Solving Quantitative Equations

Georg Ehling and Temur Kutsia

RISC, Johannes Kepler University Linz, Austria
{gehling,kutsia}@risc.jku.at

Abstract. Quantitative equational reasoning provides a framework that extends equality to an abstract notion of proximity by endowing equations with an element of a quantale. In this paper, we discuss the unification problem for a special class of shallow subterm-collapse-free quantitative equational theories. We outline rule-based algorithms for solving such equational unification problems over generic as well as idempotent Lawvereian quantales and study their properties.

Keywords: Quantitative equational reasoning · Lawvereian quantales · Equational unification.

1 Introduction

Extending the equality predicate to a notion that expresses similarity or proximity is a task that has been addressed in various ways. While fuzzy reasoning [8,22] approaches this endeavor by equipping equations with real numbers between 0 and 1 to express the degree to which they hold true, quantitative algebraic reasoning [4,16] follows a more proximity-oriented approach, attempting to establish a notion of distance between two terms.

Recently, Gavazzo and Di Florio [11] introduced a framework of metric and quantitative equational reasoning that generalizes these approaches (with a slight modification). It is based on the idea of modeling abstract quantities in quantales [21], following Lawvere's fundamental work [15]. In this framework, equations between terms are endowed with an element of a Lawvereian quantale that expresses, in one sense or another, the degree to which they hold true. The exact meaning of this degree depends on the choice of the quantale; for instance, it could correspond to the distance of two terms in a metric space, or to the probability that the terms are equal. This approach is quite general and includes various known quantitative theories as special cases.

In recent years, quantitative and approximate techniques have become increasingly popular due to various applications. In these applications, e.g., in those related to reasoning about probabilistic computations [5], reasoning about privacy and security of systems [1, 20], reasoning about resource consumption during computation [7], approximate program transformations [12], etc. equalities are replaced with their quantitative approximations to model distances between programs, processes, or systems, resulting in metric-based approximate relations. Various techniques have been used to model such metric reasoning principles, among them quantitative equational logic, discussed in [3,4,16,17].

In this paper, we address one of the central problems in equational reasoning: unification (or solving equations). We approach this problem in the framework as described in [11], studying a generalization of classical unification to solving equations in a quantitative equational theory in Lawvereian quantales.

The theories we consider in this paper are induced by shallow subterm-collapse-free equations of a special form, between terms whose arguments are the same sequence of variables, e.g. $f(x_1, \dots, x_n)$ and $g(x_1, \dots, x_n)$. This is a natural first step toward investigating quantitative equational unification since such quantitative equations generalize the principle of "extending proximity between function symbols to proximity between terms" of unification in the fuzzy quantale to unification in an arbitrary Lawvereian quantale. Despite their simple form, such theories still pose several challenges (originating from, e.g., tensor-based transitivity or extending proximity between arguments to whole terms), which affect the notions of completeness and minimality of unifier sets. We redefine these notions and show that unification (modulo the abovementioned theories) in arbitrary Lawvereian quantale is finitary, while for idempotent Lawvereian quantales, it becomes unitary. We develop the corresponding unification algorithms and study their properties.

2 Preliminaries

We start by introducing the basic notions and fixing the terminology.

Quantales. For the notions in this part, we follow [10,11].

Definition 1 (Quantale). *A (unital)* quantale $\mathbb{\Omega} = (\Omega, \precsim, \otimes, \kappa)$ *consists of a monoid* $(\Omega, k, \otimes)$ *and a complete lattice* $(\Omega, \precsim)$ *(with join* $\vee$ *and meet* $\wedge$*) satisfying the following distributivity laws:* $\delta \otimes \left(\bigvee_{i \in I} \varepsilon_i\right) = \bigvee_{i \in I}(\delta \otimes \varepsilon_i)$ *and* $\left(\bigvee_{i \in I} \varepsilon_i\right) \otimes \delta = \bigvee_{i \in I}(\varepsilon_i \otimes \delta)$.

The element κ is called the *unit* of the quantale, and $\otimes$ is called its *tensor* (or multiplication). Besides κ, we use Greek letters $\varepsilon, \delta, \eta, \zeta, \iota$, and ω to denote elements of Ω. The *top* and *bottom* elements of a quantale are denoted by $\top$ and $\bot$, respectively. Quantales in which the unit κ coincides with $\top$ are called *integral quantales.* A quantale is *commutative* if its underlying monoid is. It is *non-trivial* if $\kappa \neq \bot$. It is *cointegral* if $\varepsilon \otimes \delta = \bot$ implies either $\varepsilon = \bot$ or $\delta = \bot$.

We assume our quantales are commutative, integral, cointegral, and nontrivial. Such quantales are called *Lawvereian.* (Note that the fuzzy quantale $\mathbb{I}$ is Lawvereian for the Gödel and product T-norms, but not for the Łukasiewicz T-norm.)

Tensors in quantales always have left and right *adjoints.* For commutative quantales, these adjoints are the same, defined as $\varepsilon \multimap \delta := \bigvee\{\eta \mid \varepsilon \otimes \eta \precsim \delta\}$.

An element $\iota \in \Omega$ is an *idempotent element* (or simply an *idempotent*) of a quantale $\mathbb{\Omega}$ if it satisfies $\iota \otimes \iota = \iota$. A quantale is called *idempotent* if every

Table 1. Correspondence between quantales Ω (generic), 2 (Boolean), $\mathbb{L}$ (Lawvere), $\mathbb{L}^{\max}$ (strong Lawvere), and $\mathbb{I}$ (fuzzy).

	Ω	2	$\mathbb{L}$	$\mathbb{L}^{\max}$	$\mathbb{I}$
Carrier	Ω	$\{0,1\}$	$[0,\infty]$	$[0,\infty]$	$[0,1]$
Order	$\precsim$	$\leqslant$	$\geqslant$	$\geqslant$	$\leqslant$
Join	$\vee$	$\exists$	inf	inf	sup
Meet	$\wedge$	$\forall$	sup	sup	inf
Tensor	$\otimes$	$\wedge$	$+$	max	left-continuous T-norm
Unit	κ	1	0	0	1

element is idempotent. Among the quantales in Table 1, the idempotent ones are 2, $\mathbb{L}^{\max}$, and $\mathbb{I}$ for the minimum (Gödel) T-norm.

In any Lawvereian quantale, (i) $\otimes$ is monotonous: $\alpha \precsim \beta \Rightarrow \alpha \otimes \gamma \precsim \beta \otimes \gamma$ (using distributivity: $(\alpha \otimes \gamma) \vee (\beta \otimes \gamma) = (\alpha \vee \beta) \otimes \gamma = \beta \otimes \gamma$); (ii) $\alpha \otimes \beta \precsim \alpha \wedge \beta$ (using monotonicity and integrality: $\alpha \otimes \beta \precsim \alpha \otimes \top = \alpha$).

Given a quantale Ω and $\varepsilon, \delta \in \Omega$, the *way-below relation* $\ll$ is defined as $\delta \ll \varepsilon$ iff for every $\Psi \subseteq \Omega$, if $\varepsilon \precsim \bigvee \Psi$ then there exists a finite subset $\Psi_0 \subseteq \Psi$ such that $\delta \precsim \bigvee \Psi_0$. A quantale Ω is called *continuous* if $\varepsilon = \bigvee_{\delta \ll \varepsilon} \delta$ for all $\varepsilon \in \Omega$.

Definition 2 (Ω-relations, Ω-ternary relations). *An* Ω-relation R *between sets A and B is a function $R\colon A \times B \to \Omega$. For any set A, the identity* Ω-relation *$\Delta_A\colon A \times A \to \Omega$ maps the diagonal elements (a,a) to κ, and all other elements to $\bot$. The composition $(R;S)\colon A \times C \to \Omega$ of two Ω-relations $R\colon A \times B \to \Omega$ and $S\colon B \times C \to \Omega$ is defined as $(R;S)(a,c) \coloneqq \bigvee_{b \in B} \big(R(a,b) \otimes S(b,c)\big)$.*

An Ω-ternary relation *over $A \times B$ is a ternary relation $R \subseteq A \times \Omega \times B$ such that $R(a,\varepsilon,b)$ implies $R(a,\delta,b)$ for any $\delta \precsim \varepsilon$.*

Any Ω-ternary relation R induces an Ω-relation $R^\bullet(a,b) \coloneqq \bigvee_{R(a,\varepsilon,b)} \varepsilon$, and any Ω-relation R induces an Ω-ternary relation $R^\circ(a,\varepsilon,b) :\Longleftrightarrow \varepsilon \precsim R(a,b)$. Moreover, we have $R^{\bullet\circ} = R^{\circ\bullet} = R$, and we can freely switch between Ω-ternary relations and Ω-relations.

The complete lattice structure of Ω lifts to Ω-relations pointwise, and we can say that an Ω-relation R on $A \times A$ is reflexive if $\Delta_A \precsim R$; transitive if $(R;R) \precsim R$; symmetric if $R^- \precsim R$ (where R^- is defined as $R^-(b,a) \coloneqq R(a,b)$). Thus, we get the notions of a preorder (i.e., reflexive and transitive) and equivalence (i.e., reflexive, transitive, and symmetric) Ω-relation.

Terms and substitutions. We assume that the reader is familiar with the standard notions of unification theory, see, e.g., [2]. A *signature* $\mathcal{F}$ is a set of function symbols, each equipped with a fixed nonnegative arity. The set of *terms* over a signature $\mathcal{F}$ and a set of variables $\mathcal{V}$ is denoted by $T(\mathcal{F},\mathcal{V})$. Given a term $t \in T(\mathcal{F},\mathcal{V})$, we denote by $\mathcal{V}(t)$ the set of variables appearing in t. A term is *ground* if it contains no variables. The notion of a *position* in a term is defined in the standard way.

The set of *leaves* of a term is defined as $\ell(e) \coloneqq \{e\}$, if e is a constant symbol or a variable, and $\ell(f(t_1,\ldots,t_n)) \coloneqq \bigcup_{i=1}^{n} \ell(t_i)$. (The leaves of a term t correspond to the leaves of the tree representing t.) If s is a subterm of t, then the *depth of s in t* as the minimal length of a position at which s occurs in t.

A *substitution* is a map $\sigma\colon \mathcal{V} \to T(\mathcal{F},\mathcal{V})$ which maps all but finitely many variables to themselves. Greek letters $\sigma, \varphi, \vartheta, \tau$ are used for them, while Id denotes the identity substitution. The set of substitutions is denoted by Sub. We use the set notation for substitutions, writing σ explicitly as a finite set $\{x \mapsto \sigma(x) \mid x \neq \sigma(x)\}$. The *domain* of σ is defined as $dom(\sigma) \coloneqq \{x \mid x \neq \sigma(x)\}$. A substitution σ extends naturally to an endomorphism on $T(\mathcal{F},\mathcal{V})$. The image of a term t under this endomorphism is denoted $t\sigma$.

3 Quantitative equational theories

We now fix a signature $\mathcal{F}$, a set of variables $\mathcal{V}$, and a Lawvereian quantale Ω.

Let $\approx_E$ be an Ω-ternary relation, assumed to be induced from a given set E of triples (t,ε,s), which we write as $\varepsilon \Vdash t \approx_E s$ (called Ω-equalities). A *quantitative equational theory* (or *Ω-equational theory*) $=_E$ is an Ω-ternary relation generated from $\approx_E$ by the rules in Fig. 1. We call E a *presentation* of $=_E$.

$$\text{(Ax)}\ \frac{\varepsilon \Vdash t \approx_E s}{\varepsilon \Vdash t =_E s} \qquad \text{(Refl)}\ \frac{}{\kappa \Vdash t =_E t} \qquad \text{(Sym)}\ \frac{\varepsilon \Vdash t =_E s}{\varepsilon \Vdash s =_E t} \qquad \text{(Trans)}\ \frac{\varepsilon \Vdash t =_E s \quad \delta \Vdash s =_E r}{\varepsilon \otimes \delta \Vdash t =_E r}$$

$$\text{(NExp)}\ \frac{\varepsilon_1 \Vdash t_1 =_E s_1 \quad \cdots \quad \varepsilon_n \Vdash t_n =_E s_n}{\varepsilon_1 \otimes \cdots \otimes \varepsilon_n \Vdash f(t_1,\dots,t_n) =_E f(s_1,\dots,s_n)} \qquad \text{(Subst)}\ \frac{\varepsilon \Vdash t =_E s}{\varepsilon \Vdash t\sigma =_E s\sigma}$$

$$\text{(Ord)}\ \frac{\varepsilon \Vdash t =_E s \quad \delta \precsim \varepsilon}{\delta \Vdash t =_E s} \qquad \text{(Join)}\ \frac{\varepsilon_1 \Vdash t =_E s \quad \cdots \quad \varepsilon_n \Vdash t =_E s}{\varepsilon_1 \vee \cdots \vee \varepsilon_n \Vdash t =_E s}$$

$$\text{(Arch)}\ \frac{\forall \delta \ll \varepsilon.\, \delta \Vdash t =_E s}{\varepsilon \Vdash t =_E s}$$

Fig. 1. Quantitative equational theory

Informally, we read $\varepsilon \Vdash t =_E s$ as "t and s are at most ε-apart modulo E" or "t and s are equal modulo E with degree ε".

