\implies	$\text{Pr} \cup \{t \approx^? (a \ b) \cdot s, (c_1 \ c_2) \lambda_C^? s, (c_1 \ c_2) \lambda_C \text{Var}(s)\}$
$(\stackrel{\wedge}{\approx}_{\alpha, C} \text{var})$	$\text{Pr} \uplus \{\pi \cdot X \approx^? \pi' \cdot X\}$	\implies	$\text{Pr} \cup \{(\pi')^{-1} \circ \pi \lambda_C^? X\}$
$(\stackrel{\wedge}{\approx}_{\alpha, C} \text{inst1})$	$\text{Pr} \uplus \{\pi \cdot X \approx^? t\}$	$\xRightarrow{[X \mapsto \pi^{-1}.t]}$	$\text{Pr}\{X \mapsto \pi^{-1}.t\}$, if $X \notin \text{Var}(t)$
$(\stackrel{\wedge}{\approx}_{\alpha, C} \text{inst2})$	$\text{Pr} \uplus \{t \approx^? \pi \cdot X\}$	$\xRightarrow{[X \mapsto \pi^{-1}.t]}$	$\text{Pr}\{X \mapsto \pi^{-1}.t\}$, if $X \notin \text{Var}(t)$

Table 2: Simplification Rules for C-unification problems via C-fixed-point constraints. In rules $(\lambda_C \text{abs})$ and $(\stackrel{\wedge}{\approx}_{\alpha, C} \text{abs2})$, c_1 and c_2 are newly generated names.

We write $\text{Pr} \implies_C \text{Pr}'$ when Pr' is obtained from Pr by applying a simplification rule from Table 2 and we write \implies_C^* for the reflexive and transitive closure of \implies_C . We omit the subindex when it is clear from the context.

Lemma 5.12 (Termination of simplification for C-unification problems). *There is no infinite chain of reductions \Longrightarrow_C starting from a C-unification problem Pr .*

Proof. Termination of the simplification rules follows directly from the fact that the following measure of the size of Pr is strictly decreasing: $[\text{Pr}] = (n_1, M)$ where n_1 is the number of different variables used in Pr , and M is the multiset of *heights* of equality constraints and non-primitive fixed-point constraints occurring in Pr .

Each simplification step either eliminates one variable (when an instantiation rule is used) and therefore decreases the first component of the interpretation, or leaves the first component unchanged but replaces a constraint with primitive ones and/or constraints where terms have smaller height. \square

The simplification rules (Table 2) specify a *C-unification algorithm*: we apply the simplification rules in a problem Pr until we reach normal forms. In the case of a term rooted by a commutative symbol, two rules can be applied, so a tree of derivations is built. The termination property (Lemma 5.12) guarantees the tree is finite.

For the leaves in the tree (i.e., normal forms), the notions of *consistency*, *failure*, *correctness* can be defined as in Section 4 (see Definition 4.8). So, if a normal form contains equality constraints, or inconsistent fixed-point constraints of the form $\pi \lambda_C^? a$ such that $\pi(a) \neq a$ then this normal form is a failure. Only leaves containing consistent fixed-point constraints produce solutions.

We now prove that the C-unification algorithm is sound and complete. The proof is done in two stages, first we show that the non-instantiating rules preserve solutions if we consider all the branches of the derivation tree (here it is important to consider all the branches: due to the non-deterministic application of rules involving commutative operators, if we consider just one branch we may loose solutions). Then we show that the set of solutions computed from all the successful leaves is a complete set of solutions for the initial problem.

Lemma 5.13 (Correctness of non-instantiating rules). *Let Pr be a C-unification problem and n a natural number. Assume $\text{Pr} \xrightarrow{n}_C \text{Pr}'_i$ ($i \in I$) are all the reduction sequences of length smaller than or equal to n starting from Pr that do not use instantiating rules ($\stackrel{\wedge}{\approx}_{\alpha, C} \text{inst1}$) and ($\stackrel{\wedge}{\approx}_{\alpha, C} \text{inst2}$). Then*

- (1) $\mathcal{U}_C(\text{Pr}) = \bigcup_{i \in I} \mathcal{U}_C(\text{Pr}'_i)$, and
- (2) if Pr'_i contains inconsistent reduced fixed-point constraints then $\mathcal{U}_C(\text{Pr}'_i) = \emptyset$.

Proof. For part (1), as for the proof of Lemma 4.11, we proceed by induction on n , but here we need to consider all the branches of length smaller than or equal to n in the derivation tree to ensure completeness.

