



HAL
open science

Computing most general unifiers in Euclidean modal logics

Quentin Gougeon

► **To cite this version:**

| Quentin Gougeon. Computing most general unifiers in Euclidean modal logics. 2024. hal-04244893v2

HAL Id: hal-04244893

<https://hal.science/hal-04244893v2>

Preprint submitted on 21 May 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution 4.0 International License

Computing most general unifiers in Euclidean modal logics

Quentin Gougeon

CNRS-INPT-UT3, Toulouse University, F-31062 Toulouse, France
quentin.gougeon@irit.fr

Abstract

We prove that all extensions of **K5** have unary unification, even with parameters. Our proof is constructive in the sense that we can effectively compute, in 4-exponential space, a most general unifier for any unifiable formula. In particular, this proves that unification and admissibility are decidable. We also investigate special unification types: we show that **K5** and **KD5** are transparent, and we characterize the projective extensions of **K5**.

Keywords: Unification. Duality theory. Euclidean relations.

1 Introduction

In propositional logics, the problem of unification asks whether an arbitrary formula can be turned into a tautology, by uniformly replacing its variables with other formulas. Formally, given a logic **L** and a formula φ , we call a *unifier* of φ any substitution σ such that $\sigma(\varphi)$ is a theorem of **L**. A central question in unification theory is that of determining the *type* of a unifiable formula φ , a concept that summarizes the structural information of its space of unifiers and, more specifically, that indicates to which extent it admits a ‘compact’ representation. The different possible types are unary, finitary, infinitary and nullary, and depend on the existence and cardinality of a minimal complete set of unifiers of φ . The best of all four types is the unary one, which means that every unifier of φ can be written as an instance of one specific unifier, called a *most general unifier* of φ . Some most general unifiers have additional properties that make them easier to identify, such as Wronski’s *transparent* unifiers [Wro95], and Ghilardi’s *projective unifiers* [Ghi97]. In all cases, the type of the logic **L** is defined as the worst type among all the types of its unifiable formulas.

Unification is also closely related to the problem of *admissibility*: an inference rule $\frac{\varphi_1 \dots \varphi_n}{\psi}$ is said to be admissible for **L** if every unifier of $\varphi_1, \dots, \varphi_n$ is also a unifier of ψ – which means, intuitively, that the rule turns theorem of **L** into theorems of **L**. If $\varphi_1 \wedge \dots \wedge \varphi_n$ admits a most general unifier σ , then checking that $\frac{\varphi_1 \dots \varphi_n}{\psi}$ is admissible amounts to check that σ is a unifier of ψ . For details we refer to [BG11, BS01, Ghi99].

Unification in modal logics is still poorly understood, despite being studied since decades, and having witnessed a number of milestones: in [Ghi00], Ghilardi proved that many transitive logics like **K4**, **S4** or **GL** are finitary; in [GS04], Ghilardi and Sacchetti proved that all extensions of **K4.2**⁺ are *filtering* and, in particular, unary; in [Jeř15], Jeřábek proved that **K** is nullary; in [Kos18], Kost proved that all the extensions of **K4D1** are projective. Yet, to this day, the type of standard logics such as **KD**, **KT** and **KB** remains unknown. In *parametric unification*, substitutions are

not allowed to affect some variables, called *parameters* (or *constants* in [BS01]). In this setting, the logics **KD**, **KT**, **KB**, **KDB**, **KTB** and many others turn out to be nullary [BG17, Bal19, BG20].

While results in unification theory are usually derived via syntactic methods, semantics-driven idea have also been applied with success. Those typically exploit a *duality* between modal algebras and (general) Kripke frames, from which one can derive a correspondence between unifiers and a certain type of bounded morphisms between canonical frames. Unification results can then be proven by relying solely on the relational properties of these frames, thus favoring visual intuition. This approach was used by Ghilardi for intuitionistic logic [Ghi04], and more recently by Balbiani and Gougeon for modal logics [BG22]. In this work, we aim to apply these techniques to the extensions of the logic **K5** := **K** + $\diamond x \rightarrow \Box \diamond x$. These logics present key strong points, listed below.

1. Their canonical frames are Euclidean, and thus have a very simple form.
2. They are popular and well understood since years, with a complete characterization worked out by Nagel and Thomason as early as the 80's [NT85].
3. They are all locally tabular. Duality-based techniques were found to work best for logics with this property, as it greatly simplifies the structure of their dual frames.

Unification for these logics has already been partially investigated in recent work, though in the non-parametric setting only. First, the associated decision problem is readily proven to be in NP, seeing that every variable-free formula is equivalent to one of \perp , \top , $\Box \perp$ or $\diamond \top$.¹ In [BG22], it was proven that the projective extensions of **K5** are exactly the extensions of **K45**. On the other hand, the authors of [AABM23] exhibited infinitely many unary logics between **K5** and **K45** (including **K5** itself), but left open the question of whether *all* extensions of **K5** are unary. In the present paper, we answer this question in the affirmative, even in the case of parametric unification. Even better, we exhibit an upper bound on the number of extra variables needed to construct a unifier. More precisely, we show that if φ contains n variables, then φ admits a complete set of unifiers that uses a number of variables 2-exponential in n . Combining this result with the fact that extensions of **K5** have filtering unification [AABM23], we show that one can effectively compute a most general unifier for φ in 4-exponential space. From this, we also deduce that unification and admissibility are decidable in 4-exponential space. Finally, we attempt to characterize the special types of concise, transparent and projective unification – where *concise* means that a unifiable formula always admits a most general unifier that does not introduce new variables. We show that the projective extensions of **K5** (in the parametric case) are exactly the extensions of the logic **Kt5** := **K5** + $x \wedge \diamond \top \rightarrow \diamond x$. We also exhibit infinitely many transparent Euclidean logics that are not projective – including **K5** and **KD5** – as well as infinitely many non-concise extensions of **K5**.

The paper is structured as follows. In Section 2, we introduce the necessary material. In Section 3, we investigate the problem dual to unification and proceed to break it down into ‘atomic’ sub-problems that are easier to address. In Section 4, we present our algorithm for computing a most general unifier. In Section 5, we deal with concise, transparent and projective unification. In Section 6, we discuss potential directions for future work. For readability, we have deferred some proofs to the appendix.

¹This kind of result is folklore, and the proof sketch goes as follows. We call a *ground substitution* a substitution that replaces variables by variable-free formulas only. We can see that if φ admits a unifier, then φ admits a ground unifier, obtained by composing σ on the left with any ground substitution. Hence, we can decide whether φ is unifiable by non-deterministically selecting a ground substitution σ , and checking whether $\sigma(\varphi)$ is valid.

2 Background

2.1 Modal logic

Let \mathcal{X} be a countable set of *variables*. We fix a set $\mathcal{P} \subseteq \mathcal{X}$ of *parameters*, such that both \mathcal{P} and $\mathcal{X} \setminus \mathcal{P}$ are infinite. If $X \subseteq \mathcal{X}$, we define the modal language \mathcal{L}_X over X by the following grammar:

$$\varphi ::= x \mid \neg\varphi \mid (\varphi \wedge \psi) \mid \Box\varphi$$

where $x \in X$. We write $\mathcal{L} := \mathcal{L}_X$. The abbreviations $\perp, \top, \varphi \vee \psi, \varphi \rightarrow \psi, \varphi \leftrightarrow \psi, \Diamond\varphi$ are defined as usual, and we follow the standard rules for omission of parentheses. If $\varphi \in \mathcal{L}$, we denote by $\text{Var}(\varphi)$ the set of all variables occurring in φ . The formula φ is said to be *parameter-free* if $\text{Var}(\varphi) \cap \mathcal{P} = \emptyset$. The *size* of φ is the number of symbols occurring in φ , and is denoted by $|\varphi|$. A *substitution* is a map $\sigma : \mathcal{L}_X \rightarrow \mathcal{L}_Y$ with $X, Y \subseteq \mathcal{X}$ finite, satisfying $\sigma(\neg\varphi) = \neg\sigma(\varphi)$, $\sigma(\varphi \wedge \psi) = \sigma(\varphi) \wedge \sigma(\psi)$ and $\sigma(\Box\varphi) = \Box\sigma(\varphi)$ for all $\varphi, \psi \in \mathcal{L}_X$. If $\sigma : \mathcal{L}_X \rightarrow \mathcal{L}_Y$ and $\tau : \mathcal{L}_Y \rightarrow \mathcal{L}_Z$ are two substitutions, the composition $\tau\sigma : \mathcal{L}_X \rightarrow \mathcal{L}_Z$ of σ and τ is the substitution defined by $\tau\sigma(\varphi) := \tau(\sigma(\varphi))$ for all $\varphi \in \mathcal{L}_X$.