Observe that the Ω-relation $=^\bullet_E$ induced from $=_E$ (i.e., $t =^\bullet_E s \coloneqq \bigvee_{\varepsilon \Vdash t =_E s} \varepsilon$) is a reflexive, symmetric, transitive quantitative relation that contains $\approx^\bullet_E$ and where function symbols and substitutions behave in a non-expansive way:

$$\big(t_1 =^\bullet_E s_1 \otimes \cdots \otimes t_n =^\bullet_E s_n\big) \precsim \big(f(t_1,\dots,t_n) =^\bullet_E f(s_1,\dots,s_n)\big),$$
$$\big(t =^\bullet_E s\big) \precsim \big(t\sigma =^\bullet_E s\sigma\big).$$

We will often slightly abuse terminology by calling both $=^\bullet_E$ and a presentation E a quantitative equational theory.

The rules in Figure 1 were introduced in [10] with the aim of generalizing previous approaches to quantitative reasoning [4, 16]. This generalization is achieved up to a slight modification of the (NExp) rule, whose analogue in [16] would feature the join of $\varepsilon_1, \ldots, \varepsilon_n$ rather than their tensor product.[1]

It should further be remarked that the (Join) rule also applies to an empty hypothesis, whence $\bot \Vdash t =_E s$ holds for any t and s. The infinitary (Arch) rule is needed to guarantee the semantic completeness of the deduction system in [16], but has no effect on $=_E$ whenever the presentation E is finite, whence it can be safely ignored in that case.

Analogous to classical equational theories, an Ω-equation $\varepsilon \Vdash t = s$, where s is a proper subterm of t, is called a *subterm-collapse equation*. A quantitative equational theory E is said to be *simple* (or *subterm-collapse-free*) if whenever $\varepsilon \Vdash t =_E s$ with $\varepsilon \neq \bot$ holds, the equation $\varepsilon \Vdash t = s$ is not subterm-collapsing. An equation $\varepsilon \Vdash t = s$ is called *shallow* [6] if the depth of each variable occurrence in t or in s is at most 1. An equational theory is called shallow if each equation in its presentation is shallow.

Definition 3. *Let E be an Ω-equational theory and $\mathcal{X}$ be a set of variables. A ternary relation $\lessapprox_{E,\mathcal{X}} \subseteq Sub \times \Omega \times Sub$ is defined as*

$$\lessapprox_{E,\mathcal{X}} (\sigma, \varepsilon, \vartheta) \text{ iff there exists a } \varphi \text{ such that } \varepsilon \Vdash x\sigma\varphi =_E x\vartheta \text{ for all } x \in \mathcal{X}.$$

In this case, we say that the substitution σ is more general than ϑ modulo E on $\mathcal{X}$ up to ε. *We shortly write $\sigma \lessapprox_{E,\mathcal{X},\varepsilon} \vartheta$ and call ϑ an* $(E, \mathcal{X}, \varepsilon)$-instance *of σ (or an $(E, \mathcal{X})$-instance with degree ε).*

It is not hard to see that $\lessapprox_{E,\mathcal{X}}$ is an Ω-ternary relation over $Sub \times Sub$. To show this, we need to prove that $\sigma \lessapprox_{E,\mathcal{X},\varepsilon} \vartheta$ implies $\sigma \lessapprox_{E,\mathcal{X},\delta} \vartheta$ for any $\delta \precsim \varepsilon$, which follows from the definition of $=_E$.

Lemma 1. *If $\sigma \lessapprox_{E,\mathcal{X},\varepsilon} \vartheta$ and $\vartheta \lessapprox_{E,\mathcal{Y},\delta} \psi$, then $\sigma \lessapprox_{E,\mathcal{X}\cap\mathcal{Y},\varepsilon\otimes\delta} \psi$.*

Proof. By definition of $\lessapprox_{E,\mathcal{X}}$ and $\lessapprox_{E,\mathcal{Y}}$, we have $\varepsilon \Vdash x\sigma\varphi_1 =_E x\vartheta$ for all $x \in \mathcal{X}$, and $\delta \Vdash y\vartheta\varphi_2 =_E y\psi$ for all $y \in \mathcal{Y}$. From these equalities, for all $z \in \mathcal{X} \cap \mathcal{Y}$ by the Subst rule we get $\varepsilon \Vdash z\sigma\varphi_1\varphi_2 =_E z\vartheta\varphi_2$ and $\delta \Vdash z\vartheta\varphi_2 =_E z\psi$. Therefore, by $\otimes$-transitivity (the Trans rule) we obtain $\varepsilon\otimes\delta \Vdash z\sigma\varphi_1\varphi_2 =_E z\psi$ for all $z \in \mathcal{X}\cap\mathcal{Y}$, which, by definition of $\lessapprox_{E,\mathcal{X}\cap\mathcal{Y}}$, gives $\sigma \lessapprox_{E,\mathcal{X}\cap\mathcal{Y},\varepsilon\otimes\delta} \psi$. □

Corollary 1. *If $\sigma \lessapprox_{E,\mathcal{X},\varepsilon} \vartheta$ and $\vartheta \lessapprox_{E,\mathcal{X},\delta} \psi$, then $\sigma \lessapprox_{E,\mathcal{X},\varepsilon\otimes\delta} \psi$.*

Theorem 1. *Given a set of Ω-equalities E and a set of variables $\mathcal{X}$, the Ω-relation $\lessapprox^{\bullet}_{E,\mathcal{X}}$ induced by $\lessapprox_{E,\mathcal{X}}$ is a preorder on Sub.*

Proof. We should show that $\lessapprox^{\bullet}_{E,\mathcal{X}}$ is reflexive and transitive.

- Reflexivity: We need to show $\kappa \precsim \sigma \lessapprox^{\bullet}_{E,\mathcal{X}} \sigma$ for all σ, which follows directly from the definitions of $\lessapprox_{E,\mathcal{X}}$ and $=_E$.

[1] The reason for this modification in [10] is that (NExp) should be compatible with (Trans), which is based on the tensor rather than the join. Without it, one would obtain a system where performing various transformation steps one after the other would lead to a different distance than performing the same steps in parallel.

– Transitivity: We should prove $\big(\sigma \lessapprox^{\bullet}_{E,\mathcal{X}} \vartheta \otimes \vartheta \lessapprox^{\bullet}_{E,\mathcal{X}} \psi\big) \precsim \big(\sigma \lessapprox^{\bullet}_{E,\mathcal{X}} \psi\big)$ for all σ, ϑ, and ψ. This statement can be inferred from Corollary 1. □

The equivalence relation on substitutions induced by $\lessapprox_{E,\mathcal{X}}$ is denoted by $\cong_{E,\mathcal{X}}$. It is an Ω-ternary relation. We write $\sigma \cong_{E,\mathcal{X},\varepsilon} \vartheta$ if $\varepsilon \precsim (\sigma \cong^{\bullet}_{E,\mathcal{X}} \vartheta)$.

Example 1. Let Ω be the Lawvere quantale $\mathbb{L} = ([0,\infty], \geqslant, +, 0)$ and consider $E = \{1 \Vdash a \approx b, 1 \Vdash b \approx c\}$, $\varepsilon = 2$ and $\mathcal{X} = \{x\}$. Let $\sigma = \{x \mapsto a\}$, $\vartheta = \{x \mapsto b\}$, and $\varphi = \{x \mapsto c\}$. Then we have:

– $\sigma \lessapprox_{E,\mathcal{X},\varepsilon} \vartheta$, because $x\sigma Id = a$, $x\vartheta = b$, and $1 \Vdash a =_E b$;
– $\vartheta \lessapprox_{E,\mathcal{X},\varepsilon} \varphi$, because $x\vartheta Id = b$, $x\varphi = c$, and $1 \Vdash b =_E c$;
– $\varphi \lessapprox_{E,\mathcal{X},\varepsilon} \sigma$, because $x\varphi Id = c$, $x\sigma = a$, and $2 \Vdash c =_E a$.

Hence, $\sigma \cong_{E,\mathcal{X},\varepsilon} \vartheta$, $\vartheta \cong_{E,\mathcal{X},\varepsilon} \varphi$, and $\sigma \cong_{E,\mathcal{X},\varepsilon} \varphi$. ■

Theorem 2. *Given E, $\mathcal{X}$, t, and s such that $\mathcal{V}(t) \cup \mathcal{V}(s) \subseteq \mathcal{X}$, let R denote $=^{\bullet}_E$ and S denote $\lessapprox^{\bullet}_{E,\mathcal{X}}$. Assume σ and ϑ are substitutions such that $R(t\sigma, s\sigma) = \varepsilon$ and $S(\sigma, \vartheta) = \delta$. Then $\varepsilon \otimes \bigotimes_{i=1}^{n+m} \delta \precsim R(t\vartheta, s\vartheta)$, where n and m are the number of occurrences of variables from $\mathcal{X}$ in t and s, respectively.*

Proof. From $S(\sigma, \vartheta) = \delta$ we know that there exists φ such that $\delta \precsim R(x\sigma\varphi, x\vartheta)$ holds for all $x \in \mathcal{X}$. From this, by structural induction over terms, we can prove $\bigotimes_{i=1}^{n} \delta \precsim R(t\sigma\varphi, t\vartheta)$ and $\bigotimes_{i=1}^{m} \delta \precsim R(s\sigma\varphi, s\vartheta)$. From $R(t\sigma, s\sigma) = \varepsilon$ we get $\varepsilon \precsim R(t\sigma\varphi, s\sigma\varphi)$. Applying transitivity twice we get $\varepsilon \otimes \bigotimes_{i=1}^{n+m} \delta \precsim R(t\vartheta, s\vartheta)$. □

Example 2. In the Boolean quantale $\mathbb{2}$, this theorem implies the well-known fact that if σ is a unifier of t and s and ϑ is an instance of σ, then ϑ is also a unifier of t and s. (In that case, $\varepsilon = \delta = 1$.)

Consider the fuzzy quantale $\mathbb{I}$ with the minimum T-norm, $E = \{0.5 \Vdash a \approx b, 0.7 \Vdash b \approx c\}$, $t = f(x, x, y)$, $s = f(y, b, c)$, $\mathcal{X} = \{x, y\}$, $\sigma = \{x \mapsto b, y \mapsto b\}$, and $\vartheta = \{x \mapsto a, y \mapsto c\}$. Then $0.5 \Vdash a =_E c$ and

$t\sigma = f(b, b, b)$, $s\sigma = f(b, b, c)$, $\varepsilon = 0.7$, $0.7 \Vdash t\sigma =_E s\sigma$;
$\delta = 0.5$, $\sigma \lessapprox_{E,\{x,y\},0.5} \vartheta$ (actually, $\sigma \cong_{E,\{x,y\},0.5} \vartheta$);
$t\vartheta = f(a, a, c)$, $s\vartheta = f(c, b, c)$;
Variables of $\mathcal{X}$ occur in t and s in total 4 times;
$\min(0.7, 0.5, 0.5, 0.5, 0.5) = 0.5$; $0.5 \Vdash t\vartheta =_E s\vartheta$.

Now consider the Lawvere quantale $\mathbb{L}$ with $E = \{1 \Vdash a \approx b, 1 \Vdash b \approx c, 1 \Vdash c \approx d, 1 \Vdash f(x) \approx g(x)\}$, $t = x$, $s = f(y)$, $\mathcal{X} = \{x, y\}$, $\sigma = \{x \mapsto g(y)\}$, and $\vartheta = \{x \mapsto g(a), y \mapsto d\}$. Besides, $\precsim = \geqslant$ and we have

$t\sigma = g(y)$, $s\sigma = f(y)$, $\varepsilon = 1$, $1 \Vdash t\sigma =_E s\sigma$;
$\delta = 2$, $\sigma \lessapprox_{E,\{x,y\},2} \vartheta$;
$t\vartheta = g(a)$, $s\vartheta = f(d)$;
Variables of $\mathcal{X}$ occur in t and s in total twice;
$\varepsilon + 2\delta = 5$;
$5 \Vdash t\vartheta =_E s\vartheta$. (In fact, $=^{\bullet}_E (t\vartheta, s\vartheta) = 1 + 3 = 4$ and $5 \precsim 4$.) ■

Theorem 2 implies that if δ is an idempotent element of the quantale and it can be absorbed by ε (i.e., $\varepsilon \otimes \delta = \varepsilon$), then $\varepsilon \Vdash t\vartheta =_E s\vartheta$. Obviously, this will be fulfilled if $\delta = \kappa$. In idempotent quantales, it will also hold when $\varepsilon \precsim \delta$. For idempotent elements, a stronger version of transitivity holds. Namely, if ι is an idempotent element of a quantale, then we have for all $E, t, s, r, \mathcal{X}, \sigma, \vartheta, \varphi$:

- $\iota \Vdash t =_E s$ and $\iota \Vdash s =_E r$ imply $\iota \Vdash t =_E r$,
- $\sigma \lessapprox_{E,\mathcal{X},\iota} \vartheta$ and $\vartheta \lessapprox_{E,\mathcal{X},\iota} \varphi$ imply $\sigma \lessapprox_{E,\mathcal{X},\iota} \varphi$.