The interesting cases are for the rules involving C function symbols, all the other cases are very similar to the proof of Lemma 4.11.

- Suppose that the last step of a simplification chain has the form

$$\text{Pr}_{n-1} = \text{Pr}' \uplus \{ \pi \lambda_C^? f^C(s_0, s_1) \} \Longrightarrow \text{Pr}' \cup \{ \pi \cdot s_0 \stackrel{\wedge}{\approx}_{\alpha, C}^? s_i, \pi \cdot s_1 \stackrel{\wedge}{\approx}_{\alpha, C} s_{(i+1) \bmod 2} \} = \text{Pr}_n^i.$$

In this case, the rule used is either $(\lambda_C f^C 1)$ or $(\lambda_C f^C 2)$.

Assume $i = 0$ and $(\lambda_C f^C 1)$ was used (the case $i = 1$ is identical).

There is another reduction sequence of the same length using $(\lambda_C f^C 2)$ and ending on Pr_n^1 (the same problem with $i = 1$).

Let $\langle \Psi, \sigma \rangle \in \mathcal{U}(\text{Pr}_{n-1})$ be a solution for Pr_{n-1} :

- (1) $\Psi \vdash \pi' \lambda_C t \sigma$, for all $\pi' \lambda^? t \in \text{Pr}'$ and $\Psi \vdash \pi \lambda_C \text{f}^C(s_0, s_1)\sigma$.
- (2) $\Psi \vdash t \sigma \stackrel{\wedge}{\approx}_{\alpha, C} s \sigma$, for all $t \approx^? s \in \text{Pr}'$.

Since $\Psi \vdash \pi \lambda_C \text{f}^C(s_0, s_1)\sigma$ and $\text{f}^C(s_0, s_1)\sigma = \text{f}^C(s_0\sigma, s_1\sigma)$, it follows that $\Psi \vdash \pi \lambda_C \text{f}^C(s_0\sigma, s_1\sigma)$. From $(\lambda_E \text{f}^C)$, one has that either there exist a proof for $\Psi \vdash \pi \cdot s_0\sigma \stackrel{\wedge}{\approx}_{\alpha, C} s_i\sigma$ and $\Psi \vdash \pi \cdot s_1\sigma \stackrel{\wedge}{\approx}_{\alpha, C} s_{(i+1) \bmod 2}\sigma$, for $i = 0$ or for $i = 1$. Therefore, $\langle \Psi, \sigma \rangle \in \mathcal{U}(\text{Pr}_n^0)$ or $\langle \Psi, \sigma \rangle \in \mathcal{U}(\text{Pr}_n^1)$ and the result follows.

The other direction is similar: if $\langle \Psi, \sigma \rangle \in \mathcal{U}(\text{Pr}_n^i)$, then $\langle \Psi, \sigma \rangle \in \mathcal{U}(\text{Pr}_{n-1})$.

- Suppose the last step of the simplification chain has the form

$$\text{Pr}_{n-1} = \text{Pr}' \uplus \{\text{f}^C(s_0, s_1) \stackrel{\wedge}{\approx}_{\alpha, C} \text{f}^C(t_0, t_1)\} \Longrightarrow \text{Pr}_n^i$$

where $\text{Pr}_n^i = \text{Pr}' \cup \{s_0 \stackrel{\wedge}{\approx}_{\alpha, C} t_i, s_1 \stackrel{\wedge}{\approx}_{\alpha, C} t_{(i+1) \bmod 2}\}$.

As above, it follows that there two branches of this form, for $i = 0$ and $i = 1$.

Let $\langle \Psi, \sigma \rangle \in \mathcal{U}(\text{Pr}_{n-1})$ be a solution for Pr_{n-1} . In particular, $\Psi \vdash \text{f}^C(s_0\sigma, s_1\sigma) \stackrel{\wedge}{\approx}_{\alpha, C} \text{f}^C(t_0\sigma, t_1\sigma)$. By applying rule $(\stackrel{\wedge}{\approx}_{\alpha, E} \text{f}^C)$, one has that there exist proofs of $\Psi \vdash s_0\sigma \stackrel{\wedge}{\approx}_{\alpha, C} t_i\sigma$ and $\Psi \vdash s_1\sigma \stackrel{\wedge}{\approx}_{\alpha, C} t_{(i+1) \bmod 2}\sigma$, for $i = 0$ or $i = 1$. Therefore, $\langle \Psi, \sigma \rangle \in \mathcal{U}(\text{Pr}_n^i)$ for $i = 0$ or $i = 1$, and the result follows.