A *normal modal logic* is a set $\mathbf{L} \subseteq \mathcal{L}$ of formulas containing the axioms and closed under the inference rules described in Table 1. An *extension* of \mathbf{L} is any normal modal logic \mathbf{L}' such that $\mathbf{L} \subseteq \mathbf{L}'$. Given $\Gamma \subseteq \mathcal{L}$, we denote by $\mathbf{L} + \Gamma$ the smallest extension of \mathbf{L} containing the elements of Γ . If Γ is a singleton $\{\theta\}$, we write $\mathbf{L} + \theta := \mathbf{L} + \{\theta\}$. For convenience, we let $\vdash_{\mathbf{L}} \varphi$ stand for $\varphi \in \mathbf{L}$. We then write $\varphi \equiv_{\mathbf{L}} \psi$ whenever $\vdash_{\mathbf{L}} \varphi \leftrightarrow \psi$. Given two substitutions $\sigma : \mathcal{L}_X \rightarrow \mathcal{L}_Y$ and $\tau : \mathcal{L}_X \rightarrow \mathcal{L}_Z$, we write $\sigma \equiv_{\mathbf{L}} \tau$ if $\sigma(x) \equiv_{\mathbf{L}} \tau(x)$ for all $x \in X$. Below we also define the relation of *global consequence* relatively to \mathbf{L} , adapted from [CZ97, Ch. 3].

All propositional tautologies $\Box(x \rightarrow y) \rightarrow (\Box x \rightarrow \Box y)$
From φ infer $\sigma(\varphi)$ From φ and $\varphi \rightarrow \psi$ infer ψ From φ infer $\Box\varphi$

Table 1: Axioms and rules of normal modal logics. Here σ is a substitution.

Definition 2.1. Let \mathbf{L} be a normal modal logic. Given $\varphi, \psi \in \mathcal{L}$, we write $\varphi \vdash_{\mathbf{L}} \psi$ if there exist $\theta_0, \dots, \theta_n \in \mathcal{L}$ such that $\theta_n = \psi$ and for all $i \in [0, n]$, one of the following holds:

- $\theta_i = \varphi$,
- $\theta_i \in \mathbf{L}$,
- there exist $j, k < i$ such that $\theta_k = \theta_j \rightarrow \theta_i$,
- there exists $j < i$ such that $\theta_i = \Box\theta_j$.

We denote by **K5** the smallest modal logic containing the axiom $\Diamond x \rightarrow \Box\Diamond x$. We will also consider the following extensions of **K5**:

$$\begin{aligned} \mathbf{KD5} &:= \mathbf{K5} + \Diamond\top, \\ \mathbf{K45} &:= \mathbf{K5} + \Diamond x \rightarrow \Diamond\Diamond x, \\ \mathbf{Kt5} &:= \mathbf{K5} + x \wedge \Diamond\top \rightarrow \Diamond x, \\ \mathbf{S5} &:= \mathbf{K5} + x \rightarrow \Diamond x. \end{aligned}$$

The axiom $x \wedge \diamond \top \rightarrow \diamond x$ is a weaker variation of the axiom $x \rightarrow \diamond x$, traditionally named \top [BdRV01]; this explains the name **Kt5**.

Finally, if \mathbf{L} is a normal modal logic, we call a substitution $\sigma : \mathcal{L}_X \rightarrow \mathcal{L}_Y$ *parametric* (with respect to \mathbf{L}) if for all $p \in X \cap \mathcal{P}$ we have $p \in Y$ and $\sigma(p) \equiv_{\mathbf{L}} p$.²

2.2 Kripke frames and duality

A *Kripke frame* is a pair $\mathfrak{F} = (W, R)$ with W a non-empty set of *points* and $R \subseteq W^2$ a binary relation. Abusing notations, we will often identify \mathfrak{F} to the set W . We define the relation R^n on W by induction on n as follows:

- wR^0u iff $w = u$, and
- $wR^{n+1}u$ if there exists $v \in W$ such that wRv and vR^nu .

We then write $wR^{\leq n}u$ if wR^ku for some $k \leq n$.

A *generated subframe* of \mathfrak{F} is a frame $\mathfrak{F}' = (W', R')$ such that $W' \subseteq W$, $R' = R \cap W'^2$, and $w \in W'$ and wRu implies $u \in W'$. In this case \mathfrak{F} is said to be *generated* by W' . The frame \mathfrak{F}' is called a *closed subframe* of \mathfrak{F} if in addition, $w \in W'$ and uRw implies $u \in W'$. If $\mathfrak{F} = (W, R)$ and $\mathfrak{F}' = (W', R')$ are two Kripke frames, a *bounded morphism* from \mathfrak{F} to \mathfrak{F}' is a map $f : \mathfrak{F} \rightarrow \mathfrak{F}'$ such that:

- if wRu then $f(w)R'f(u)$, *(forward condition)*
- if $f(w)R'u'$ then there exists $u \in W$ such that wRu and $f(u) = u'$. *(backward condition)*

We define the *image* of f as the set $\text{Im } f := f[\mathfrak{F}] = \{f(w) : w \in \mathfrak{F}\}$.

Definition 2.2. Let $\mathfrak{F} = (W, R)$ be a Kripke frame. Let $X \subseteq \mathcal{X}$ and let $V : X \rightarrow 2^W$ be a *valuation* on \mathfrak{F} . Given $w \in W$, we define $\mathfrak{F}, V, w \models \varphi$ by induction on $\varphi \in \mathcal{L}_X$:

- $\mathfrak{F}, V, w \models x$ if $w \in V(x)$,
- $\mathfrak{F}, V, w \models \neg\varphi$ if $\mathfrak{F}, V, w \not\models \varphi$,
- $\mathfrak{F}, V, w \models \varphi \wedge \psi$ if $\mathfrak{F}, V, w \models \varphi$ and $\mathfrak{F}, V, w \models \psi$,
- $\mathfrak{F}, V, w \models \Box\varphi$ if $\mathfrak{F}, V, u \models \varphi$ for all $u \in W$ such that wRu .

We write $\mathfrak{F} \models \varphi$ in case $\mathfrak{F}, V, w \models \varphi$ for all valuations V on \mathfrak{F} and all $w \in W$. If $\Gamma \subseteq \mathcal{L}_X$, we write $\mathfrak{F} \models \Gamma$ in case $\mathfrak{F} \models \varphi$ for all $\varphi \in \Gamma$.

For the rest of the paper, we fix a normal modal logic \mathbf{L} .

Proposition 2.3 ([BdRV01, Th. 3.14]).

1. If $\mathfrak{F} \models \mathbf{L}$ and \mathfrak{F}' is a generated subframe of \mathfrak{F} , then $\mathfrak{F}' \models \mathbf{L}$.
2. If f is a surjective bounded morphism from \mathfrak{F} to \mathfrak{F}' and $\mathfrak{F} \models \mathbf{L}$, then $\mathfrak{F}' \models \mathbf{L}$.

Definition 2.4. Let \mathbf{L} be a normal modal logic, and X be a finite set of variables. A set $\Gamma \subseteq \mathcal{L}_X$ of formulas is said to be *consistent* if there are no $\varphi_1, \dots, \varphi_n \in \Gamma$ such that $\vdash_{\mathbf{L}} \neg \bigwedge_{i=1}^n \varphi_i$. We say that Γ is *maximal consistent* if Γ is consistent, and there is no consistent set $\Delta \subseteq \mathcal{L}_X$ such that $\Gamma \subset \Delta$. The *canonical frame* of \mathbf{L} over X is then the frame $\mathfrak{F}_X = (W_X, R_X)$ where:

- W_X is the set of all maximal consistent sets of \mathbf{L} over X , and
- $\Gamma R_X \Delta$ iff for all $\Box\varphi \in \Gamma$, we have $\varphi \in \Delta$.