4 Quantitative equational unification

Definition 4 (Quantitative equational unification). *A quantitative equational unification problem is formulated as follows:*

Given:	*A quantale Ω, $\varepsilon \in \Omega$ (called the threshold) with $\varepsilon \neq \bot$, a set of Ω-equalities E (a presentation of an equational theory), and two terms t and s.*
Find:	*A substitution σ such that $\varepsilon \Vdash t\sigma =_E s\sigma$.*

We call this problem (E,ε)-unification problem *over Ω. The quantale name is usually skipped when it does not cause confusion. For unification problems, we use the notation $\varepsilon \Vdash t =^?_E s$, called a* unification equation*, where the question mark indicates that the equation is supposed to be solved. For simplicity, we write the problem as $t =^?_{E,\varepsilon} s$, and further skip E if it is clear from the context.*

The substitution σ, if it exists, is called an (E,ε)-unifier *of t and s (alternatively, a unifier or a solution of $t =^?_{E,\varepsilon} s$) over Ω. In such a case we say that the given unification problem is* solvable*, or that the terms t and s are (E,ε)*-unifiable *over Ω. The set of all unifiers of $t =^?_{E,\varepsilon} s$ is denoted by $\mathfrak{U}_{E,\varepsilon}(t,s)$.*

For a given presentation E, a function symbol from $\mathcal{F}$ is called *free* if it does not appear in E. If $\mathcal{F}$ contains a free m-ary function symbol for some $m > 1$, the unification problem formulated above is equivalent to a problem of finding a solution of a system of unification equations (instead of a single equation) formulated as a constrained problem:

Given:	A quantale Ω, the threshold value $\varepsilon \in \Omega$ with $\varepsilon \neq \bot$, a set of Ω-equalities E (a presentation of an equational theory), and a set of term-pairs t_i, s_i, $1 \le i \le n$.
Find:	A substitution σ such that $\delta_i \Vdash t_i\sigma =_E s_i\sigma$, $1 \le i \le n$, for some δ_i with $\varepsilon \precsim \delta_1 \otimes \cdots \otimes \delta_n$.

In such a case we can write the unification problem as a pair of a set of unification equations and a $\precsim$-constraint: $\{t_1 =^?_{E,\boldsymbol{\alpha}_1} s_1, \ldots, t_n =^?_{E,\boldsymbol{\alpha}_n} s_n\}; \varepsilon \precsim \boldsymbol{\alpha}_1 \otimes \cdots \otimes \boldsymbol{\alpha}_n$, where $\boldsymbol{\alpha}_i$ are new metavariables whose values are also to be found alongside with variables that appear in the given t's and s's. This problem can be transformed to a single-equation problem. For instance, the constrained problem

$\{t_1 =^?_{E,\boldsymbol{\alpha}_1} s_1,\, t_2 =^?_{E,\boldsymbol{\alpha}_2} s_2,\, t_3 =^?_{E,\boldsymbol{\alpha}_3} s_3\}; \varepsilon \precsim \boldsymbol{\alpha}_1 \otimes \boldsymbol{\alpha}_2 \otimes \boldsymbol{\alpha}_3$ can be transformed to $f(f(t_1, t_2), t_3) =^?_{E,\varepsilon} f(f(s_1, s_2), s_3)$, where f is a free binary function symbol. The two problems are equivalent in the sense that they have exactly the same set of (E, ε)-unifiers. If the arity of f is bigger than the number of equations, the missing arguments will be filled in by fresh variables. For instance, for a quaternary f, the problem above can be encoded as $f(t_1, t_2, t_3, x) =^?_{E,\varepsilon} f(s_1, s_2, s_3, x)$, where x is fresh.

In classical unification, an important property of the instantiation relation is that any substitution that is less general than a given unifier of two terms will still be a unifier. In the quantitative case, we should take into account the approximation, which complicates things. First, we have the following fact as a consequence of Theorem 2:

Fact 1 *If σ is an (E, ε)-unifier of t and s and $\sigma \lessapprox_{E,\mathcal{V}(t,s),\delta} \vartheta$ for some δ, then ϑ is an $(E, \varepsilon \otimes \delta^k)$-unifier of t and s, where k is the total number of occurrences of variables in t and s.*

Some results become simpler for idempotent elements:

Lemma 2. *Let ι be an idempotent element of Ω, and σ and ϑ be substitutions.*

(i) If $\iota \Vdash y\sigma =_E y\vartheta$ holds for every $y \in \mathcal{Y} \subseteq \mathcal{V}$, then $\iota \Vdash r\sigma =_E r\vartheta$ holds for every term $r \in T(\mathcal{F}, \mathcal{Y})$.
(ii) Suppose that $\varepsilon \in \Omega$ satisfies $\varepsilon \otimes \iota = \varepsilon$. If $\sigma \lessapprox_{E,\mathcal{V}(t,s),\iota} \vartheta$ and σ is an (E, ε)-unifier of the terms t and s, then so is ϑ.

Proof. Part (i) is proved by structural induction, using idempotence. For (ii), as $\sigma \lessapprox_{E,\mathcal{V}(t,s),\iota} \vartheta$, there exists a substitution φ such that $\iota \Vdash x\sigma\varphi =_E x\vartheta$ holds for every variable $x \in \mathcal{V}(t, s)$. Thus, $\iota \Vdash r\sigma\varphi =_E r\vartheta$ holds for any term $r \in T(\mathcal{F}, \mathcal{V}(t, s))$ by (i). Hence, we have $\iota \Vdash s\vartheta =_E s\sigma\varphi$, $\varepsilon \Vdash s\sigma\varphi =_E t\sigma\varphi$, and $\iota \Vdash t\sigma\varphi =_E t\vartheta$. From these equalities, using $\varepsilon \otimes \iota = \varepsilon$, we get $\varepsilon \Vdash s\vartheta =_E t\vartheta$. □

This lemma implies that $(E, \mathcal{V}(t, s), \kappa)$-instances of (E, ε)-unifiers of t and s are still their (E, ε)-unifiers. Besides, it has the following corollary:

Corollary 2. *Let Ω be an idempotent quantale, in which σ is an (E, ε)-unifier of t and s, $\sigma \lessapprox_{E,\mathcal{V}(t,s),\delta} \vartheta$, and $\varepsilon \precsim \delta$. Then ϑ is an (E, ε)-unifier of t and s in Ω.*

These results motivate a specialized version of the notion of a minimal complete set of unifiers that we use in this paper:

Definition 5 (Minimal ι-complete set of unifiers). *Let P be an (E, ε)-unification problem over a quantale Ω and signature $\mathcal{F}$. Let $\mathcal{X} = \mathcal{V}(P)$ be the set of all variables of P, and let ι be an idempotent element of Ω such that $\varepsilon \otimes \iota = \varepsilon$. An* ι-complete set of (E, ε)-unifiers *of P is a set $\mathcal{C}$ of substitutions such that*

(1) $\mathcal{C} \subseteq \mathfrak{U}_{E,\varepsilon}(P)$, i.e., each element of $\mathcal{C}$ is an (E, ε)-unifier of P,
(2) for each $\vartheta \in \mathfrak{U}_{E,\varepsilon}(P)$ there exists $\sigma \in \mathcal{C}$ such that $\sigma \lessapprox_{E,\mathcal{X},\iota} \vartheta$.

The set $\mathcal{C}$ *is a* minimal ι-complete set of (E,ε)-unifiers *of* P *iff it is an* ι*-complete set that satisfies the following minimality property:*

(3) for all $\sigma, \sigma' \in \mathcal{C}$*, if* $\sigma \lesssim_{E,\mathcal{X},\iota} \sigma'$*, then* $\sigma = \sigma'$*.*

We denote a minimal ι*-complete set of* (E,ε)*-unifiers of* P *by* $mcsu_{E,\varepsilon,\iota}(P)$*.*

Given a unification problem with threshold ε, in order to make use of this definition, one first needs to find an idempotent element ι such that $\varepsilon \otimes \iota = \varepsilon$. In an arbitrary quantale, we can always take $\iota = \kappa$. If ε is idempotent itself, then we can also choose $\iota = \varepsilon$. [2]

Let P be an (E,ε)-unification problem. If it is unsolvable, then for any idempotent ι with $\varepsilon \otimes \iota = \varepsilon$ we have $mcsu_{E,\varepsilon,\iota}(t,s) = \emptyset$. Depending on E, ε, and ι, minimal ι-complete sets of (E,ε)-unifiers may not always exist. Even if they do, they may be infinite. When they exist, they are unique modulo the instantiation equivalence relation $\cong_{E,\mathcal{X},\iota}$.[3]

Example 3. Let Ω be the Lawvere quantale $\mathbb{L}$, E be the set of equations $E = \{1 \Vdash a \approx b, 1 \Vdash b \approx c, 1 \Vdash c \approx d\}$, $\varepsilon = 1$, $t = f(x,b)$, and $s = f(c,x)$.

The substitutions $\sigma = \{x \mapsto b\}$, $\vartheta = \{x \mapsto c\}$ are (E,ε)-unifiers of t and s:

- $t\sigma = f(b,b)$, $s\sigma = f(c,b)$ and $1 \Vdash f(b,b) =_E f(c,b)$,
- $t\vartheta = f(c,b)$, $s\vartheta = f(c,c)$ and $1 \Vdash f(c,b) =_E f(c,c)$.

In fact, $\{\sigma,\vartheta\} = mcsu_{E,\varepsilon,0}(t,s)$. Note that we have $\sigma \cong_{E,\mathcal{X},1} \vartheta$, but also $\sigma \cong_{E,\mathcal{X},1} \{x \mapsto a\}$ and $\vartheta \cong_{E,\mathcal{X},1} \{x \mapsto d\}$. However, neither $\{x \mapsto a\}$ nor $\{x \mapsto d\}$ is an (E,ε)-unifier of t and s (but they are $(E,3)$-unifiers of t and s). ■

The notion of an *occurrence cycle* will be needed later on.

Definition 6 (Occurrence cycle). *A set of unification equations* $\{x_1 \approx^?_{\varepsilon_1} t_1, \ldots, x_n \approx^?_{\varepsilon_n} t_n\}$ *constitutes an* occurrence cycle *if* t_i *is a non-variable term for at least one* i*,* $x_i \in \mathcal{V}(t_{i-1})$ *for* $1 < i \leqslant n$ *and* $x_1 \in \mathcal{V}(t_n)$*.*

4.1 Unification: simple shallow theories (special form)

In this section, we will consider Ω-equational theories that admit a presentation consisting of a finite number of equations of the form $\gamma \Vdash f(x_1,\ldots,x_n) \approx g(x_1,\ldots,x_n)$, where $n \geqslant 0$, $f \neq g$, and all x's are pairwise distinct. Also, the γ's in different equations can be different. This is the very basic form of quantitative axioms. The quantale Ω, as said above, is an arbitrary Lawvereian quantale. In the next subsection we consider the special case of idempotent Lawvereian quantales.

The presentation is shallow. One can easily show that the theory generated from such a presentation is simple. Hence, we consider simple shallow quantitative equational theories (of a special form). We will refer to such theories by E_{ssh}.

[2] This is how it is done, e.g., for fuzzy proximity/similarity relations [9,13,14,18,19,22].
[3] $\cong_{E,\mathcal{X},\iota}$ is a standard binary relation on *Sub* induced by the Ω-ternary relation $\cong_{E,\mathcal{X}}$ for a fixed ι.

In them, it makes sense to speak about the *approximation degree of two function symbols*, which is defined as $\mathfrak{d}_{E_{\mathsf{ssh}}}(f,g) \coloneqq \big(f(x_1,\dots,x_n) \overset{\bullet}{=}_{E_{\mathsf{ssh}}} g(x_1,\dots,x_m)\big)$ for an n-ary f and m-ary g. (Obviously, $\mathfrak{d}_{E_{\mathsf{ssh}}}(f,g) = \bot$ if $n \neq m$.) We say that f and g are $(E_{\mathsf{ssh}},\varepsilon)$-proximal, if $\varepsilon \precsim \mathfrak{d}_{E_{\mathsf{ssh}}}(f,g)$.

Remark 1. Since the equational theories we consider here are finitely presented, the degree of two function symbols of the same arity can be effectively computed as $\mathfrak{d}_{E_{\mathsf{ssh}}}(f,g) = \bigvee\{\gamma \mid \gamma \Vdash f(x_1,\dots,x_n) = g(x_1,\dots,x_n) \in C\}$, where C is the closure of the presentation of E_{ssh} under the (Trans) and (Sym) rules.

Theorem 3. *E_{ssh}-unification is finitary in a Lawvereian quantale Ω in the sense that for any $\varepsilon \in \Omega$, every $(E_{\mathsf{ssh}},\varepsilon)$-unification problem $t =^?_{E_{\mathsf{ssh}},\varepsilon} s$ has a finite minimal κ-complete set of unifiers.*

Proof (Sketch). Let $N_{E_{\mathsf{ssh}},\varepsilon}(r)$ denote an ε-neighborhood of a term r with respect to E_{ssh}, defined as the set of all terms obtained from r by replacing some function symbols by their $(E_{\mathsf{ssh}},\varepsilon)$-proximal ones. Then $mcsu_{E_{\mathsf{ssh}},\varepsilon,\kappa}(t,s) \subseteq \cup_{t'\in T, s'\in S}\{mgu(t' =^? s')\}$, where $T = N_{E_{\mathsf{ssh}},\varepsilon}(t)$, $S = N_{E_{\mathsf{ssh}},\varepsilon}(s)$, and $mgu(t' =^? s')$ is a most general unifier of the syntactic unification problem $t' =^? s'$. Since the presentation of theories of the form E_{ssh} is finite, the set $N_{E_{\mathsf{ssh}},\varepsilon}(r)$ is finite for $\varepsilon \neq \bot$ for any r. Hence, the set $\cup_{t'\in T, s'\in S}\{mgu(t' =^? s')\}$ is finite, which implies that $mcsu_{E_{\mathsf{ssh}},\varepsilon,\kappa}(t,s)$ is finite as well. □

Remark 2. In principle, the above proof already outlines an E_{ssh}-unification algorithm. However, there are several reasons for not using it: first, it would be a brute-force approach blindly replacing symbols with all their proximal ones in all possible ways. Second, it would not be sound because non-unifier answers would be returned and we would have to clean the computed set afterwards. Third, we want to keep our approach flexible, leaving equations between variables as a part of the output instead of forcing them to have only a syntactic solution.