Similarly, if $\langle \Psi, \sigma \rangle \in \mathcal{U}(\text{Pr}_n^i)$, then $\langle \Psi, \sigma \rangle \in \mathcal{U}(\text{Pr}_{n-1})$.

The second part of the lemma follows directly from the fact that inconsistent constraints are not derivable, therefore have no solutions. \square

As in Section 4, when Pr is a successful leaf, $\langle \text{Pr} \rangle_{\text{sol}}$ consists of the composition σ of all substitutions applied through the simplification steps and the fixed point context obtained.

Theorem 5.14 (Soundness and Completeness). *Let $\text{Pr} = \langle \Upsilon, P \rangle$ be a C-unification problem and let $\{\text{Pr}'_i \mid \text{Pr} \xrightarrow{*}_C \text{Pr}'_i$ and Pr'_i successful normal form $\}$ be the set of all the successful normal forms of Pr (i.e., leaves in the derivation tree without equality constraints or inconsistent fixed-point constraints).*

- (1) *If $\langle \Phi, \sigma \rangle \in \bigcup \langle \text{Pr}'_i \rangle_{\text{sol}}$ then $\langle \Phi, \sigma \rangle \in \mathcal{U}(\text{Pr})$, and*
- (2) *If $\langle \Phi, \sigma \rangle \in \mathcal{U}(\text{Pr})$, there exists $\langle \Phi', \sigma' \rangle$ such that $\langle \Phi', \sigma' \rangle \in \bigcup \langle \text{Pr}'_i \rangle_{\text{sol}}$ and $\langle \Phi', \sigma' \rangle \preceq \langle \Phi, \sigma \rangle$, that is, the set $\bigcup \langle \text{Pr}'_i \rangle_{\text{sol}}$ is a complete set of solutions.*

Proof. The proof is by induction on the length of a derivation $\text{Pr} \xrightarrow{*}_C \text{Pr}'_i$, distinguishing cases according to the first rule used. By Lemma 5.13, it is sufficient to check generality of solutions after each application of instantiation rules. If the first step uses a non-instantiating rule, then the previous lemma, together with the induction hypothesis, ensures that the set of solutions of Pr is exactly the set of solutions of its children. If the first step is instantiating, we proceed as in the proof of Theorem 4.12. \square

Observation 5.15. Using the approach to nominal α -equivalence via freshness, a nominal C-unification algorithm was presented in [ARdCSFNS18, ARdCSFNS17], which outputs solutions represented as triples $\langle \nabla, \sigma, P \rangle$ consisting of a freshness context ∇ , a substitution σ and a set P of fixed-point equations of the form $\pi \cdot X \stackrel{?}{\approx}_{\{\alpha, C\}} X$.

As with standard nominal unification, one can use the functions $[-]_{\#}$ and $[-]_{\lambda}$ to translate solutions $\langle \nabla, \sigma, P \rangle$ of nominal C-unification problems with freshness constraints as solutions $\langle [\nabla]_{\lambda} \cup \{P_{\lambda_C}\}, \sigma \rangle$ of nominal C-unification problems via C-fixed-point constraints, where $P_{\lambda_C} = \{\pi \lambda_C X \mid \pi \cdot X \stackrel{?}{\approx}_{\{\alpha, C\}} X \in P\}$.

A set of simplification rules generalising the C-unification algorithm to take into account A and AC symbols was proposed in Ribeiro’s thesis [dCS19] following the freshness constraint approach. The main difficulty is in the rules to deal with the AC symbols and with treatment of fixed-point equations of the form $\pi \cdot X \approx_{\{\alpha, AC\}}^? X$. Analytical proofs of soundness and completeness of such rules were given, and a formalisation in Coq was developed for the C-unification algorithm presented in [ARdCSFNS18]. In future work we will consider and relate AC-nominal unification with freshness and fixed-point constraints.

Regarding the complexity of the C-unification algorithm based on fixed-point constraints, we observe that in the syntactic case (i.e., without α -equivalence rules), the C-unification problem is NP-complete so it is expected that the algorithm will be exponential (Chapter 10 on Equational Unification in [BN98] surveys in detail the cases of C and AC unification). Comparing the nominal unification modulo C based on freshness and on fixed-point constraints, we can again notice that there is a one-to-one correspondence in the simplification rules, and thus the algorithms have the same behaviour during the simplification phase. The main difference is that using the freshness approach, a second algorithm is needed to solve fixed-point equations (generating an infinite number of solutions in general), which is avoided with the fixed-point approach.