²Traditionally, parametric unification requires that $\sigma(p) = p$ instead of $\sigma(p) \equiv_{\mathbf{L}} p$. This condition, however, is too strictly syntactic to admit a dual counterpart, as in Proposition 2.9. This is why we favor our version, which presents the advantage of being invariant under equivalence of substitutions. As we will see in Remark 2.20, this does not alter unification in \mathbf{L} in any significant way.

We also introduce the *canonical valuation* V_X defined by $V_X(x) := \{\Gamma \in W_X : x \in \Gamma\}$. We denote by $\langle \varphi \rangle_X := \{\Gamma \in W_X : \varphi \in \Gamma\}$ the *extension* of φ in \mathfrak{F}_X . We then denote by $\langle \varphi \rangle_X^\infty := \bigcap_{n \in \mathbb{N}} \langle \Box^n \varphi \rangle_X$ the *tight extension* of φ in \mathfrak{F}_X .

The following results are known as the *Lindenbaum Lemma* and *Truth Lemma*, respectively. In the sequel, we will use them without justification.

Proposition 2.5 ([BdRV01, Lemma 4.17]). *If $\Gamma \subseteq \mathcal{L}_X$ is consistent, then there exists $\Delta \in W_X$ such that $\Gamma \subseteq \Delta$.*

Proposition 2.6 ([BdRV01, Lemma 4.21]). *For all $w \in W_X$ and $\varphi \in \mathcal{L}_X$, we have $\varphi \in w$ iff $\mathfrak{F}_X, V_X, w \models \varphi$.*

The logic \mathbf{L} is said to be *locally tabular* if for every $X \subseteq \mathcal{X}$ finite, there are only finitely many formulas in \mathcal{L}_X up to $\equiv_{\mathbf{L}}$. In particular, this implies that \mathfrak{F}_X is finite, and that $\mathfrak{F}_X \models \mathbf{L}$, as a consequence of the Truth Lemma and [BdRV01, Lemma 3.27]. In what follows, we assume that \mathbf{L} is locally tabular.

As we have seen, the canonical frame offers a bridge between logics and frames, between syntax and semantics. Duality theory appears to further strengthen this link, by establishing a correspondence between substitutions and bounded morphisms. Indeed, if $\sigma : \mathcal{L}_X \rightarrow \mathcal{L}_Y$ is a substitution, the *dual* of σ is the bounded morphism $\sigma^* : \mathfrak{F}_Y \rightarrow \mathfrak{F}_X$ defined by $\sigma^*(\Gamma) := \sigma^{-1}[\Gamma]$ for all $\Gamma \in W_Y$. This correspondence is one-to-one modulo $\equiv_{\mathbf{L}}$, which means that (1) for every bounded morphism $f : \mathfrak{F}_Y \rightarrow \mathfrak{F}_X$ there exists a substitution $\sigma : \mathcal{L}_X \rightarrow \mathcal{L}_Y$ such that $\sigma^* = f$, and (2) if $\sigma^* = \tau^*$ then $\sigma \equiv_{\mathbf{L}} \tau$ (see [BG22]).³ It can then be specialized to a correspondence between parametric substitutions and *parametric bounded morphisms*.

Definition 2.7. Let \mathfrak{G} be a generated subframe of \mathfrak{F}_Y and \mathfrak{F} be a generated subframe of \mathfrak{F}_X . A bounded morphism $f : \mathfrak{G} \rightarrow \mathfrak{F}$ is said to be *parametric* if for all $p \in X \cap \mathcal{P}$ we have $p \in Y$ and $f^{-1}[\langle p \rangle_X \cap \mathfrak{F}] = \langle p \rangle_Y \cap \mathfrak{G}$.

Remark 2.8. Alternatively, f is parametric if and only if for all $p \in X \cap \mathcal{P}$ we have $p \in Y$, and then for all $w \in \mathfrak{G}$ we have $w \in \langle p \rangle_Y \iff f(w) \in \langle p \rangle_X$. The point is that parametric bounded morphisms ‘preserve’ parameters in a similar way as parametric substitutions do.

Proposition 2.9. *Let $\sigma : \mathcal{L}_X \rightarrow \mathcal{L}_Y$ be a substitution. Then σ is a parametric substitution if and only if σ^* is a parametric bounded morphism.*

Proof. Suppose that σ is parametric, and let $p \in X \cap \mathcal{P}$. By assumption we have $p \in Y$. Then $\Gamma \in \sigma^{*-1}[\langle p \rangle_X]$ iff $\sigma^*(\Gamma) \in \langle p \rangle_X$, iff $\sigma(p) \in \Gamma$, iff $p \in \Gamma$ by assumption, iff $\Gamma \in \langle p \rangle_Y$. Conversely, suppose that σ is not parametric. Then there exists $p \in X \cap \mathcal{P}$ such that $p \notin Y$, or else $\sigma(p) \not\equiv_{\mathbf{L}} p$. If $p \notin Y$ we are done, so assume that $\sigma(p) \not\equiv_{\mathbf{L}} p$. Then $\{\neg(\sigma(p) \leftrightarrow p)\}$ is consistent, and so there exists $\Gamma \in W_Y$ such that either $\sigma(p) \wedge \neg p \in \Gamma$ or $\neg\sigma(p) \wedge p \in \Gamma$. In the former case, we have $\Gamma \in \sigma^{*-1}[\langle p \rangle_X]$ and $\Gamma \notin \langle p \rangle_Y$, and in the latter case we have $\Gamma \notin \sigma^{*-1}[\langle p \rangle_X]$ and $\Gamma \in \langle p \rangle_Y$. In both cases, this proves that $\sigma^{*-1}[\langle p \rangle_X] \neq \langle p \rangle_Y$. \square

³In the general case, one also has to endow \mathfrak{F}_X and \mathfrak{F}_Y with a collection of admissible sets, and to restrict the definition of a bounded morphism accordingly. However, when \mathfrak{F}_X and \mathfrak{F}_Y are finite (as this is the case here), this condition collapses and can simply be forgotten. For details we refer to [BG22].

2.3 Euclidean frames

The logic **K5** is known to be the logic of *Euclidean frames*. A Kripke frame $\mathfrak{F} = (W, R)$ is said to be *Euclidean* if wRu and wRv implies uRv , for all $w, u, v \in W$ [CZ97, Cor. 3.37]. It is then widely known that \mathfrak{F} is Euclidean if and only if $\mathfrak{F} \models \mathbf{K5}$. From now on, **we assume that \mathbf{L} is an extension of $\mathbf{K5}$** . In this case, \mathbf{L} is locally tabular [NT85, Cor. 5], and it follows that \mathfrak{F}_X is a finite Euclidean frame, for all $X \subseteq \mathcal{X}$ finite. In this section, we provide a comprehensive description of Euclidean frames and bounded morphisms between them, as these will be of central interest in the sequel. The following property is known under the name of *2-transitivity*.

Proposition 2.10. *Let $\mathfrak{F} = (W, R)$ be an Euclidean frame. If $wR^{\leq 3}u$ then $wR^{\leq 2}u$.*

Proof. Suppose that $wR^{\leq 3}u$ and *not* $wR^{\leq 2}u$. Then there exist $v, v' \in W$ such that $wRvRv'Ru$. We use the fact that \mathfrak{F} is Euclidean. First, from wRv we obtain vRv . From vRv' and vRv we then obtain $v'Rv$. Finally, from $v'Rv$ and $v'Ru$ we obtain vRu . Thus $wRvRu$, a contradiction. This proves the claim. \square

Now, let $\mathfrak{F} = (W, R)$ be an Euclidean frame. A *pre-cluster* in \mathfrak{F} is a non-empty set $c \subseteq W$ such that for all $w, u \in c$ we have wRu . A pre-cluster c is called a *cluster* if there exists no pre-cluster c' in \mathfrak{F} such that $c \subset c'$; or alternatively, if whenever $w \in c$ and wRu we have $u \in c$. A pre-cluster which is not a cluster is called a *proper* pre-cluster. If W is a cluster, the frame \mathfrak{F} itself will be called a cluster, and typically named \mathfrak{C} .