In the following, we use bold-face upright Greek letters $\boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\gamma}$ for metavariables that range over the domain of the quantale. The rules constituting our unification method operate on configurations whose form is stated below.

Definition 7 (Configuration). *A* configuration *is either a special symbol* $\boldsymbol{F}$ *or a quadruple* $P; C; \delta; \sigma$, *where*

- *P is a set of unification equations of the form $t =^?_{\boldsymbol{\alpha}} s$, where $\boldsymbol{\alpha}$ is a metavariable (intuitively, P is the remaining problem to be solved),*
- *C is a constraint of the form $\varepsilon \precsim \boldsymbol{\alpha}_1 \otimes \cdots \otimes \boldsymbol{\alpha}_n$, where $\boldsymbol{\alpha}_1,\dots,\boldsymbol{\alpha}_n$ are metavariables and $\varepsilon \in \Omega$,*
- *δ is an element of the quantale domain (the current approximation degree),*
- *σ is a substitution (part of the unifier computed so far).*

In C, we also allow for the case where $n = 0$, in which the empty product on the right-hand side is κ by convention.

The solving rules for a theory E_{ssh} are given below. They operate on configurations and are formulated modulo associativity and commutativity of $\otimes$. We use f and g to denote (not necessarily distinct) n-ary function symbols, t, t_i and s_i for terms and x to denote a variable symbol. The symbol Δ denotes the tensor product of finitely many metavariables.

Tri: **Trivial**

$$\{t =^?_{\boldsymbol{\alpha}} t\} \uplus P; \zeta \precsim \boldsymbol{\alpha} \otimes \Delta; \delta; \sigma \Longrightarrow P; \zeta \precsim \Delta; \delta; \sigma.$$

Dec: **Decompose**

$$\{f(t_1, \dots, t_n) =^?_{\boldsymbol{\alpha}} g(s_1, \dots, s_n)\} \uplus P; \zeta \precsim \boldsymbol{\alpha} \otimes \Delta; \delta; \sigma \Longrightarrow$$
$$\{t_1 =^?_{\boldsymbol{\beta}_1} s_1, \dots, t_n =^?_{\boldsymbol{\beta}_n} s_n\} \cup P;$$
$$\mathfrak{d}_{E_{\mathsf{ssh}}}(f, g) \multimap \zeta \precsim \boldsymbol{\beta}_1 \otimes \cdots \otimes \boldsymbol{\beta}_n \otimes \Delta; \delta \otimes \mathfrak{d}_{E_{\mathsf{ssh}}}(f, g); \sigma,$$

where $\boldsymbol{\beta}_1, \dots, \boldsymbol{\beta}_n$ are new metavariables and $\zeta \precsim \mathfrak{d}_{E_{\mathsf{ssh}}}(f, g)$.

Cla: **Clash**

$$\{f(t_1, \dots, t_n) =^?_{\boldsymbol{\alpha}} g(s_1, \dots, s_m)\} \uplus P; \zeta \precsim \boldsymbol{\alpha} \otimes \Delta; \delta; \sigma \Longrightarrow \mathbf{F}, \text{ if } \zeta \not\precsim \mathfrak{d}_{E_{\mathsf{ssh}}}(f, g).$$

L-Sub: **Substitute (lazy)**

$$\{x =^?_{\boldsymbol{\alpha}} f(s_1, \dots, s_n)\} \uplus P; \zeta \precsim \boldsymbol{\alpha} \otimes \Delta; \delta; \sigma \Longrightarrow$$
$$\{x_1 =^?_{\boldsymbol{\beta}_1} s_1, \dots, x_n =^?_{\boldsymbol{\beta}_n} s_n\} \cup P\rho;$$
$$\mathfrak{d}_{E_{\mathsf{ssh}}}(f, g) \multimap \zeta \precsim \boldsymbol{\beta}_1 \otimes \cdots \otimes \boldsymbol{\beta}_n \otimes \Delta; \delta \otimes \mathfrak{d}_{E_{\mathsf{ssh}}}(f, g); \sigma\rho,$$

where x does not appear in an occurrence cycle in $\{x =^?_{\boldsymbol{\alpha}} f(s_1, \dots, s_n)\} \cup P$, and $\rho = \{x \mapsto g(x_1, \dots, x_n)\}$ with $x_1, \dots, x_n$ being fresh variables and $\zeta \precsim \mathfrak{d}_{E_{\mathsf{ssh}}}(f, g)$.

CCh: **Cycle check**

$$\{x =^?_{\boldsymbol{\alpha}} t\} \uplus P; C; \delta; \sigma \Longrightarrow \mathbf{F},$$

if x appears in an occurrence cycle in $\{x =^?_{\boldsymbol{\alpha}} t\} \uplus P$.

Ori: **Orient**

$$\{t =^?_{\boldsymbol{\alpha}} x\} \uplus P; C; \delta; \sigma \Longrightarrow P \cup \{x =^?_{\boldsymbol{\alpha}} t\}; C; \delta; \sigma, \text{ where } t \notin \mathcal{V}.$$

To solve an $(E_{\mathsf{ssh}}, \varepsilon)$-unification problem between terms t and s, we create the initial configuration $\{t =^?_{\boldsymbol{\alpha}} s\}; \varepsilon \precsim \boldsymbol{\alpha}; \kappa; Id$ and start applying the rules as long as possible. The equation to be transformed is chosen arbitrarily ("don't care nondeterminism"). We call the obtained algorithm QUnif.

Note that a configuration $P; C; \sigma$ obtained from an $(E_{\mathsf{ssh}}, \varepsilon)$-unification problem satisfies the following properties:

- Any metavariable occurring in P also occurs in C and vice versa.
- No metavariable appears more than once in P or C.
- The domain of σ is disjoint from the set of variables occurring in P.

We will refer to such configurations as *admissible*.

To prove termination of QUNIF, we introduce some terminology.

Definition 8. *Let P be a set of quantitative equations and let P_{st} be the set of standard equations obtained from P by ignoring the indices: $P_{\mathsf{st}} \coloneqq \{t = s \mid \varepsilon \Vdash t = s \in P\}$. Then $\mathsf{DecNF}_{E_{\mathsf{ssh}}}(P)$ denotes the decomposition normal form of P with respect to E_{ssh}, which is the set of standard equations obtained from P_{st} by applying the following version of the decomposition rule as long as possible:*

$$\{f(t_1, \dots, t_n) = g(s_1, \dots, s_n)\} \uplus S \Longrightarrow \{t_1 = s_1, \dots, t_n = s_n\} \cup S,$$

where $\mathfrak{d}_{E_{\mathsf{ssh}}}(f, g) \neq \bot$.

It is easy to see that every equation in $\mathsf{DecNF}_{E_{\mathsf{ssh}}}(P)$ is of the form $x = s$ where s is an arbitrary term, or $t = x$ where t is not a variable.

For a set of (quantitative) equations P, the *variable dependency graph* $\Gamma(P)$ is constructed as follows:

- For each variable x appearing in P, add a node with label x to $\Gamma(P)$.
- Add a node with label G (the "ground node").
- For every equation $x = y \in \mathsf{DecNF}(P)$ between variables x and y, merge the nodes corresponding to x and y.
- In order to construct the set of edges of $\Gamma(P)$, we consider all equations of the form $x = t$ (or $t = x$) in $\mathsf{DecNF}(P)$, where t is a non-variable term. For such an equation, we consider the set of leaves of t. For each element $l \in \ell(t)$, if d is the depth in which l appears in t, we add an weighted edge to $\Gamma(P)$:
  - If l is a constant, then we add an edge $x \to_{d+1} G$ (with weight $d + 1$).
  - If l is a variable y, then we add an edge $x \to_d y$ (with weight d).

In this way, we obtain a directed, weighted graph $\Gamma(P)$, which is acyclic (hence, a *dag*) if and only if P does not contain any occurrence cycles.

For any variable x occurring in P, we define now the level $lev_P(x)$ of x with respect to P as the maximal weight of a walk in $\Gamma(P)$ starting in x. Here, the weight of a walk is defined as the sum of the weights of its edges. Note that lev_P may take the value ∞ if P contains occurrence cycles.

We now consider the multiset $\lambda(P) \coloneqq \{lev_P(x) \mid x \in \mathcal{V}(P)\}$. (It will be used as a component of a termination measure below.) We compare such multisets via the multiset extension $>_m$ of the standard order on $\mathbb{N} \cup \{\infty\}$, which is well-founded. The following lemma is the main ingredient for the termination proof.

Lemma 3. *Let $\mathfrak{C} = P; C; \delta; \sigma$ be a configuration.*

(i) If $P'; C'; \delta'; \sigma'$ is obtained from $\mathfrak{C}$ by $\mathsf{L\text{-}Sub}$*, then $\lambda(P) >_m \lambda(P')$.*
(ii) If $P'; C'; \delta'; \sigma'$ is obtained from $\mathfrak{C}$ by Tri*,* Dec*, or* Ori*, then $\lambda(P) \geqslant_m \lambda(P')$.*

Proof. For (i), we shall investigate how the variable dependency graphs of P and P' differ. We can write $P = P_0 \cup \{x =^?_{\boldsymbol{\alpha}} f(t_1, \dots, t_n)\}$ and $P' = P_0\rho \cup \{x_1 =^?_{\boldsymbol{\beta}_1} t_1, \dots, x_n =^?_{\boldsymbol{\beta}_n} t_n\}$, where $\rho = \{x \mapsto g(x_1, \dots, x_n)\}$ and $\mathfrak{d}_{E_{\mathsf{ssh}}}(f, g) \neq \bot$. Note that the set of nodes of $\Gamma(P')$ differs from the set of nodes of $\Gamma(P)$ only in as far as the node corresponding to x has been removed and new nodes have been added for $x_1, \ldots, x_n$. As for the edges, note that any edge of $\Gamma(P)$ that connects two nodes with labels different from x will be present also in $\Gamma(P')$. Considering now the incoming edges of x in $\Gamma(P)$, if $y \to_d x$ is such an edge, then there is some equation $y =^?_{\boldsymbol{\beta}} t$ in P, where x appears in depth d in t. Thus, $y =^?_{\boldsymbol{\beta}} t\rho$ is an equation in P', and $x_1, \ldots, x_n$ appear in depth $d+1$ in $t\rho$; so $\Gamma(P')$ features edges $y \to_{d+1} x_1, \ldots, y \to_{d+1} x_n$. If $x \to_d z$ is an outgoing edge from x in $\Gamma(P)$, then there P contains an equation of the form $x =^?_{\boldsymbol{\gamma}} s$, where z appears in depth d in s. As $d \geqslant 1$, we can write $s = h(s_1, \ldots, s_m)$. Then P' contains the equation $g(x_1, \ldots, x_n) =^? h(s_1, \ldots, s_m)$. If $\mathfrak{d}_{E_{\mathsf{ssh}}}(h, g) \neq \bot$, then $\Gamma(P')$ contains an edge $x \to_{d-1} y$ as $y \in \mathcal{V}(s_i)$ holds for some i. If $\mathfrak{d}_{E_{\mathsf{ssh}}}(h, g) = \bot$, then $\Gamma(P')$ does not contain such an edge. The same holds also if y is the ground node. Summing up:

- Edges not involving x are the same in $\Gamma(P)$ and $\Gamma(P')$.
- Edges $y \to_d x$ in $\Gamma(P)$ correspond to edges $y \to_{d+1} x_i$ in $\Gamma(P')$.
- Edges $x \to_d y$ in $\Gamma(P)$ are either removed, or correspond to edges $x_i \to_{d-1} y$ in $\Gamma(P')$.

This means that the walks in $\Gamma(P)$ not passing through x are the same as in $\Gamma(P')$; and for every walk in $\Gamma(P')$ passing through x_i, one can find a walk of the same length in $\Gamma(P)$ (which will pass through x one step before). Thus, we have $lev_{P'}(y) = lev_P(y)$ for every variable y different from x, $x_1, \ldots, x_n$. As for walks in $\Gamma(P')$ starting in some x_i, note first that any such walk must be finite as $\Gamma(P')$ is finite and x does not appear in any occurrence cycle, whence x_i cannot be part of a cycle in $\Gamma(P')$. Moreover, for any walk in $\Gamma(P')$ starting in x_i of weight w, one can find a walk in $\Gamma(P)$ starting in x of weight $w+1$ (because the first edge of the walk can be replaced by an edge of weight one greater). Therefore, we have $lev_P(x) = lev_P(x_i) + 1$ for each i, and thus, we obtain $\lambda(P') <_m \lambda(P)$.