6. CONCLUSIONS AND FUTURE WORK

The notion of fixed-point constraint allowed us to obtain a finite representation of solutions for nominal C-unification problems. This brings a novel alternative to standard nominal unification approaches in which just the algebra of atom permutations and the logic of freshness constraints are used to implement equational reasoning (e.g., [AK16, Cal13, CF08, Che10, FGM04]), and in particular to their extensions modulo commutativity, for which only infinite representations were possible in the standard approach. We have shown that with the new proposed approach the development of an algorithm to solve nominal equational problems modulo C is simpler, avoiding, thanks to the use of fixed-point constraints, the development of procedures for the generation of infinite independent sets of solutions.

In future work we plan to extend this approach to matching and unification modulo different equational theories as well as to the treatment of equational problems in nominal rewriting modulo. Future study will also address handling the case of Mal’cev permutative theories, which include n -ary functions with permutative arguments [Com93], as well as the more general and complex case of permutative equational theories [Sch89]. Finally, exploring the relation between Higher-Order Pattern unification in the style of Levy and Villaret [LV08] and nominal unification with fixed-point constraints would be also of great interest.

REFERENCES

- [AdCSFN17] M. Ayala-Rincón, W. de Carvalho Segundo, M. Fernández, and D. Nantes-Sobrinho. A formalisation of nominal α -equivalence with A and AC function symbols. *Electronic Notes in Theoretical Computer Science*, 332:21–38, 2017.
- [AdCSMFR19] M. Ayala-Rincón, W. de Carvalho Segundo, D. Nantes-Sobrinho M. Fernández, and A. C. Rocha-Oliveira. A formalisation of nominal α -equivalence with A, C, and AC function symbols. *Theor. Comput. Sci.*, 781:3–23, 2019.

- [AFN18] M. Ayala-Rincón, M. Fernández, and D. Nantes-Sobrinho. Fixed-point constraints for nominal equational unification. In *Proc. of the 3rd International Conference on Formal Structures for Computation and Deduction (FSCD)*, volume 108 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 7:1–7:16, 2018.
- [AK16] T. Aoto and K. Kikuchi. A Rule-Based Procedure for Equivariant Nominal Unification. In *Pre-proc. of Higher-Order Rewriting (HOR)*, pages 1–5, 2016.
- [ARdCSFNS17] M. Ayala-Rincón, W. de Carvalho Segundo, M. Fernández, and D. Nantes-Sobrinho. On Solving Nominal Fixpoint Equations. In *Proc. of the 11th Int. Symp. on Frontiers of Combining Systems (FroCoS)*, volume 10483 of *Lecture Notes in Computer Science*, pages 209–226. Springer Verlag, 2017.
- [ARdCSFNS18] M. Ayala-Rincón, W. de Carvalho Segundo, M. Fernández, and D. Nantes-Sobrinho. Nominal C-Unification. In *Proc. of the 27th Int. Symp. Logic-based Program Synthesis and Transformation (LOPSTR 2017), Revised Selected Papers*, volume 10855 of *Lecture Notes in Computer Science*, pages 235–251. Springer, 2018.
- [BKN87] D. Benanav, D. Kapur, and P. Narendran. Complexity of Matching Problems. *J. of Sym. Computation*, 3(1/2):203–216, 1987.
- [BN98] F. Baader and T. Nipkow. *Term Rewriting and All That*. Cambridge University Press, 1998.
- [BS94] F. Baader and J. H. Siekmann. Unification theory. In *Handbook of Logic in Artificial Intelligence and Logic Programming, Volume 2, Deduction Methodologies*, pages 41–126. Oxford University Press, 1994.
- [BSN⁺01] Franz Baader, Wayne Snyder, Paliath Narendran, Manfred Schmidt-Schauß, and Klaus U. Schulz. Unification theory. In *Handbook of Automated Reasoning (in 2 volumes)*, pages 445–532. Elsevier and MIT Press, 2001.
- [Cal10] C. F. Calvès. *Complexity and implementation of nominal algorithms*. PhD Thesis, King’s College London, 2010.
- [Cal13] C. F. Calvès. Unifying Nominal Unification. In *24th International Conference on Rewriting Techniques and Applications (RTA)*, volume 21 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 143–157, 2013.
- [CF08] C. F. Calvès and M. Fernández. A Polynomial Nominal Unification Algorithm. *Theor. Comput. Sci.*, 403(2-3):285–306, 2008.
- [Che05] J. Cheney. A Simpler Proof Theory for Nominal Logic. In *Proc. of the 8th Int. Conference on the Foundations of Software Science and Computational Structures (FOSSACS)*, volume 3441 of *Lecture Notes in Computer Science*, pages 379–394. Springer Verlag, 2005.
- [Che10] J. Cheney. Equivariant unification. *Journal of Automated Reasoning*, 45(3):267–300, 2010.
- [CM17] J. Cheney and A. Momigliano. α check: A mechanized metatheory model checker. *Theory and Practice of Logic Programming (TPLP)*, 17(3):311–352, 2017.
- [Com93] H. Comon. Complete Axiomatizations of Some Quotient Term Algebras. *Theor. Comput. Sci.*, 118(2):167–191, 1993.
- [CU08] J. Cheney and C. Urban. Nominal Logic Programming. *ACM Transactions on Programming Languages*, 30(5):1–47, 2008.
- [dCS19] W. de Carvalho Segundo. *Nominal Equational Problems Modulo Associativity, Commutativity and Associativity-Commutativity*. PhD thesis, Universidade de Brasília, Brazil, 2019.
- [Fag87] F. Fages. Associative-Commutative Unification. *J. of Sym. Computation*, 3:257–275, 1987.
- [FG07] M. Fernández and M. J. Gabbay. Nominal Rewriting. *Information and Computation*, 205(6):917–965, 2007.
- [FGM04] M. Fernández, M. J. Gabbay, and I. Mackie. Nominal Rewriting Systems. In *Proc. of the 6th Int. Conf. on Principles and Practice of Declarative Programming (PPDP)*, pages 108–119. ACM Press, 2004.
- [Gab00] M. J. Gabbay. *A Theory of Inductive Definitions With α -equivalence*. PhD thesis, DPMMS and Trinity College, University of Cambridge, 2000.
- [GM08] M. J. Gabbay and A. Mathijssen. Capture-avoiding substitution as a nominal algebra. *Formal Aspects of Computing*, 20(4-5):451–479, 2008.
- [GP02] M. J. Gabbay and A. M. Pitts. A New Approach to Abstract Syntax with Variable Binding. *Formal Aspects of Computing*, 13(3-5):341–363, 2002.