A point w in \mathfrak{F} is *reflexive* if wRw , *initial* if there is no u such that uRw , and *isolated* if there is no u such that wRu . Note that since \mathfrak{F} is Euclidean, every isolated point is also initial. Indeed, by contraposition, if w is non-isolated then we have uRw for some u , and it follows that wRw . We write $w \triangleright c$ whenever w is initial, c is a pre-cluster and $c = \{u \in W \mid wRu\}$. In this case we call w a *predecessor* of c . We call w a *quasi-predecessor* of c , and write $w \trianglerighteq c$, if either w is a predecessor of c , or c is a cluster and $w \in c$. A pre-cluster with no predecessor is said to be *degenerated*. If C is a cluster we then write $\hat{C} := C \cup \{w \in W \mid \exists u \in C, wRu\}$.

We denote by \bar{R} the smallest equivalence relation on W containing R . The frame \mathfrak{F} is said to be *primitive* if for all $w, u \in W$ we have $w\bar{R}u$. Primitive frames present a fairly simple structure.

Proposition 2.11 ([NT85, Lemma 1]). *A non-empty Euclidean frame $\mathfrak{F} = (W, R)$ is primitive if and only if one of the following holds:*

1. *we have $W = \{w\}$ with w isolated,*
2. *there exists a cluster C in \mathfrak{F} such that $W = \hat{C}$.*

Proof. The direction from right to left is clear. Conversely, assume that \mathfrak{F} is primitive. Since \mathfrak{F} is non-empty, there exists $w \in W$. First suppose that w is isolated. For all $u \in W$, we have $w\bar{R}u$, but then the only possibility is $w = u$. Thus \mathfrak{F} falls in case 1. Otherwise, there exists $u \in W$ such that wRu . Then we also have uRu . Let $C := \{t \in W : wR^2t\}$, which is non-empty by assumption. First, let $t, t' \in C$. Then $wRvRt$ and $wRv'Rt'$ for some $v, v' \in W$. Since \mathfrak{F} is Euclidean, we obtain vRv' , and then $v'Rt$, and finally tRt' and $t'Rt$. This proves that C is a pre-cluster. By Proposition 2.10, it is then clear that C is a cluster. Now let $v \in W$. We have $w\bar{R}v$, but by Proposition 2.10 and the Euclidean character of R , we see that this amounts to vRt for some $t \in C$. Therefore $W = \hat{C}$, and we are done. \square

Primitive frames that fall in case 1 will be called *trivial*. If \mathfrak{F} falls in case 2, then it is clear that the cluster C is the unique cluster in \mathfrak{F} . If \mathfrak{C} is the subframe of \mathfrak{F} generated by C , then we write $\mathfrak{F} = \hat{\mathfrak{C}}$ and we call \mathfrak{F} a *pinned cluster*. Thus, Proposition 2.11 can also be understood as the property that every Euclidean frame is a disjoint union of trivial frames and pinned clusters. Examples of primitive frames are depicted in Figure 1.




Figure 1: Some primitive Euclidean frames. From left to right: a trivial frame; a degenerated two-element cluster; a one-element cluster with two predecessors; a non-degenerated two-element cluster, containing both a degenerated and a non-degenerated one-element pre-cluster; a degenerated cluster containing two overlapping two-element pre-clusters with one predecessor each.

Proposition 2.12. *Let $\mathfrak{F} = (W, R)$ and $\mathfrak{F}' = (W', R')$ be two Euclidean frames and $f : W \rightarrow W'$ be a map. Then f is a bounded morphism from \mathfrak{F} to \mathfrak{F}' if and only if the following four conditions are satisfied.*

1. *If w is isolated in \mathfrak{F} then $f(w)$ is isolated in \mathfrak{F}' .*
2. *If C is a cluster in \mathfrak{F} then $f[C]$ is a cluster in \mathfrak{F}' .*
3. *Suppose that c is a pre-cluster in \mathfrak{F} . If $w \triangleright c$, then $f(w) \triangleright f[c]$.*
4. *Suppose that c is a pre-cluster in \mathfrak{F} , and that $f[c]$ is a proper pre-cluster in \mathfrak{F}' . If $w \triangleright c$, then $f(w) \triangleright f[c]$.*

We now prove a number of important results related to the canonical frame of \mathbf{L} . Lemma 2.13 and Lemma 2.14 state that primitive subframes can be ‘copied’ from one canonical frame to another, while forcing some variables to be true or false on the copy. Lemma 2.15 determines the number of predecessors of a pre-cluster. First, we set some notations. Let $Y, Z \subseteq \mathcal{X}$ be finite, and $P \subseteq Y \cap Z$. Given $w \in \mathfrak{F}_Y$ and $u \in \mathfrak{F}_Z$, we write $w \sim_P u$ if we have $w \in \langle p \rangle_Y \iff u \in \langle p \rangle_Z$ for all $p \in P$.

Lemma 2.13. *Let $Y, Z \subseteq \mathcal{X}$ be finite, and $P \subseteq Y \cap Z$. Let $\{w\}$ be a trivial subframe of \mathfrak{F}_Z . Then there exists a trivial subframe $\{w'\}$ of \mathfrak{F}_Y such that $w \sim_P w'$.*

Proof. Let us write $\rho := \left(\bigwedge_{p \in P \text{ and } w \in \langle p \rangle_Z} p \right) \wedge \left(\bigwedge_{p \in P \text{ and } w \notin \langle p \rangle_Z} \neg p \right)$. By construction we have $w \in \langle \square \perp \wedge \rho \rangle_Z$, and so $\{\square \perp \wedge \rho\}$ is consistent. Hence, there exists $w' \in \mathfrak{F}_Y$ such that $w' \in \langle \square \perp \wedge \rho \rangle_Y$, and $\{w'\}$ is then the desired frame. \square

Lemma 2.14. *Let $Y, Z \subseteq \mathcal{X}$ be finite, and $P \subseteq Y \cap Z$. Let \mathfrak{C} be a cluster in \mathfrak{F}_Z , and \sim be an equivalence relation on \mathfrak{C} such that $w \sim u$ implies $w \sim_P u$ for all $w, u \in \mathfrak{C}$. Assume that $|\mathfrak{C}/\sim| \leq 2^{|Y \setminus P|}$. Then there exists a cluster \mathfrak{C}' in \mathfrak{F}_Y and a surjective map $f : \mathfrak{C} \rightarrow \mathfrak{C}'$ such that for all $w, u \in \mathfrak{C}$:*

1. *$w \sim u$ if and only if $f(w) = f(u)$,*
2. *$w \sim_P f(w)$.*

Lemma 2.15. *Let $X \subseteq \mathcal{X}$ be finite and \mathfrak{C} be a cluster in \mathfrak{F}_X .*

1. *If \mathfrak{C} is non-degenerated then \mathfrak{C} has $2^{|X|} - |\mathfrak{C}|$ predecessors.*
2. *Every non-degenerated proper pre-cluster of \mathfrak{C} has $2^{|X|}$ predecessors.*

More precisely, for all $c \subset \mathfrak{C}$ and $Y \subseteq X$, there is exactly one quasi-predecessor w of c such that $Y = \{x \in X : w \in \langle x \rangle_X\}$.

2.4 The landscape of extensions of **K5**

Investigating unification in *all* Euclidean logics is an ambitious undertaking, but luckily we do not start from scratch. In their 1985 paper [NT85], Nagel and Thomason studied these logics in depth, and provided a complete characterization of the extensions of **K5**. In particular, they show that every one of them is finitely axiomatizable, and that their axiomatization can be put in some ‘normal form’. We briefly present their results.