For (ii), if the applied rule was Tri or Ori, then $P' \subseteq P$, whence every walk in $\Gamma(P')$ is also present in $\Gamma(P)$, thus yielding $\lambda(P') \leqslant \lambda(P)$. If the applied rule was Dec, we can write $P = P_0 \cup \{f(t_1, \ldots, t_n) =^?_{\boldsymbol{\alpha}} g(s_1, \ldots, s_n)\}$ and $P' = P_0 \cup \{t_1 =^?_{\boldsymbol{\beta}_1} s_1, \ldots, t_n =^?_{\boldsymbol{\beta}_1} s_n\}$. It is easy to see that we have $C_E(P) = C_E(P') \cup \{f(t_1, \ldots, t_n) =^?_{\boldsymbol{\alpha}} g(s_1, \ldots, s_n)\}$, so that $\Gamma(P') = \Gamma(P)$ and $\lambda(P') = \lambda(P)$. □

Theorem 4 (Termination of QUnif). *For a given $(E_{\mathsf{ssh}}, \varepsilon)$-unification problem, the algorithm* QUnif *terminates either with the configuration* $\boldsymbol{F}$ *(indicating failure) or with a configuration of the form $V; C; \delta; \sigma$ (indicating success), where V is a set of unification equations between variables.*

Proof. A simple analysis of the rules of QUnif shows that all terminal configurations are of the form described above. In order to prove that the algorithm terminates, first note that the Cla and CCh rules terminate the derivation immediately, so it suffices to show that the remaining rules cannot yield an infinite derivation. For this purpose, we consider the measures λ, n_2 and n_3, where n_2 is the size of P and n_3 is the number of equations of the form $t =^?_{\boldsymbol{\alpha}} x$ in P such that

t is a non-variable term. By Lemma 3, L-Sub decreases λ while all other rules do not increase it; Dec and Tri decrease n_2, and Ori decreases n_3 while leaving n_2 invariant. Hence, the lexicographical combination of λ with n_2 and n_3 yields a measure that strictly decreases upon each of the aforementioned rules with respect to a well-founded order, thus proving termination. ☐

Proceeding now to the soundness and completeness proofs for QUnif, we fix a notion of solution of a configuration.

Definition 9 (Solution of a configuration). *A substitution τ is a* solution *of the configuration $P; \zeta \precsim \boldsymbol{\alpha}_1 \otimes \boldsymbol{\alpha}_2 \otimes \cdots \otimes \boldsymbol{\alpha}_n; \delta; \sigma$ if there exists a function μ mapping metavariables to elements of Ω such that*

(S1) $\zeta \precsim \mu(\boldsymbol{\alpha}_1) \otimes \mu(\boldsymbol{\alpha}_2) \otimes \cdots \otimes \mu(\boldsymbol{\alpha}_n)$ is valid,
(S2) $\mu(\boldsymbol{\beta}) \Vdash s\tau =_E t\tau$ holds for every equation $s =^?_{\boldsymbol{\beta}} t$ in P.
(S3) $x\tau = x\sigma\tau$ (syntactic equality) holds for every variable $x \in dom(\sigma)$.

The configuration $\boldsymbol{F}$ has no solutions.

This definition is compatible with Definition 4 in the following sense:

Lemma 4. *Let $\varepsilon \in \Omega$. A substitution τ is an (E, ε)-unifier of t and s if and only if τ is a solution for the corresponding initial configuration $\{t =^?_{\boldsymbol{\alpha}} s\}; \varepsilon \precsim \boldsymbol{\alpha}; \kappa; Id$.*

Proof. By definition, τ solves $\{t =^?_{\boldsymbol{\alpha}} s\}; \varepsilon \precsim \boldsymbol{\alpha}; \kappa; Id$ iff there exists μ such that $\varepsilon \precsim \mu(\boldsymbol{\alpha})$ and $\mu(\boldsymbol{\alpha}) \Vdash t\tau =_E s\tau$, which is equivalent to $\varepsilon \Vdash t\tau =_E s\tau$. ☐

The lemmas below are needed to show soundness and completeness of QUnif.

Lemma 5. *Let $\varepsilon \in \Omega$. Then $\varepsilon \Vdash f(t_1, \ldots, t_n) =_{E_{\mathsf{ssh}}} g(s_1, \ldots, s_m)$ if and only if $n = m$ and there exist $\varepsilon_1, \ldots, \varepsilon_n \in \Omega$ such that (i) $\varepsilon_i \Vdash t_i =_{E_{\mathsf{ssh}}} s_i$ for every $i \in \{1, \ldots, n\}$; (ii) $\varepsilon \precsim \varepsilon_1 \otimes \cdots \otimes \varepsilon_n \otimes \mathfrak{d}_{E_{\mathsf{ssh}}}(f, g)$.*

Proof. Write t for $f(t_1, \ldots, t_n)$ and s for $g(s_1, \ldots, s_m)$ as shorthands. We proceed by induction on the length of the proof for $\varepsilon \Vdash s =_{E_{\mathsf{ssh}}} t$ (using the rules in Figure 1). By definition, no non-empty statement has a proof of length 0. So suppose that the statement above holds for any equation between non-variable terms that has a proof of length k, and suppose we have a proof of length $k+1$ for $\varepsilon \Vdash s =_{E_{\mathsf{ssh}}} t$. We distinguish cases depending on the last rule that has been applied.

- If $\varepsilon \Vdash t =_{E_{\mathsf{ssh}}} s$ has been obtained via Ax, then we have $m = n$, $t_i = s_i = x_i$ for $1 \leqslant i \leqslant n$ and $\varepsilon \Vdash f(x_1, \ldots, x_n) = g(x_1, \ldots, x_n) \in E_{\mathsf{ssh}}$, so $\varepsilon \precsim \mathfrak{d}_{E_{\mathsf{ssh}}}(f, g)$. Taking $\varepsilon_i \coloneqq \kappa$ for $1 \leqslant i \leqslant n$, condition (i) is satisfied by Refl, and condition (ii) holds as $\varepsilon_1 \otimes \cdots \otimes \varepsilon_n \otimes \mathfrak{d}_{E_{\mathsf{ssh}}}(f, g) = \mathfrak{d}_{E_{\mathsf{ssh}}}(f, g) \succsim \varepsilon$.
- If $\varepsilon \Vdash t =_{E_{\mathsf{ssh}}} s$ has been obtained via Refl, then $f = g$ and $t_i = s_i$ for $1 \leqslant i \leqslant n$, so $\mathfrak{d}_{E_{\mathsf{ssh}}}(f, g) = \kappa$, and as in the previous case, one can take $\varepsilon_i = \kappa$ for every i.

– If $\varepsilon \Vdash t =_{E_{\mathsf{ssh}}} s$ has been obtained via Sym, then there is a proof of $\varepsilon \Vdash g(s_1,\dots,s_m) =_{E_{\mathsf{ssh}}} f(t_1,\dots,t_n)$ of length k, so by induction hypothesis, $m = n$ and there exist $\varepsilon_1,\dots,\varepsilon_n$ such that $\varepsilon \precsim \varepsilon_1 \otimes \cdots \otimes \varepsilon_n \otimes \mathfrak{d}_{E_{\mathsf{ssh}}}(f,g)$ and $\varepsilon_i \Vdash s_i =_{E_{\mathsf{ssh}}} t_i$ for every i. Thus, $\varepsilon_i \Vdash t_i =_{E_{\mathsf{ssh}}} s_i$ by Sym, and condition (ii) holds as $\mathfrak{d}_{E_{\mathsf{ssh}}}(f,g) = \mathfrak{d}_{E_{\mathsf{ssh}}}(g,f)$.

– If $\varepsilon \Vdash t =_{E_{\mathsf{ssh}}} s$ has been obtained via Trans, say from $\zeta \Vdash t =_{E_{\mathsf{ssh}}} r$ and $\eta \Vdash r =_{E_{\mathsf{ssh}}} s$, where $\zeta \otimes \eta = \varepsilon$, then r cannot be a variable as E_{ssh} is collapse-free. Writing $r = h(r_1,\dots,r_l)$, by induction hypothesis, $n = r = m$ and there exist $\zeta_1,\dots,\zeta_n,\eta_1,\dots,\eta_n$ such that $\zeta_i \Vdash t_i =_{E_{\mathsf{ssh}}} r_i$, $\eta_i \Vdash r_i =_{E_{\mathsf{ssh}}} s_i$ for all i and $\zeta \precsim \zeta_1 \otimes \dots \zeta_n \otimes \mathfrak{d}_{E_{\mathsf{ssh}}}(f,h)$, $\eta \precsim \eta_1 \otimes \dots \eta_n \otimes \mathfrak{d}_{E_{\mathsf{ssh}}}(h,g)$ Taking $\varepsilon_i := \zeta_i \otimes \eta_i$, condition (i) holds by Trans. As for condition (ii), note that we have $\mathfrak{d}_{E_{\mathsf{ssh}}}(f,g) \succsim \mathfrak{d}_{E_{\mathsf{ssh}}}(f,h) \otimes \mathfrak{d}_{E_{\mathsf{ssh}}}(h,g)$, and thus, $\varepsilon_1 \otimes \cdots \otimes \varepsilon_n \otimes \mathfrak{d}_{E_{\mathsf{ssh}}}(f,g) \succsim \zeta \otimes \eta = \varepsilon$.

– If $\varepsilon \Vdash t =_{E_{\mathsf{ssh}}} s$ has been obtained via NExp, then $f = g$ and we have $\delta_1,\dots,\delta_n \in \Omega$ such that $\delta_1 \otimes \cdots \otimes \delta_n$ and $\delta_i \Vdash t_i =_{E_{\mathsf{ssh}}} s_i$ for $1 \leqslant i \leqslant n$. As moreover $\mathfrak{d}_{E_{\mathsf{ssh}}}(f,g) = \kappa$, on can take $\varepsilon_i := \delta_i$.

– If $\varepsilon \Vdash t =_{E_{\mathsf{ssh}}} s$ has been obtained via Subst, say from $\varepsilon \Vdash t' =_{E_{\mathsf{ssh}}} s'$, then neither t' nor s' can be variables, so they are of the form $t' = f(t'_1,\dots,t'_n)$ and $s' = g(s'_1,\dots,s'_m)$. By induction hypothesis, $n = m$ and there $\varepsilon_1,\dots,\varepsilon_n$ which satisfy (ii), and, using Subst, also (i).

– If $\varepsilon \Vdash t =_{E_{\mathsf{ssh}}} s$ has been obtained via Ord, say from $\delta \Vdash t =_{E_{\mathsf{ssh}}} s$, where $\varepsilon \precsim \delta$, then induction hypothesis yields the desired elements $\varepsilon_1,\dots,\varepsilon_n$ of Ω.

– If $\varepsilon \Vdash t =_{E_{\mathsf{ssh}}} s$ has been obtained via Join, say from $\delta_1 \Vdash t =_{E_{\mathsf{ssh}}} s, \dots, \delta_l \Vdash t =_{E_{\mathsf{ssh}}} s$, where $\varepsilon = \bigvee_i \delta_i$ then for each i, there are $\zeta^i_1,\dots,\zeta^i_l$ such that $\zeta^i_j \Vdash t_i =_{E_{\mathsf{ssh}}} s_i$ holds for $1 \leqslant i \leqslant n$, $1 \leqslant j \leqslant l$. Setting $\varepsilon_i := \bigvee_{j=1}^n \zeta^j_i$, we have $\varepsilon_i \Vdash t_i =_{E_{\mathsf{ssh}}} s_i$ by Join, verifying (i). Moreover, writing $\omega = \mathfrak{d}_{E_{\mathsf{ssh}}}(f,g)$, we have

$$\varepsilon_1 \otimes \cdots \otimes \varepsilon_n \otimes \omega = (\bigotimes_{i=1}^{n}(\bigvee_{j=1}^{l} \zeta^j_i)) \otimes \omega = \bigvee_{j=1}^{l}(\bigotimes_{i=1}^{n} \zeta^j_i \otimes \omega) \succsim \delta_1 \vee \cdots \vee \delta_n = \varepsilon$$

by distributivity of Ω, yielding (ii).

For the converse, if $\varepsilon_1,\dots,\varepsilon_n$ with properties (i) and (ii) are given, then we have $\varepsilon_1 \otimes \cdots \otimes \varepsilon_n \Vdash f(t_1,\dots,t_n) = f(s_1,\dots,s_n)$ by (NExp) and $\mathfrak{d}_{E_{\mathsf{ssh}}}(f,g) \Vdash f(s_1,\dots,s_n) =_{E_{\mathsf{ssh}}} g(s_1,\dots,s_n)$ by (Subst). As $\varepsilon \precsim \varepsilon_1 \otimes \cdots \otimes \varepsilon_n \otimes \mathfrak{d}_{E_{\mathsf{ssh}}}(f,g)$, (Trans) and (Ord) yield $\varepsilon \Vdash t =_{E_{\mathsf{ssh}}} s$. □

Lemma 6. *Let $\mathfrak{C}$ be an admissible configuration.*

(i) If $\mathfrak{C} \Longrightarrow \mathfrak{C}'$ and τ solves $\mathfrak{C}'$, then τ solves $\mathfrak{C}$.
(ii) If τ solves $\mathfrak{C}$, then either $\mathfrak{C}$ is terminal, or there exist a configuration $\mathfrak{C}'$ and a substitution τ' such that $\mathfrak{C} \Longrightarrow \mathfrak{C}'$, $\tau'|_{dom(\tau)} = \tau$ and τ' solves $\mathfrak{C}'$.

Proof. For the proof of (i), suppose that τ solves $\mathfrak{C}'$, and let μ be a witness. We distinguish cases according to the rule that has been applied to obtain $\mathfrak{C}'$ from $\mathfrak{C}$. For each rule, we may assume that $\mathfrak{C}$ and $\mathfrak{C}'$ have the form described in the definition of the rule.