- [Hun80] T. W. Hungerford. *Algebra*. Number 73 in Graduate Texts in Mathematics. Springer Verlag, second edition, 1980.
- [LV08] J. Levy and M. Villaret. Nominal Unification from a Higher-Order Perspective. In *Proceedings 19th International Conference on Rewriting Techniques and Applications, RTA*, volume 5117 of *Lecture Notes in Computer Science*, pages 246–260. Springer, 2008.
- [LV10] J. Levy and M. Villaret. An Efficient Nominal Unification Algorithm. In *Proc. of the 21st Int. Conf. on Rewriting Techniques and Applications (RTA)*, volume 6 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 209–226, 2010.
- [Pit03] A. M. Pitts. Nominal Logic, a First Order Theory of Names and Binding. *Information and Computation*, 186(2):165–193, 2003.
- [Pit13] A. M. Pitts. *Nominal Sets: Names and Symmetry in Computer Science*. Cambridge University Press, 2013.
- [Sch89] M. Schmidt-Schauß. Unification in Permutative Equational Theories is Undecidable. *J. of Sym. Computation*, 8(4):415–421, 1989.
- [Sie90] J. H. Siekmann. Unification theory. *Decision Support Systems*, 6(4):315–337, 1990.
- [SKLV17] M. Schmidt-Schauß, T. Kutsia, J. Levy, and M. Villaret. Nominal Unification of Higher Order Expressions with Recursive Let. In *Proc. of the 26th International Symposium on Logic-Based Program Synthesis and Transformation (LOPSTR 2016), Revised Selected Papers*, volume 10184 of *Lecture Notes in Computer Science*, pages 328–344. Springer, 2017.
- [Sti81] M. E. Stickel. A Unification Algorithm for Associative-Commutative Functions. *J. of the ACM*, 28(3):423–434, 1981.
- [UPG04] C. Urban, A. M. Pitts, and M. J. Gabbay. Nominal Unification. *Theor. Comput. Sci.*, 323(1-3):473–497, 2004.