Let $m \geq 1$. Given $n \geq 0$ we define the frame $\mathfrak{F}_{n,m} = (W, R)$ with $W = \{w\} \cup A \cup B$ and $R = (\{w\} \times A) \cup (A \cup B)^2$, where A and B are disjoint and $|A| = m$ and $|B| = n$. We define the frame $\mathfrak{F}_{-1,m} = (W, R)$ with $|W| = m$ and $R = W^2$. We also introduce the set

$$\mathbf{N}_2 := \{(n, m) : n \geq -1 \text{ and } m \geq 1\}.$$

Pairs in \mathbf{N}_2 are compared coordinate-wise, that is, $(n, m) \leq (n', m')$ stands for $n \leq m$ and $n' \leq m'$. Nagel and Thomason then introduced, for all $(n, m) \in \mathbf{N}_2$, an axiom $\mathbf{Ax}_{n,m}$ expressing the fact that the frame $\mathfrak{F}_{n,m}$ is a ‘forbidden pattern’. To be precise, we have for all $(n', m') \in \mathbf{N}_2$,

$$\mathfrak{F}_{n',m'} \models \mathbf{Ax}_{n,m} \iff (n, m) \not\leq (n', m')$$

and so in particular $\mathfrak{F}_{n,m} \not\models \mathbf{Ax}_{n,m}$. Then every extension \mathbf{L} of **K5** admits a finite axiomatization of the form

$$\mathbf{L} = \mathbf{K5} + \{\mathbf{Ax}_{n,m} : (n, m) \in \Delta\} + \mathbf{D}^? \quad (1)$$

where $\mathbf{D}^? \in \{\top, \diamond\top\}$ and $\Delta \subseteq \mathbf{N}_2$ is finite. Basically, the axiom $\mathbf{D}^?$ is either vacuous (in case $\mathbf{D}^? = \top$), or else enforces the absence of any isolated point in the canonical frame. In the sequel, we assume that \mathbf{L} is presented as in 1, and consider Δ and $\mathbf{D}^?$ to be fixed from now on. Let us write $\uparrow\Delta := \{(n', m') \in \mathbf{N}_2 : \exists (n, m) \in \Delta, (n, m) \leq (n', m')\}$. In Proposition 2.16, we prove that $\uparrow\Delta$ characterizes the degenerated pre-clusters of the canonical frame. We give in Table 2 a presentation of standard Euclidean logics in terms of Δ of $\mathbf{D}^?$. We also provide a schematic depiction of $\uparrow\Delta$ in Figure 2.

L	Δ	$\mathbf{D}^?$
K5	\emptyset	\top
KD5	\emptyset	$\diamond\top$
K45	$\{(1, 1)\}$	\top
Kt5	$\{(1, 0)\}$	\top
S5	$\{(1, 0)\}$	$\diamond\top$

Table 2: Standard characterizations of some logics

Proposition 2.16. *Let $X \subseteq \mathcal{X}$ be finite.*

1. *Given $m' \geq 1$, there is a cluster in \mathfrak{F}_X of size m' if and only if $2^{|X|} \geq m'$ and $(-1, m') \notin \uparrow\Delta$.*
2. *Let \mathcal{C} be a cluster in \mathfrak{F}_X , and $c \subseteq \mathcal{C}$. Let $n' := |\mathcal{C} \setminus c|$ and $m' := |c|$. Then c is degenerated if and only if $(n', m') \in \uparrow\Delta$.*




Figure 2: The space \mathbf{N}_2 and the set Δ , whose elements are identified by circles. The colored area corresponds to $\uparrow\Delta$.

2.5 Unification and admissibility

If $\sigma : \mathcal{L}_X \rightarrow \mathcal{L}_Y$ and $\tau : \mathcal{L}_X \rightarrow \mathcal{L}_Z$ are two substitutions, we write $\sigma \simeq \tau$ in case $\vdash_{\mathbf{L}} \sigma(x) \leftrightarrow \tau(x)$ for all $x \in X$. We say that σ is *at least as general* as τ , and write $\sigma \preceq \tau$, if there exists a substitution $\mu : \mathcal{L}_Y \rightarrow \mathcal{L}_{Z'}$ such that $\tau \simeq \mu\sigma$. The following lemma states that we can in fact always take $Z' = Z$ in the previous definition.

Lemma 2.17. *Let $\sigma : \mathcal{L}_X \rightarrow \mathcal{L}_Y$ and $\tau : \mathcal{L}_X \rightarrow \mathcal{L}_Z$ be two substitutions. Then $\sigma \preceq \tau$ if and only if there exists a substitution $\mu : \mathcal{L}_Y \rightarrow \mathcal{L}_Z$ such that $\tau \simeq \mu\sigma$.*

Proof. From right to left, this is clear. Conversely, suppose that $\sigma \preceq \tau$. Then there exists a substitution $\mu' : \mathcal{L}_Y \rightarrow \mathcal{L}_{Z'}$ such that $\tau \simeq \mu'\sigma$. Let $\lambda : \mathcal{L}_{Z'} \rightarrow \mathcal{L}_Z$ be the substitution defined by $\lambda(z) := z$ for all $z \in Z \cap Z'$, and $\lambda(z) := \perp$ for all $z \in Z' \setminus Z$. Then from $\tau \simeq \mu'\sigma$ we obtain $\lambda\tau \simeq \lambda\mu'\sigma$, and since τ has codomain \mathcal{L}_Z , we have $\lambda\tau = \tau$. Setting $\mu := \lambda\mu' : \mathcal{L}_Y \rightarrow \mathcal{L}_Z$, we finally obtain $\tau \simeq \mu\sigma$. \square

When a substitution $\sigma : \mathcal{L}_X \rightarrow \mathcal{L}_Y$ is parametric, we have in particular $X \cap \mathcal{P} \subseteq Y \cap \mathcal{P}$. The following lemma allows to strengthen this inclusion into an equality, which in some cases is more convenient.

Lemma 2.18. *Let $\sigma : \mathcal{L}_X \rightarrow \mathcal{L}_Y$ be a unifier of φ . Then there exists a unifier $\tau : \mathcal{L}_X \rightarrow \mathcal{L}_Z$ of φ such that $\tau \preceq \sigma$ and $X \cap \mathcal{P} = Z \cap \mathcal{P}$.*

Proof. Let $P := Y \cap \mathcal{P} \setminus X$ be the set of ‘excessive’ parameters in Y . For every $p \in P$, we introduce a fresh variable z_p (not a parameter). Let $Z := (Y \setminus P) \cup \{z_p : p \in P\}$, which satisfies $X \cap \mathcal{P} = Z \cap \mathcal{P}$

by construction. Let $\lambda : Y \rightarrow Z$ be the substitution defined by $\lambda(p) := z_p$ for all $p \in P$, and $\lambda(y) := y$ for all $y \in Y \setminus P$. Likewise, let $\mu : Z \rightarrow Y$ be defined by $\mu(z_p) := p$ for all $p \in P$, and $\mu(z) := z$ for all other $z \in Z$. Writing $\tau := \lambda\sigma$, it is then clear that τ is a unifier of φ and that $\mu\tau = \sigma$, whence $\tau \preceq \sigma$. \square

Lemma 2.19. *The relation \preceq is transitive.*

Proof. Consider three substitutions $\sigma : \mathcal{L}_X \rightarrow \mathcal{L}_Y$, $\sigma' : \mathcal{L}_X \rightarrow \mathcal{L}_{Y'}$, $\sigma'' : \mathcal{L}_X \rightarrow \mathcal{L}_{Y''}$ such that $\sigma \preceq \sigma'$ and $\sigma' \preceq \sigma''$. Then there exist two substitutions μ and μ' such that $\sigma' \simeq \mu\sigma$ and $\sigma'' \simeq \mu'\sigma'$. Let $x \in X$. We have $\vdash_{\mathbf{L}} \sigma'(x) \leftrightarrow \mu\sigma(x)$, and by uniform substitution it follows that $\vdash_{\mathbf{L}} \mu'\sigma'(x) \leftrightarrow \mu'\mu\sigma(x)$. Since $\vdash_{\mathbf{L}} \sigma''(x) \leftrightarrow \mu'\sigma'(x)$, it follows that $\vdash_{\mathbf{L}} \sigma''(x) \leftrightarrow \mu'\mu\sigma(x)$. If $\mu'' := \mu'\mu$, we then obtain $\sigma'' \simeq \mu''\sigma$, and thus $\sigma \preceq \sigma''$. \square

Now, let $\varphi \in \mathcal{L}$ and $X := \text{Var}(\varphi)$. A *unifier* of φ is a parametric substitution $\sigma : \mathcal{L}_X \rightarrow \mathcal{L}_Y$ such that $\vdash_{\mathbf{L}} \sigma(\varphi)$. If φ admits a unifier, then φ is said to be *unifiable*. A set Σ of unifiers of φ is said to be *complete* if for all unifiers τ of φ , there exists $\sigma \in \Sigma$ such that $\sigma \preceq \tau$. In this case, when Σ is a singleton $\{\sigma\}$, we call σ a *most general unifier* of φ (or *MGU* for short). If every unifiable formula admits a MGU, then \mathbf{L} is said to be *unary*.