- If $\mathfrak{C}'$ has been obtained via Tri, say by removing the equation $t =^?_{\boldsymbol{\alpha}} t$, then we can define $\tilde{\mu}$ by setting $\tilde{\mu}(\boldsymbol{\alpha}) := \kappa$ and $\tilde{\mu}(\boldsymbol{\beta}) := \mu(\boldsymbol{\beta})$ for $\boldsymbol{\beta} \neq \boldsymbol{\alpha}$. Then $\tilde{\mu}$ witnesses that τ solves $\mathfrak{C}$.
- If $\mathfrak{C}'$ has been obtained via Dec, then we define $\tilde{\mu}$ by $\tilde{\mu}(\boldsymbol{\alpha}) := \mu(\boldsymbol{\beta}_1) \otimes \cdots \otimes \mu(\boldsymbol{\beta}_n) \otimes \omega$ and $\tilde{\mu}(\boldsymbol{\gamma}) := \mu(\boldsymbol{\gamma})$ for $\boldsymbol{\gamma} \neq \boldsymbol{\alpha}$. Then $\tilde{\mu}$ satisfies the constraint of $\mathfrak{C}$. Moreover, by Lemma 5, $\tilde{\mu}(\boldsymbol{\alpha}) \Vdash f(t_1, \ldots, t_n)\tau =_{E_{\mathsf{ssh}}} g(s_1, \ldots, s_n)\tau$, so $\tilde{\mu}$ witnesses that τ solves $\mathfrak{C}$.
- If $\mathfrak{C}'$ has been obtained via Cla or CCh, then there is nothing to show as $\mathbf{F}$ does not have a solution.
- If $\mathfrak{C}'$ has been obtained via L-Sub, then note that $\mathfrak{C}'$ can also be obtained from the configuration $P\rho \cup \{g(x_1, \ldots, x_n) =^?_{\boldsymbol{\alpha}} f(s_1, \ldots, s_n)\}; \zeta \precsim \boldsymbol{\alpha} \otimes \Delta; \delta; \sigma\rho$ via Dec. Thus, as we have already shown, τ is also a solution of this configuration, which we denote by $\mathfrak{C}''$.

  We now claim that μ witnesses that τ is a solution of $\mathfrak{C}$. We check the conditions in Definition 9. (S1) is trivial. For (S2), note that $x\tau = x\sigma\rho\tau = g(x_1, \ldots, x_n)\tau = x\rho\tau$ (as $x \notin dom(\sigma)$ by admissibility. Since $dom(\rho) = \{x\}$, this means that $\rho\tau = \tau$. Since τ solves $\mathfrak{C}''$, we have $\mu(\boldsymbol{\alpha}) \Vdash x\rho\tau =_{E_{\mathsf{ssh}}} f(s_1, \ldots, s_n)\rho\tau$, and therefore $\mu(\boldsymbol{\alpha}) \Vdash x\tau =_{E_{\mathsf{ssh}}} f(s_1, \ldots, s_n)\tau$. Similarly, for every equation $l =^?_{\boldsymbol{\beta}} r$ in P, we have $\mu(\boldsymbol{\alpha}) \Vdash l\rho\tau =_{E_{\mathsf{ssh}}} r\rho\tau$, and therefore, $\mu(\boldsymbol{\alpha}) \Vdash l\tau =_{E_{\mathsf{ssh}}} r\tau$. For (S3), it needs to be verified that $y\tau = y\sigma\tau$ holds for every $y \in dom(\sigma)$. But as τ solves $\mathfrak{C}'$, we know that $y\tau = y\sigma\rho\tau$ holds for every $y \in dom(\sigma\rho) = dom(\sigma) \cup \{x\}$, so we can conclude as we have already established that $\rho\tau = \tau$.
- If $\mathfrak{C}'$ has been obtained via Ori, then $\mu(\boldsymbol{\alpha}) \Vdash x\tau =_{E_{\mathsf{ssh}}} t\tau$ implies $\mu(\boldsymbol{\alpha}) \Vdash t\tau =_{E_{\mathsf{ssh}}} x\tau$, so μ witnesses that τ solves $\mathfrak{C}$.

For the proof of (ii), suppose that τ solves $\mathfrak{C}$, where $\mathfrak{C}$ is not terminal, and let μ be a witness. We distinguish cases according to the rule that can be applied to $\mathfrak{C}$. As before, we assume for each rule that $\mathfrak{C}$ has the form described in the definition of that rule.

- If $\mathfrak{C}'$ is obtained from $\mathfrak{C}$ via Tri, then μ witnesses that τ solves $\mathfrak{C}'$.
- If $\mathfrak{C}'$ is obtained from $\mathfrak{C}$ via Dec, then we have $\mu(\boldsymbol{\alpha}) \Vdash f(t_1, \ldots, t_n)\tau =_{E_{\mathsf{ssh}}} g(s_1, \ldots, s_n)\tau$, so by Lemma 5, there exist $\varepsilon_1, \ldots, \varepsilon_n \in \Omega$ such that $\varepsilon_i \Vdash s_i\tau =_{E_{\mathsf{ssh}}} t_i\tau$ for $1 \leqslant i \leqslant n$ and $\mu(\boldsymbol{\alpha}) \precsim \varepsilon_1 \otimes \cdots \otimes \varepsilon_n \otimes \mathfrak{d}_{E_{\mathsf{ssh}}}(f, g)$. Extending μ to $\tilde{\mu}$ by setting $\tilde{\mu}(\boldsymbol{\beta}_i) := \varepsilon_i$, we obtain a map that witnesses that τ solves $\mathfrak{C}'$.
- If $\mathfrak{C}' = \mathbf{F}$ is obtained from $\mathfrak{C}$ via Cla, then we have $\zeta \not\precsim \mathfrak{d}_{E_{\mathsf{ssh}}}(f, g)$. Therefore, there cannot exist $\varepsilon_1, \ldots, \varepsilon_n \in \Omega$ such that $\mu(\boldsymbol{\alpha}) \precsim \varepsilon_1 \otimes \cdots \otimes \varepsilon_n \otimes \omega$, as $\zeta \precsim \mu(\boldsymbol{\alpha})$. Thus, by Lemma 5, we have $\mu(\boldsymbol{\alpha}) \not\Vdash f(s_1, \ldots, s_n)\tau =_{E_{\mathsf{ssh}}} g(t_1, \ldots, t_n)\tau$, contradicting the assumption that τ is a solution of $\mathfrak{C}$.
- If L-Sub can be applied to $\mathfrak{C}$, then we have $\mu(\boldsymbol{\alpha}) \Vdash x\tau =_{E_{\mathsf{ssh}}} f(s_1, \ldots, s_n)\tau$. Since E_{ssh} is collapse-free, $x\tau$ cannot be a variable. So we can write $x\tau = g(t_1, \ldots, t_m)$. By Lemma 5, we have $m = n$ and $\zeta \precsim \mathfrak{d}_{E_{\mathsf{ssh}}}(f, g)$. We define τ' by setting $\tau'(x_i) := t_i$ and $\tau'(y) := \tau(y)$ for $y \in dom(\tau)$. Then μ witnesses that τ' solves the configuration $\mathfrak{C}'' := (P \cup \{x =^?_{\boldsymbol{\alpha}} f(t_1, \ldots, t_n)\})\rho; C; \delta; \sigma\rho$

obtained from $\mathfrak{C}$ by applying the substitution $\rho = \{x \mapsto f(x_1, \ldots, x_n)\}$. Since $\mathfrak{C}'$ can be obtained from $\mathfrak{C}''$ via a Dec step, τ' also solves $\mathfrak{C}'$.

- If CCh can be applied to $\mathfrak{C}$, then P contains an occurrence cycle, say $\{x_1 \approx^?_{\boldsymbol{\alpha}_1} t_1, \ldots, x_m \approx^?_{\boldsymbol{\alpha}_m} t_m\} \subseteq P$, where not all t_i are variables, $x_i \in \mathcal{V}(t_{i-1})$ for $1 < i \leqslant m$ and $x_1 \in \mathcal{V}(t_m)$. Then we have $\mu(\boldsymbol{\alpha}_1) \otimes \cdots \otimes \mu(\boldsymbol{\alpha}_m) \Vdash x_1\tau =_{E_{\mathsf{ssh}}} (t_1\{x_2 \mapsto t_2\} \ldots \{x_m \mapsto t_m\})\tau$, and $x_1\tau$ appears in the right-hand side in a non-trivial position. Thus, E_{ssh} entails a subterm-collapse equation, contradicting the fact that E_{ssh} is a simple equational theory.
- If $\mathfrak{C}'$ can be obtained from $\mathfrak{C}$ via Ori, then μ witnesses that τ solves $\mathfrak{C}'$. □

Theorem 5 (Soundness and completeness of QUnif). *Consider an $(E_{\mathsf{ssh}}, \varepsilon)$-unification problem between terms t and s.*

Soundness: *If* QUnif *terminates in a configuration $V; C; \delta; \sigma$ starting from the initial configuration $\{t =^?_{\boldsymbol{\alpha}} s\}; \varepsilon \precsim \boldsymbol{\alpha}; \kappa; Id$, then any solution of $V; C; \delta; \sigma$ is an $(E_{\mathsf{ssh}}, \varepsilon)$-unifier of t and s.*

Completeness: *If τ is an (E, ε)-unifier of t and s, then there is a run of* QUnif *starting from the initial configuration $\{t =^?_{\boldsymbol{\alpha}} s\}; \varepsilon \precsim \boldsymbol{\alpha}; \kappa; Id$ that terminates in a configuration $\{x_1 =_{\boldsymbol{\alpha}_1} y_1, \ldots, x_n =_{\boldsymbol{\alpha}_n} y_n\}; \zeta \precsim \boldsymbol{\alpha}_1 \otimes \cdots \otimes \boldsymbol{\alpha}_n; \delta; \sigma$ such that there exist a substitution φ and a map μ satisfying the following conditions:*

(i) $\zeta \precsim \mu(\boldsymbol{\alpha}_1) \otimes \cdots \otimes \mu(\boldsymbol{\alpha}_n)$;
(ii) $\mu(\boldsymbol{\alpha}_i) \Vdash x_i\varphi =_{E_{\mathsf{ssh}}} y_i\varphi$ for all $1 \leqslant i \leqslant n$;
(iii) $x\sigma\varphi = x\tau$ for all $x \in \mathcal{V}(s, t)$.

Proof. For soundness, suppose that QUnif produces a derivation $\mathfrak{C}_0 \Longrightarrow \ldots \Longrightarrow \mathfrak{C}_m$, where $\mathfrak{C}_0$ is the initial configuration $\{t =^?_{\boldsymbol{\alpha}} s\}; \varepsilon \precsim \boldsymbol{\alpha}; \kappa; Id$ and $\mathfrak{C}_m$ is a terminal configuration given by $V; S; \delta; \sigma$. If τ is a solution of $\mathfrak{C}_m$ then τ is also a solution of $\mathfrak{C}_0$ (by Lemma 6(i)), and therefore, τ is an $(E_{\mathsf{ssh}}, \varepsilon)$-unifier of t and s (by Lemma 4).

For completeness, suppose that τ is an (E, ε)-unifier of t and s. Then τ solves the corresponding initial configuration $\mathfrak{C}_0$ (by Lemma 4). If $\mathfrak{C}_0$ is not terminal, then there exists a rule application $\mathfrak{C}_0 \Longrightarrow \mathfrak{C}_1$ and a substitution τ_1 such that $\tau_1|_{dom(\tau)} = \tau$ and τ_1 solves $\mathfrak{C}_1$ (by Lemma 6(ii)). Iterating this argument, we obtain a derivation $\mathfrak{C}_0 \Longrightarrow \mathfrak{C}_1 \Longrightarrow \ldots$ and a sequence of substitutions $\tau, \tau_1, \ldots$. After a finite number of steps, this derivation reaches a terminal configuration $\mathfrak{C}_m$ by Theorem 4, and with it, we obtain a solution τ_m such that $\tau_m|_{\mathcal{V}(s,t)} = \tau$. Since τ_m solves $\mathfrak{C}_m$, there exist φ and μ satisfying (i) and (ii), as well as $x\sigma\varphi = x\tau_m$ for all $x \in \mathcal{V}(\mathfrak{C}_m)$, yielding (iii). □

Remark 3. In particular, a κ-complete set of solutions for the problem $t =^?_{\varepsilon} s$ can be obtained by determining for every terminal configuration obtained via QUnif the set of substitutions that meet conditions (i)-(iii) above. If one is just interested in finding some solution, it suffices to compute a terminal configuration $V; \zeta \precsim \Delta; \delta; \sigma$ and compose σ with a substitution that maps all variables in V to a fresh variable. The value of δ corresponds to the "degree" to which such a solution τ solves the unification problem, i.e. $\delta = (t\tau =^{\bullet}_{E_{\mathsf{ssh}}} s\tau)$.