Remark 2.20. Due to technical motivations, our definition of unifiers does not exactly agree with the state of the art. As explained in Footnote 2, the preservation of parameters is generally taken to be purely syntactic. So let us call a *standard unifier* of φ any substitution $\tau : \mathcal{L}_X \rightarrow \mathcal{L}_Y$ such that $\vdash_{\mathbf{L}} \tau(\varphi)$, and such that for all $p \in X \cap \mathcal{P}$ we have $p \in Y$ and $\sigma(p) = p$. Though the two presentations differ, we claim that they are equivalent as far as the structural properties of the space of unifiers are concerned. Indeed, given a unifier $\sigma : \mathcal{L}_X \rightarrow \mathcal{L}_Y$ of φ , we can define $\tau : \mathcal{L}_X \rightarrow \mathcal{L}_Y$ by $\tau(p) := p$ for all $p \in X \cap \mathcal{P}$ and $\tau(x) := \sigma(x)$ for all $x \in X \setminus \mathcal{P}$. Since σ is parametric, it is clear that $\sigma \simeq \tau$. Since σ is a unifier of φ , it follows that τ is a standard unifier of φ . Hence, every unifier of φ is equivalent to some standard unifier of φ (and the converse is obviously true). In particular, this implies that φ admits a MGU iff φ admits a standard MGU.

Unification is closely related to the problem of *admissibility*, which we introduce in its parametric variant.

Definition 2.21. Let $\varphi_1, \dots, \varphi_n, \psi \in \mathcal{L}$. The rule $\frac{\varphi_1 \dots \varphi_n}{\psi}$ is said to be *admissible* if every unifier of $\varphi_1 \wedge \dots \wedge \varphi_n$ is also a unifier of ψ . The decision problem of *admissibility* has inputs of the form $(\varphi_1, \dots, \varphi_n, \psi) \in \mathcal{L}^{n+1}$ for some $n \in \mathbb{N}$, and an input $(\varphi_1, \dots, \varphi_n, \psi)$ is positive if and only if the rule $\frac{\varphi_1 \dots \varphi_n}{\psi}$ is admissible.

The traditional account of admissibility corresponds to the case where $\varphi_1, \dots, \varphi_n, \psi$ are parameter-free. Proposition 2.22 below was first pointed out in [Ghi99].

Proposition 2.22. *Let $\varphi_1, \dots, \varphi_n, \psi \in \mathcal{L}$, and suppose that $\varphi_1 \wedge \dots \wedge \varphi_n$ admits a most general unifier σ . Then $\frac{\varphi_1 \dots \varphi_n}{\psi}$ is admissible if and only if σ is a unifier of ψ .*

In [BG22], it was proved that unifiers correspond to a certain type of bounded morphisms.

Definition 2.23. A *dual unifier* of φ is a parametric bounded morphism $f : \mathfrak{F}_Y \rightarrow \mathfrak{F}_X$ such that $\text{Im } f \subseteq \langle \varphi \rangle_X^\infty$.

Proposition 2.24 ([BG22, Th. 4.5]). *A parametric substitution $\sigma : \mathcal{L}_X \rightarrow \mathcal{L}_Y$ is a unifier of φ if and only if σ^* is a dual unifier of φ .*

We also proceed to dualize the relation \preceq , namely, if σ and τ are two substitutions, we introduce conditions on σ^* and τ^* that are necessary and sufficient for $\sigma \preceq \tau$. So consider two bounded morphisms $g : \mathfrak{F}_Y \rightarrow \mathfrak{F}_X$ and $h : \mathfrak{F}_Z \rightarrow \mathfrak{F}_X$. We write $g \preceq h$ in case there exists a bounded morphism $f : \mathfrak{F}_Z \rightarrow \mathfrak{F}_Y$ such that $gf = h$.

Proposition 2.25. *Let $\sigma : \mathcal{L}_X \rightarrow \mathcal{L}_Y$, $\tau : \mathcal{L}_X \rightarrow \mathcal{L}_Z$ and $\mu : \mathcal{L}_Y \rightarrow \mathcal{L}_Z$ be three substitutions. Then $\tau \simeq \mu\sigma$ iff $\sigma^*\mu^* = \tau^*$.*

Proof. Suppose that $\tau \simeq \mu\sigma$. Let $\Gamma \in \mathfrak{F}_Z$. Given $\varphi \in \mathcal{L}_Z$, the assumption $\tau \simeq \mu\sigma$ yields $\vdash_{\mathbf{L}} \mu\sigma(\varphi) \leftrightarrow \tau(\varphi)$, and it follows that $\varphi \in \tau^*(\Gamma)$ iff $\tau(\varphi) \in \Gamma$, iff $\mu\sigma(\varphi) \in \Gamma$, iff $\varphi \in \sigma^*\mu^*(\Gamma)$. Hence $\tau^*(\Gamma) = \sigma^*\mu^*(\Gamma)$, and this proves that $\sigma^*\mu^* = \tau^*$.

Conversely, suppose that $\sigma^*\mu^* = \tau^*$. We show that $\tau \simeq \mu\sigma$. For suppose not. Then there exists $x \in X$ such that $\not\vdash_{\mathbf{L}} \tau(x) \leftrightarrow \mu\sigma(x)$. Hence, there exists $\Gamma \in \mathfrak{F}_Z$ such that $\Gamma \notin \langle \mu\sigma(x) \leftrightarrow \tau(x) \rangle_Z$. Then either $\mu\sigma(x) \rightarrow \tau(x) \notin \Gamma$ or $\tau(x) \rightarrow \mu\sigma(x) \notin \Gamma$. Assume that $\mu\sigma(x) \rightarrow \tau(x) \notin \Gamma$. Then $\mu\sigma(x) \in \Gamma$ and $\tau(x) \notin \Gamma$, whence $x \in (\sigma^*\mu^*)(\Gamma)$ and $x \notin \tau^*(\Gamma)$, and thus $(\sigma^*\mu^*)(\Gamma) \neq \tau^*(\Gamma)$, contradicting $\sigma^*\mu^* = \tau^*$. In case $\tau(x) \rightarrow \mu\sigma(x) \notin \Gamma$, the reasoning is similar. \square

As a consequence of Lemma 2.17 and Proposition 2.25, we get the following result.

Corollary 2.26. *Let $\sigma : \mathcal{L}_X \rightarrow \mathcal{L}_Y$ and $\tau : \mathcal{L}_X \rightarrow \mathcal{L}_Z$ be two substitutions. Then $\sigma \preceq \tau$ if and only if $\sigma^* \preceq \tau^*$.*

We also obtain the following from Lemma 2.18.

Corollary 2.27. *Let $f : \mathfrak{F}_Y \rightarrow \mathfrak{F}_X$ be a dual unifier of φ . Then there exists a dual unifier $g : \mathfrak{F}_Z \rightarrow \mathfrak{F}_X$ of φ such that $g \preceq f$ and $X \cap \mathcal{P} = Z \cap \mathcal{P}$.*

3 Reflections

Following the preliminary work laid down in Section 2.5, we now aim to study the structure of dual unifiers. However, those are complicated objects: the domain of a dual unifier is a full canonical frame, which consists of a large union of primitive frames. To bypass this difficulty, it will be convenient to reason on a more granular scale, that is, to work with *partial* dual unifiers (in the sense of partial functions). Such objects will be called *reflections*. In this section, we fix $X \subseteq \mathcal{X}$ finite.

Definition 3.1. Let \mathfrak{F} be a generated subframe of \mathfrak{F}_X . Given $Y \subseteq \mathcal{X}$ finite, a *Y-reflection* of \mathfrak{F} is a pair (\mathfrak{G}, g) such that:

- \mathfrak{G} is a closed subframe of \mathfrak{F}_Y ,
- g is a parametric bounded morphism from \mathfrak{G} to \mathfrak{F}_X ,
- $g[\mathfrak{G}] \subseteq \mathfrak{F}$.

We call (\mathfrak{G}, g) *primitive* if \mathfrak{G} is primitive. We call (\mathfrak{G}, g) *total* if $\mathfrak{G} = \mathfrak{F}_Y$.

For convenience, any Y -reflection of \mathfrak{F} , where Y is arbitrary, will be called a *reflection* of \mathfrak{F} . Given a Y -reflection (\mathfrak{G}, g) of \mathfrak{F} , and a Z -reflection (\mathfrak{H}, h) of \mathfrak{F} , we write $(\mathfrak{G}, g) \preceq (\mathfrak{H}, h)$ if there exists a parametric bounded morphism $f : \mathfrak{H} \rightarrow \mathfrak{G}$ such that $h = gf$. We say that \mathfrak{F} has *Y-bounded* (total, primitive) reflection if for every (total, primitive) reflection (\mathfrak{H}, h) of \mathfrak{F} , there exists a (total, primitive) Y -reflection (\mathfrak{G}, g) of \mathfrak{F} such that $(\mathfrak{G}, g) \preceq (\mathfrak{H}, h)$.

If $\mathfrak{F} = \langle \varphi \rangle_X^\infty$ and $X = \text{Var}(\varphi)$ for some $\varphi \in \mathcal{L}$, it is easy to see that the total reflections of \mathfrak{F} are exactly the dual unifiers of φ . In addition, the relation \preceq , when restricted to dual unifiers, coincides with the relation \preceq defined in Section 2.5. Hence, reflections are a proper generalization of dual unifiers.

Lemma 3.2. *The relation \preceq on reflections of \mathfrak{F} is transitive.*

Proof. Suppose that $(\mathfrak{G}, g) \preceq (\mathfrak{G}', g')$ and $(\mathfrak{G}', g') \preceq (\mathfrak{G}'', g'')$. Then there exist two parametric bounded morphisms f and f' such that $g' = gf$ and $g'' = g'f'$. It follows that $g'' = g(ff')$, whence $(\mathfrak{G}, g) \preceq (\mathfrak{G}'', g'')$. \square

We first set a useful lemma, by showing that reflections can be ‘inflated’ – a process dual to how a substitution $\sigma : \mathcal{L}_X \rightarrow \mathcal{L}_Y$ can be extended to a substitution $\sigma' : \mathcal{L}_X \rightarrow \mathcal{L}_Z$ with $Y \subseteq Z$.

Lemma 3.3. *Let \mathfrak{F} be a generated subframe of \mathfrak{F}_X . Let $Y, Z \subseteq \mathcal{X}$ be finite with $X \cap \mathcal{P} = Y \cap \mathcal{P} \subseteq Z \cap \mathcal{P}$. Let $Y_0, Y_1 \subseteq Y \setminus X$ and assume that $Y_0 \cap Y_1 = \emptyset$ and $|Z| = |Y \setminus (Y_0 \cup Y_1)|$. Let (\mathfrak{H}, h) be a primitive Z -reflection of \mathfrak{F} . Then there exists a primitive Y -reflection (\mathfrak{G}, g) of \mathfrak{F} such that:*

- $(\mathfrak{G}, g) \preceq (\mathfrak{H}, h)$,
- $\mathfrak{G} \subseteq \langle y \rangle_Y$ for all $y \in Y_1$,
- $\mathfrak{G} \subseteq \langle \neg y \rangle_Y$ for all $y \in Y_0$.

Corollary 3.4. *Let $Y, Z \subseteq \mathcal{X}$ be finite with $X \cap \mathcal{P} = Y \cap \mathcal{P} \subseteq Z \cap \mathcal{P}$. If \mathfrak{F} has primitive Z -bounded reflection and $|Z| \leq |Y|$, then \mathfrak{F} has primitive Y -bounded reflection.*

Proof. Since $|Z| \leq |Y|$, we can select arbitrary sets $Y_0, Y_1 \subseteq Y$ so that $Y_0 \cap Y_1 = \emptyset$ and $|Z| = |Y \setminus (Y_0 \cup Y_1)|$. Then, it suffices to apply Lemma 3.3. \square

Now, our goal is to show that every generated subframe \mathfrak{F} of \mathfrak{F}_X has bounded total Y -reflection for a large enough set Y of variables, whose size can be computed effectively. To this end, our strategy is to fragment \mathfrak{F}_X into primitive generated subframes, and to show that each of these subframes has primitive Y -bounded reflection. By ‘gluing’ all of these Y -reflections together, we will then obtain the desired total Y -reflection of \mathfrak{F} . So first, let us prove that primitive frames have bounded primitive reflection. To this end, we will distinguish the case of trivial frames from that of pinned clusters. Observe that, by Proposition 2.12, all primitive reflections of trivial frames are trivial frames themselves, and that all primitive reflections of pinned clusters are pinned clusters.

Lemma 3.5. *Let $Y \subseteq \mathcal{X}$ be finite and such that $X \cap \mathcal{P} \subseteq Y$. Then every trivial subframe of \mathfrak{F}_X has Y -bounded reflection.*

Proof. Let $\{w\}$ be a trivial subframe of \mathfrak{F}_X . Let $(\{v\}, h)$ be a Z -reflection of \mathfrak{F} , for some Z . Then we have $Y \cap \mathcal{P} \subseteq Z$ and thus $X \cap \mathcal{P} \subseteq Z$. By applying Lemma 2.13 to $P := X \cap \mathcal{P}$, we obtain a trivial subframe $\{u\}$ of \mathfrak{F}_Y with $v \sim_{X \cap \mathcal{P}} u$. As a result, u satisfies the same parameters in X as v (and thus also w). By setting $g(u) := w$, we obtain a Y -reflection $(\{u\}, g)$ of $\{w\}$, which obviously satisfies $(\{u\}, g) \preceq (\{v\}, h)$. \square

For the next lemma, we introduce the integers $N_{\mathbf{L}} := \max \{n : (n, m) \in \Delta\}$ and $M_{\mathbf{L}} := \max \{m : (n, m) \in \Delta\}$ (see Figure 3), with the convention that $\max \emptyset = -1$. We then set $K := \max \{2^{|X|}, N_{\mathbf{L}}, M_{\mathbf{L}}\}$.




Figure 3: The space \mathbf{N}_2 with some subsets of interest.

Lemma 3.6. *Let $\hat{\mathfrak{C}}$ be a pinned cluster in \mathfrak{F}_X , with X finite. Let Y be a finite set of variables such that $X \subseteq Y$ and $X \cap \mathcal{P} = Y \cap \mathcal{P}$ and $2^{|Y \setminus X|} \geq 2K|\mathfrak{C}| + 1$. Then $\hat{\mathfrak{C}}$ has Y -bounded primitive reflection.*

Proof. Let $(\hat{\mathfrak{C}}_1, g_1)$ be a primitive Z -reflection of $\hat{\mathfrak{C}}$, for some Z . If $|Z| < |Y|$, it suffices to apply Corollary 3.4, so we assume that $|Y| \leq |Z|$. For all $w \in \mathfrak{C}$, we introduce the ‘ w -cell’ $C_w^1 := g_1^{-1}(w)$ of \mathfrak{C}_1 , which will play a key role in the present proof. It is then clear that the C_w^1 ’s cover \mathfrak{C}_1 and are pairwise disjoint. The idea is that the cells of size $> 2K$ are somehow too big, and can be safely ‘contracted’ without affecting the existence of predecessors. Hence, we will be able to map \mathfrak{C}_1 onto a cluster \mathfrak{C}_0 whose cells are all of size $\leq 2K$. So, let $w \in \mathfrak{C}$. We introduce an equivalence relation \sim_w over C_w^1 as follows.