Example 4. Consider the unification problem $f(y, g(x,x)) =^?_{E,\varepsilon} g(f(c,a), y)$, where $\Omega = \mathbb{L}$, $E = \{1 \Vdash a \approx b, 1 \Vdash b \approx c, 1 \Vdash f(x_1,x_2) \approx g(x_1,x_2)\}$ and $\varepsilon = 5$. The following derivation can be obtained by QUNIF:

$$\{f(y, g(x,x)) =^?_{\boldsymbol{\alpha}} g(f(c,a), y)\};\ 5 \geqslant \boldsymbol{\alpha};\ 0;\ Id$$
$$\Longrightarrow_{\mathsf{Dec}} \{y =^?_{\boldsymbol{\beta}_1} f(c,a),\ g(x,x) =^?_{\boldsymbol{\beta}_2} y\};\ 4 \geqslant \boldsymbol{\beta}_1 + \boldsymbol{\beta}_2;\ 1;\ Id$$
$$\Longrightarrow^{y \mapsto f(z_1,z_2)}_{\mathsf{L\text{-}Sub}} \{z_1 =^?_{\boldsymbol{\gamma}_1} c,\ z_2 =^?_{\boldsymbol{\gamma}_2} a,\ g(x,x) =^?_{\boldsymbol{\beta}_2} f(z_1,z_2)\};$$
$$4 \geqslant \boldsymbol{\gamma}_1 + \boldsymbol{\gamma}_2 + \boldsymbol{\beta}_2;\ 1;\ \{y \mapsto f(z_1,z_2)\}$$
$$\Longrightarrow^{z_1 \mapsto b}_{\mathsf{L\text{-}Sub}} \{\ z_2 =^?_{\boldsymbol{\gamma}_2} a,\ g(x,x) =^?_{\boldsymbol{\beta}_2} f(b,z_2)\};$$
$$3 \geqslant \boldsymbol{\gamma}_2 + \boldsymbol{\beta}_2;\ 2;\ \{y \mapsto f(b,z_2), z_1 \mapsto b\}$$
$$\Longrightarrow_{\mathsf{Dec}} \{z_2 =^?_{\boldsymbol{\gamma}_2} a,\ x =^?_{\boldsymbol{\delta}_1} b,\ x =^?_{\boldsymbol{\delta}_2} z_2\};$$
$$2 \geqslant \boldsymbol{\gamma}_2 + \boldsymbol{\delta}_1 + \boldsymbol{\delta}_2;\ 3;\ \{y \mapsto f(b,z_2), z_1 \mapsto b\}$$
$$\Longrightarrow^{z_2 \mapsto a}_{\mathsf{L\text{-}Sub}} \{x =^?_{\boldsymbol{\delta}_1} b,\ x =^?_{\boldsymbol{\delta}_2} a\};\ 2 \geqslant \boldsymbol{\delta}_1 + \boldsymbol{\delta}_2;\ 3;\ \{y \mapsto f(b,a), z_1 \mapsto b, z_2 \mapsto a\}$$
$$\Longrightarrow^{x \mapsto a}_{\mathsf{L\text{-}Sub}} \{a =^?_{\boldsymbol{\delta}_2} a\};\ 1 \geqslant \boldsymbol{\delta}_2;\ 4;\ \{y \mapsto f(b,a), z_1 \mapsto b, z_2 \mapsto a, x \mapsto a\}$$
$$\Longrightarrow_{\mathsf{Tri}} \emptyset; 1 \geqslant 0;\ 4;\ \{y \mapsto f(b,a), z_1 \mapsto b, z_2 \mapsto a, x \mapsto a\}$$

This leads to the solution $\{y \mapsto f(b,a), x \mapsto a\}$ (with degree 4). Further solutions can be obtained via different choices in the Subst steps. ■

Example 5. Consider $\Omega = \mathbb{L}$, $E = \{1 \Vdash f(x,y) \approx g(x,y)\}$, and the E-unification problem $g(a,x) =^?_3 f(y, g(b,z))$. A derivation of QUNIF is given below.

$$\{g(a,x) =^?_{\boldsymbol{\alpha}} f(y, g(b,z))\};\ 3 \geqslant \boldsymbol{\alpha};\ 0;\ Id$$
$$\Longrightarrow_{\mathsf{Dec}} \{a =^?_{\boldsymbol{\beta}_1} y,\ x =^?_{\boldsymbol{\beta}_2} g(b,z)\};\ 2 \geqslant \boldsymbol{\beta}_1 + \boldsymbol{\beta}_2;\ 1;\ Id$$
$$\Longrightarrow_{\mathsf{Ori}} \{y =^?_{\boldsymbol{\beta}_1} a,\ x =^?_{\boldsymbol{\beta}_2} g(b,z)\};\ 2 \geqslant \boldsymbol{\beta}_1 + \boldsymbol{\beta}_2;\ 1;\ Id$$
$$\Longrightarrow^{y \mapsto a}_{\mathsf{L\text{-}Sub}} \{x =^?_{\boldsymbol{\beta}_2} g(b,z)\};\ 2 \geqslant \boldsymbol{\beta}_2;\ 1;\ \{y \mapsto a\}$$
$$\Longrightarrow^{x \mapsto f(x_1,x_2)}_{\mathsf{L\text{-}Sub}} \{x_1 =^?_{\boldsymbol{\gamma}_1} b,\ x_2 =^?_{\boldsymbol{\gamma}_2} z\};\ 1 \geqslant \boldsymbol{\gamma}_1 + \boldsymbol{\gamma}_2;\ 2;\ \{y \mapsto a, x \mapsto f(x_1,x_2)\}$$
$$\Longrightarrow^{x_1 \mapsto b}_{\mathsf{L\text{-}Sub}} \{x_2 =^?_{\boldsymbol{\gamma}_2} z\};\ 1 \geqslant \boldsymbol{\gamma}_2;\ 2;\ \{y \mapsto a, x \mapsto f(b,x_2), x_1 \mapsto b\}$$

The computed terminal configuration still contains equations between variables. For any ψ such that $1 \Vdash x_2\psi =_E z\psi$, the substitution $\{y \mapsto a, x \mapsto f(b,x_2)\}\psi$ is an (E,ε)-unifier of the given terms. In particular, unifiers that can be obtained from this configuration include, e.g., $\{y \mapsto a, x \mapsto f(b,u), z \mapsto u\}$, where u is a fresh variable (with degree 2), and also $\{y \mapsto a, x \mapsto f(b, f(a,a)), z \mapsto g(a,a)\}$ (with degree 3). ■

4.2 Idempotent quantales

Now we consider the case where Ω is idempotent. Under this hypothesis, we can strengthen our results and show that – with the right definitions – the unification problem is unitary, and that a simplified version of QUNIF computes a most general unifier of two given terms. For the fuzzy quantale $\Omega = \mathbb{I}_{\min} = ([0,1], \leqslant, \min)$, our algorithm coincides with Sessa's weak unification algorithm [22].

Note that in any integral idempotent quantale, meet and tensor coincide. As a consequence, in an idempotent quantale, $\alpha \precsim \beta$ implies $\beta \multimap \alpha = \alpha$.

Definition 10 (Weak mgu). *A substitution σ is a* weak most general (E,ε)-unifier of t and s, *denoted $wmgu_{E,\varepsilon}(t,s)$, if $\mathfrak{U}_{E,\varepsilon}(t,s) = \{\tau \mid \sigma \lessapprox_{E,\mathcal{V}(t,s),\varepsilon} \tau\}$.*

By Lemma 2 (ii), $\sigma = wmgu_{E,\varepsilon}(t,s)$ iff $\sigma \in \mathfrak{U}_{E,\varepsilon}(t,s)$ and $\sigma \lessapprox_{E,\mathcal{V}(t,s),\varepsilon} \tau$ holds for every $\tau \in \mathfrak{U}_{E,\varepsilon}(t,s)$; that is, iff $\{\sigma\} = mcsu_{E,\varepsilon,\varepsilon}(t,s)$.

In the idempotent setting, the rules L-Sub and CCh from QUnif can be replaced by simpler versions:

E-Sub: **Substitute (eager)**

$$\{x =^?_{\boldsymbol{\alpha}} s\} \uplus P; \zeta \precsim \boldsymbol{\alpha} \otimes \Delta; \delta; \sigma \Longrightarrow P\{x \mapsto s\}; \zeta \precsim \Delta; \delta; \sigma\{x \mapsto s\}, \text{ if } x \notin \mathcal{V}(s).$$

OCh: **Occurrence check**

$$\{x =^?_{\boldsymbol{\alpha}} s\} \uplus P; C; \delta; \sigma \Longrightarrow \mathbf{F}, \quad \text{if } x \in \mathcal{V}(s) \text{ and } s \neq x.$$

Note that both of these rules constitute steps that could also be achieved by the rules from QUnif: E-Sub can be viewed as a composition of L-Sub and Dec steps, and OCh is just a restricted version of CCh. As before, we use these rules to transform the initial configuration corresponding to a given $(E_{\mathsf{ssh}}, \iota)$-unification problem. As an output, we return Failure if $\mathbf{F}$ has been obtained, or σ if a terminal configuration $P; C; \delta; \sigma$ has been reached. We denote the resulting algorithm by QUnif-id.

In order to obtain a stronger completeness theorem than in the general case, we refine the notion of a solution of a configuration.

Definition 11 (ι-solution of a configuration). *Let $\iota \in \Omega$ be idempotent. A substitution τ is an ι-*solution *of the configuration $P; \zeta \precsim \boldsymbol{\alpha}_1 \otimes \boldsymbol{\alpha}_2 \otimes \cdots \otimes \boldsymbol{\alpha}_n; \delta; \sigma$ if there exists a function μ mapping metavariables to elements of Ω such that*

(ι1) $\zeta \precsim \mu(\boldsymbol{\alpha}_1) \otimes \mu(\boldsymbol{\alpha}_2) \otimes \cdots \otimes \mu(\boldsymbol{\alpha}_n)$ is valid,
(ι2) $\mu(\boldsymbol{\beta}) \Vdash t\tau =_E s\tau$ holds for every equation $t =^?_{\boldsymbol{\beta}} s$ in P.
(ι3) $\iota \Vdash x\tau =_E x\sigma\tau$ holds for every variable $x \in dom(\sigma)$.

*The configuration **F** has no solutions.*

Note that the only difference in comparison with Definition 9 is that (ι3) features a quantitative equality over E_{ssh}, whereas in (S3) we have a syntactic equality.

The lemmas below are needed in the proof of soundness and completeness of QUnif-id.

Lemma 7. *Let Ω be a (not necessarily idempotent) quantale, $\iota \in \Omega$ be an idempotent element of Ω, τ be a substitution, and t and s be terms.*

(i) *τ is an (E,ι)-unifier of t and s iff τ is an ι-solution for the corresponding initial configuration $\{t =^?_{\boldsymbol{\alpha}} s\}; \iota \precsim \boldsymbol{\alpha}; \kappa; Id$.*

(ii) *τ is an ι-solution for an admissible configuration of the form $\emptyset; C; \delta; \sigma$ iff $\sigma \lessapprox_{E,dom(\sigma),\iota} \tau$.*

Proof. As for (i), τ is an ι-solution of $\{t =^?_{\boldsymbol{\alpha}} s\}; \iota \precsim \boldsymbol{\alpha}; \kappa; Id$ if and only if there exists $\mu(\boldsymbol{\alpha}) \succsim \iota$ such that $\mu(\boldsymbol{\alpha}) \Vdash t\tau =_E s\tau$, which is equivalent to $\iota \Vdash t\tau =_E s\tau$.

For (ii), note that C is of the form $\zeta \precsim \kappa$ by admissibility. Thus, τ is an ι-solution of $\emptyset; C; \delta; \sigma$ iff $\iota \Vdash x\tau =_E x\sigma\tau$ holds for every variable $x \in dom(\sigma)$. In that case, τ witnesses that $\sigma \lessapprox_{E,dom(\sigma),\iota} \tau$. For the converse, suppose that $\sigma \lessapprox_{E,dom(\sigma),\iota} \tau$, and let φ be the witness. Then $\iota \Vdash x\sigma\varphi =_E x\tau$ for all $x \in dom(\sigma)$, and thus, by Lemma 2(i), $\iota \Vdash x\sigma\tau =_E x\sigma\sigma\varphi$, and since σ is idempotent (by admissibility), $x\sigma\sigma\varphi = x\sigma\varphi$. Hence, (Trans) and idempotence of ι yield $\iota \Vdash x\sigma\tau =_E x\tau$ for any $x \in dom(\sigma)$. □

Lemma 8. *Let $\mathfrak{Q}$ be an idempotent quantale, $\iota \in \Omega$, and $\mathfrak{C}$ be a configuration obtained from an (E_{ssh},ι)-unification problem in $\mathfrak{Q}$ by applying rules from* QUNIF-ID*. If $\mathfrak{C}'$ is obtained from $\mathfrak{C}$ by a rule from* QUNIF-ID*, then a substitution τ is an ι-solution of $\mathfrak{C}$ iff it is an ι-solution of $\mathfrak{C}'$.*

Proof (Sketch). If $\mathfrak{C}'$ is obtained from $\mathfrak{C}$ via Tri, Dec, Cla or Ori, then the proof of the statement above is analogous to the corresponding cases in the proof of Lemma 6. In the case of CCh, where $\mathfrak{C}$ is of the form $\{x =^?_{\boldsymbol{\alpha}} s\} \cup P; C; \delta; \sigma$ and s contains the variable x as a proper subterm, τ cannot be an ι-solution of $\mathfrak{C}$ because then $\mu(\boldsymbol{\alpha}) \Vdash x\tau =_{E_{\mathsf{ssh}}} s\tau$ would contradict simplicity of E_{ssh}.

It remains to verify the statement for E-Sub and OCh.

For E-Sub, first observe that in any configuration, the constraint C will be of the form $\iota \precsim \Delta$, because $\omega \multimap \iota = \iota$ whenever $\iota \precsim \omega$. Suppose now that τ is an ι-solution of $\{x =^?_{\boldsymbol{\alpha}} s\} \uplus P; \iota \precsim \boldsymbol{\alpha} \otimes \Delta; \delta; \sigma$. We define $\tilde{\mu}$ by $\tilde{\mu}(\boldsymbol{\beta}) := \iota$ for every metavariable $\boldsymbol{\beta}$ and claim that $\tilde{\mu}$ witnesses that τ is a solution to $P\{x \mapsto s\}; \iota \precsim \Delta; \sigma\{x \mapsto s\}$, the configuration obtained via eager substitution.

Condition (ι1) is satisfied as ι is idempotent. For (ι2), it needs to be shown that $\iota \Vdash l\{x \mapsto s\}\tau =_{E_{\mathsf{ssh}}} r\{x \mapsto s\}\tau$ holds for every equation $l =^?_{\boldsymbol{\beta}} r$ in P. We know that $\iota \Vdash x\tau =_{E_{\mathsf{ssh}}} s\tau$, and therefore, we have $\tau \lessapprox_{\iota,E_{\mathsf{ssh}},dom(\tau)\cup\{x\}} \{x \mapsto s\}\tau$. As τ is an $(\mathfrak{Q},\iota,E_{\mathsf{ssh}})$-unifier of l and r, Lemma 2(ii) yields that $\tau\{x \mapsto s\}$ is an $(\mathfrak{Q},\iota,E_{\mathsf{ssh}})$-unifier of l and r, too. This means that $\iota \Vdash l\{x \mapsto s\}\tau =_{E_{\mathsf{ssh}}} r\{x \mapsto s\}\tau$, as desired. For condition (ι3), it needs to be verified that $\iota \Vdash y\tau =_{E_{\mathsf{ssh}}} y\sigma\{x \mapsto s\}\tau$ holds for any $y \in dom(\sigma) \cup \{x\}$. Since $\iota \Vdash x\tau =_{E_{\mathsf{ssh}}} s\tau$, we have $\iota \Vdash y\tau = y\{x \mapsto s\}\tau$ for every variable y. Thus, Lemma 2(i) yields $\iota \Vdash y\sigma\tau =_{E_{\mathsf{ssh}}} y\sigma\{x \mapsto s\}\tau$ for any variable y. Since we know that $\iota \Vdash y\tau =_{E_{\mathsf{ssh}}} y\sigma\tau$ holds for $y \in dom(\sigma) \cup \{x\}$, we obtain $\iota \Vdash y\tau =_{E_{\mathsf{ssh}}} y\sigma\{x \mapsto s\}\tau$ for any $y \in dom(\sigma) \cup \{x\}$, as desired.

For the converse, suppose now that τ solves $P\{x \mapsto s\}; C; \delta; \sigma\{x \mapsto s\}$. As before, we claim that the constant map $\mu \colon \boldsymbol{\beta} \mapsto \kappa$ witnesses that τ solves $P; C; \sigma$. Condition (ι1) is satisfied by the idempotence of ι. Since we know that $\iota \Vdash y\tau =_{E_{\mathsf{ssh}}} y\sigma\{x \mapsto s\}\tau$ holds for any $x \in dom(\sigma) \cup \{x\}$, we have $\iota \Vdash x\tau =_{E_{\mathsf{ssh}}} s\tau$ as a special case, using that $x \notin dom(\sigma)$. Therefore, similar as before, we have

$\iota \Vdash y\sigma\tau = y\sigma\{x \mapsto s\}\tau$ for any variable y, and thus, (ι2) and (ι3) are obtained using Lemma 2.

As for OCh, it suffices to show that no configuration of the form $\{x =_{\boldsymbol{\alpha}} s\}; C; \delta; \sigma$, where x is a proper subterm of s, has an ι-solution; but this is clear as E_{ssh} is simple.

Theorem 6 (Soundness and completeness of QUnif-id). *Consider an $(E_{\mathsf{ssh}}, \iota)$-unification problem between terms t and s in an idempotent quantale Ω, where $\iota \in \Omega$. Any run of* QUnif-id *starting from $\{t =^?_{\boldsymbol{\alpha}} s\}; \boldsymbol{\alpha} \precsim \iota; \kappa; Id$ terminates and returns $wmgu_{E_{\mathsf{ssh}},\iota}(t, s)$ if it exists, or fails otherwise.*

Proof. Termination follows from termination of QUnif (Theorem 4). For soundness and completeness, by Lemma 7(i), a substitution τ is an $(E_{\mathsf{ssh}}, \iota)$-unifier of t and s iff it is an ι-solution of the initial configuration $\mathfrak{C}_0$. By Lemma 8, the latter holds iff τ is an ι-solution for any terminal configuration $\emptyset; C; \delta; \sigma$, or equivalently, iff $\sigma \lessapprox_{E_{\mathsf{ssh}}, \mathcal{V}(t,s), \iota} \tau$ (by Lemma 7(ii)), concluding the proof. □

Example 6. We demonstrate algorithm QUnif-id for the problem $f(x, c) =^?_{\mathbb{I},0.4,E} h(a, x)$ in the (idempotent) fuzzy quantale $\mathbb{I}$ with the min T-norm modulo $E = \{0.5 \Vdash a \approx b,\ 0.5 \Vdash b \approx c,\ 0.6 \Vdash f(x_1, x_2) \approx g(x_1, x_2),\ 0.7 \Vdash g(x_1, x_2) \approx h(x_1, x_2)\}$. A derivation of QUnif-id is shown below:

$$\{f(x, c) =^?_{\boldsymbol{\alpha}} h(a, x)\}; 0.4 \leqslant \boldsymbol{\alpha}; 1; Id$$
$$\Longrightarrow_{\mathsf{Dec}} \{x =^?_{\boldsymbol{\beta}_1} a,\ c =^?_{\boldsymbol{\beta}_2} x\}; 0.4 \leqslant \min(\boldsymbol{\beta}_1, \boldsymbol{\beta}_2); 0.6; Id$$
$$\Longrightarrow^{x \mapsto a}_{\mathsf{L\text{-}Sub}} \{c =^?_{\boldsymbol{\beta}_2} a\}; 0.4 \leqslant \boldsymbol{\beta}_2; 0.6; \{x \mapsto a\}$$
$$\Longrightarrow_{\mathsf{Dec}} \emptyset; 0.4 \leqslant 1; 0.5; \{x \mapsto a\}.$$

Choosing the other equation in the L-Sub step would lead to a different unifier $\{x \mapsto c\}$ with the same degree 0.5. The solution $\{x \mapsto b\}$ (with degree 0.5) is not computed. All three solutions are 0.5-equivalent. ■

5 Conclusion

In the quantitative setting, equality is replaced by its quantitative counterpart modeling the abstract notion of proximity between terms. A quantitative unification problem asks for finding a substitution that brings the given terms close to each other within a predefined range (with respect to this abstract proximity). However, unlike the standard unification, here it is not guaranteed that an instance of a unifier is still a unifier. The reason is that the instantiation is also quantitative, and it might move the more specific substitution "too far away" from a unifier of the given problem.

In studying quantitative unification, one has to address such and related challenges. We investigated the quantitative equational unification problem in Lawvereian quantales modulo theories presented by axioms of the form $\gamma \Vdash f(x_1, \ldots, x_n) \approx g(x_1, \ldots, x_n)$. Our notion of a minimal complete set of unifiers takes into account two (abstract) distances: between terms to be unified and between substitutions via instantiation. We showed that our unification problems in arbitrary Lawvereian quantales are finitary, while for idempotent Lawvereian quantales, they are unitary. The corresponding algorithms were developed and their properties were studied.

The equational theories that we considered here are a special case of simple shallow theories. An interesting future work would be to extend this work to a larger class of shallow theories (which have some desirable properties in the standard case [6]). Further, the related problem of disunification in Lawvereian quantales is worth investigating.

Acknowledgments. Supported by the Austrian Science Fund (FWF) under project P 35530 (SQUEE).

Disclosure of Interests. The authors have no competing interests to declare that are relevant to the content of this article.

References

1. de Amorim, A.A., Gaboardi, M., Hsu, J., Katsumata, S., Cherigui, I.: A semantic account of metric preservation. In: Castagna, G., Gordon, A.D. (eds.) Proceedings of the 44th ACM SIGPLAN Symposium on Principles of Programming Languages, POPL 2017, Paris, France, January 18-20, 2017. pp. 545–556. ACM (2017). https://doi.org/10.1145/3009837.3009890
2. Baader, F., Snyder, W.: Chapter 8 - Unification Theory. In: Robinson, A., Voronkov, A. (eds.) Handbook of Automated Reasoning, p. 445–533. North-Holland, Amsterdam (2001). https://doi.org/10.1016/B978-044450813-3/50010-2
3. Bacci, G., Mardare, R., Panangaden, P., Plotkin, G.D.: An algebraic theory of markov processes. In: Dawar, A., Grädel, E. (eds.) Proceedings of the 33rd Annual ACM/IEEE Symposium on Logic in Computer Science, LICS 2018, Oxford, UK, July 09-12, 2018. pp. 679–688. ACM (2018). https://doi.org/10.1145/3209108.3209177
4. Bacci, G., Mardare, R., Panangaden, P., Plotkin, G.D.: Quantitative Equational Reasoning. In: Barthe, G., Katoen, J.P., Silva, A. (eds.) Foundations of Probabilistic Programming, p. 333–360. Cambridge University Press (2020). https://doi.org/10.1017/9781108770750.011
5. Barthe, G., Katoen, J.P., Silva, A. (eds.): Foundations of Probabilistic Programming. Cambridge University Press (2020). https://doi.org/10.1017/9781108770750
6. Comon, H., Haberstrau, M., Jouannaud, J.P.: Syntacticness, cycle-syntacticness, and shallow theories. Inf. Comput. **111**(1), 154–191 (1994). https://doi.org/10.1006/INCO.1994.1043
7. Dal Lago, U., Gavazzo, F.: A relational theory of effects and coeffects. Proc. ACM Program. Lang. **6**(POPL), 1–28 (2022). https://doi.org/10.1145/3498692
8. Dubois, D., Prade, H.: Fuzzy sets and systems: theory and applications, Mathematics in science and engineering, vol. 144. Acad. Press (1980), https://www.worldcat.org/oclc/05726778

9. Dundua, B., Kutsia, T., Marin, M., Pau, I.: Constraint solving over multiple similarity relations. In: Ariola, Z.M. (ed.) 5th International Conference on Formal Structures for Computation and Deduction, FSCD 2020, June 29-July 6, 2020, Paris, France (Virtual Conference). LIPIcs, vol. 167, pp. 30:1–30:19. Schloss Dagstuhl - Leibniz-Zentrum für Informatik (2020). https://doi.org/10.4230/LIPICS.FSCD.2020.30
10. Gavazzo, F., Di Florio, C.: Quantitative and metric rewriting: Abstract, non-expansive, and graded systems. CoRR **abs/2206.13610** (2022). https://doi.org/10.48550/ARXIV.2206.13610
11. Gavazzo, F., Di Florio, C.: Elements of quantitative rewriting. Proc. ACM Program. Lang. **7**(POPL), 1832–1863 (2023). https://doi.org/10.1145/3571256
12. Geoffroy, G., Pistone, P.: A partial metric semantics of higher-order types and approximate program transformations. In: Baier, C., Goubault-Larrecq, J. (eds.) 29th EACSL Annual Conference on Computer Science Logic, CSL 2021, January 25-28, 2021, Ljubljana, Slovenia (Virtual Conference). LIPIcs, vol. 183, pp. 23:1–23:18. Schloss Dagstuhl - Leibniz-Zentrum für Informatik (2021). https://doi.org/10.4230/LIPICS.CSL.2021.23
13. Julián-Iranzo, P., Rubio-Manzano, C.: Proximity-based unification theory. Fuzzy Sets and Systems **262**, 21–43 (2015). https://doi.org/10.1016/j.fss.2014.07.006
14. Kutsia, T., Pau, C.: A framework for approximate generalization in quantitative theories. In: Blanchette, J., Kovács, L., Pattinson, D. (eds.) Automated Reasoning - 11th International Joint Conference, IJCAR 2022, Haifa, Israel, August 8-10, 2022, Proceedings. Lecture Notes in Computer Science, vol. 13385, pp. 578–596. Springer (2022). https://doi.org/10.1007/978-3-031-10769-6_34
15. Lawvere, F.W.: Metric spaces, generalized logic, and closed categories. Rendiconti del Seminario Matematico e Fisico di Milano **43**, 135–166 (1973)
16. Mardare, R., Panangaden, P., Plotkin, G.D.: Quantitative Algebraic Reasoning. In: Proceedings of the 31st Annual ACM/IEEE Symposium on Logic in Computer Science. p. 700–709. LICS '16, Association for Computing Machinery, New York, NY, USA (2016). https://doi.org/10.1145/2933575.2934518
17. Mardare, R., Panangaden, P., Plotkin, G.D.: On the axiomatizability of quantitative algebras. In: 32nd Annual ACM/IEEE Symposium on Logic in Computer Science, LICS 2017, Reykjavik, Iceland, June 20-23, 2017. pp. 1–12. IEEE Computer Society (2017). https://doi.org/10.1109/LICS.2017.8005102
18. Pau, C.: Symbolic Techniques for Approximate Reasoning. Ph.D. thesis, RISC, Johannes Kepler University Linz (2022)
19. Pau, C., Kutsia, T.: Proximity-based unification and matching for fully fuzzy signatures. In: 30th IEEE International Conference on Fuzzy Systems, FUZZ-IEEE 2021, Luxembourg, July 11-14, 2021. pp. 1–6. IEEE (2021). https://doi.org/10.1109/FUZZ45933.2021.9494438
20. Reed, J., Pierce, B.C.: Distance makes the types grow stronger: a calculus for differential privacy. In: Hudak, P., Weirich, S. (eds.) Proceeding of the 15th ACM SIGPLAN international conference on Functional programming, ICFP 2010, Baltimore, Maryland, USA, September 27-29, 2010. pp. 157–168. ACM (2010). https://doi.org/10.1145/1863543.1863568
21. Rosenthal, K.I.: Quantales and their applications. Pitman Research Notes in Mathematics, Longman Scientific & Technical (1990)
22. Sessa, M.I.: Approximate reasoning by similarity-based SLD resolution. Theor. Comput. Sci. **275**(1-2), 389–426 (2002). https://doi.org/10.1016/S0304-3975(01)00188-8