1. If $|C_w^1| \leq 2K$, then \sim_w is just the identity relation over C_w^1 .
2. Otherwise we have $|C_w^1| > 2K$. Let $\sim'_{X \cap \mathcal{P}}$ be the restriction of $\sim_{X \cap \mathcal{P}}$ to C_w^1 . Then it is clear that $\sim'_{X \cap \mathcal{P}}$ has at most $2^{|X \cap \mathcal{P}|}$ equivalence classes. Since $2^{|X \cap \mathcal{P}|} \leq 2^{|X|} \leq K \leq 2K < |C_w^1|$, we can construct an equivalence relation \sim_w finer than $\sim'_{X \cap \mathcal{P}}$, so that \sim_w has exactly $2K$ equivalence classes.

We then consider the equivalence relation $\sim := \bigcup_{w \in \mathfrak{C}} \sim_w$ on \mathfrak{C}_1 . By construction, \sim has at most $2K|\mathfrak{C}|$ equivalence classes, and satisfies $w \sim u \implies w \sim_{X \cap \mathcal{P}} u$. In addition, we have $2K|\mathfrak{C}| \leq 2^{|Y \setminus \mathcal{P}|} \leq 2^{|Y \setminus (X \cap \mathcal{P})|}$. Since $X \cap \mathcal{P} = Y \cap \mathcal{P}$ and $Y \cap \mathcal{P} \subseteq Z \cap \mathcal{P}$, we also have $X \cap \mathcal{P} \subseteq Z \cap \mathcal{P}$. We can thus apply Lemma 2.14 to $P := X \cap \mathcal{P}$, and obtain a cluster \mathfrak{C}_0 in \mathfrak{F}_Z and a surjective map $f : \mathfrak{C}_1 \rightarrow \mathfrak{C}_0$ such that for all $w, u \in \mathfrak{C}_1$, $w \sim u$ iff $f(w) = f(u)$, and $w \sim_{X \cap \mathcal{P}} f(w)$. Then, given $u_0 \in \mathfrak{C}_0$, there exists $u_1 \in \mathfrak{C}_1$ such that $f(u_1) = u_0$, and we set $g_0(u_0) := g_1(u_1)$. By construction of f , we see that $g_0 : \mathfrak{C}_0 \rightarrow \mathfrak{C}$ is well-defined. We also have $g_0 \circ f = g_1$, and f and g_0 preserve the parameters in X . Given $w \in \mathfrak{C}$, we write $C_w^0 := g_0^{-1}(w)$, and we can check that f maps C_w^1 onto C_w^0 , that $|C_w^0| = |C_w^1|$ in case $|C_w^1| \leq 2K$, and that $|C_w^0| = 2K$ in case $|C_w^1| > 2K$. The situation is

depicted in Figure 4.




Figure 4: The three clusters under examination. Here the cluster \mathfrak{C} contains three elements and, accordingly, both \mathfrak{C}_0 and \mathfrak{C}_1 are divided into three cells.

We then extend g_0 to a parametric bounded morphism from $\hat{\mathfrak{C}}_0$ to $\hat{\mathfrak{C}}$. Consider a non-degenerated pre-cluster $c_0 \subseteq \mathfrak{C}_0$. We write $m_0 := |c_0|$ and $n_0 := |\mathfrak{C}_0 \setminus c_0|$. We write $c := g_0[c_0]$, and we construct $c_1 \subseteq \mathfrak{C}_1$ non-degenerated such that $g_1[c_1] = c$. Below we consider three cases. In case 1, \mathfrak{C}_0 and \mathfrak{C}_1 are isomorphic, so this is easily solved. In case 2, the pair (n_0, m_0) lies in the colored horizontal rectangle of Figure 3. When moving from (n_0, m_0) to the right, one is then guaranteed to stay in this area and to remain outside of $\uparrow\Delta$. This means that if we manage to construct c_1 so that $|\mathfrak{C}_1 \setminus c_1| = |\mathfrak{C}_0 \setminus c_0|$, then c_1 will be non-degenerated. Finally, in case 3, the pair (n_0, m_0) lies in the vertical rectangle, and similarly we will construct c_1 so that $|c_1| = |c_0|$.

1. Suppose that $n_0 \leq K$ and $m_0 \leq K$. This means that $|\mathfrak{C}_0| \leq 2K$. Then for all $w \in \mathfrak{C}$ we have $|C_w^0| \leq 2K$, and thus $|C_w^0| = |C_w^1|$ by construction. Hence, f is a bijection from \mathfrak{C}_1 to \mathfrak{C}_0 . Then, since c_0 is non-degenerated, so is $c_1 := f^{-1}[c_0]$, as a consequence of Proposition 2.16. We also have $g_1[c_1] = g_0 \circ f[c_1] = g_0[c_0] = c$.
2. Suppose that $n_0 \geq K$. Let $c_1 \subseteq \mathfrak{C}_1$ be such that $|c_1| = |c_0|$ and $f[c_1] = c_0$. Let $n_1 := |\mathfrak{C}_1 \setminus c_1|$. Note that since $|c_1| = |c_0|$ and $|\mathfrak{C}_0| \leq |\mathfrak{C}_1|$, we have $n_0 \leq n_1$. Suppose that c_1 is degenerated. Then by Proposition 2.16, there exists $(n, m) \in \Delta$ such that $(n, m) \leq (n_1, m_0)$. Since $n_0 \geq K \geq N_{\mathbf{L}} \geq n$, it follows that $(n, m) \leq (n_0, m_0)$, contradicting the fact that c_0 is non-degenerated. Thus c_1 is non-degenerated. We also have $g_1[c_1] = g_0 \circ f[c_1] = g_0[c_0] = c$.
3. Otherwise we have $n_0 \leq K$ and $m_0 \geq K$. Given $w \in \mathfrak{C}$, we are going to define $A_w \subseteq C_w^1$ such that $|A_w| = |C_w^0 \setminus c_0|$. If $w \notin c$, then we have $C_w^0 \cap c_0 = \emptyset$, whence $|C_w^0| = |C_w^0 \setminus c_0| \leq n_0 \leq 2K$. By construction, it follows that $|C_w^1| = |C_w^0|$, and we set $A_w := C_w^1$. If instead $w \in c$, then we select any $A_w \subseteq C_w^1$ such that $|A_w| = |C_w^0 \setminus c_0|$ – which exists since $|C_w^1| \geq |C_w^0|$. Now let $c_1 := \mathfrak{C}_1 \setminus (\bigcup_{w \in \mathfrak{C}} A_w)$ and $m_1 = |c_1|$. By construction we have

$$|\mathfrak{C}_1 \setminus c_1| = \left| \bigcup_{w \in \mathfrak{C}} A_w \right| = \sum_{w \in \mathfrak{C}} |A_w| = \sum_{w \in \mathfrak{C}} |C_w^0 \setminus c_0| = |\mathfrak{C}_0 \setminus c_0| = n_0.$$

Suppose that c_1 is degenerated. Then by Proposition 2.16 there exists $(n, m) \in \Delta$ such that $(n, m) \leq (n_0, m_1)$. Since $m_0 \geq K \geq M_{\mathbf{L}} \geq m$, it follows that $(n, m) \leq (n_0, m_0)$, contradicting the fact that c_0 is non-degenerated. Thus, c_1 is non-degenerated.

We prove that $g_1[c_1] = c$. Let $w \in c$. Then $C_w^0 \cap c_0 \neq \emptyset$, or equivalently $C_w^0 \setminus c_0 \subset C_w^0$. It follows that $|A_w| = |C_w^0 \setminus c_0| < |C_w^0| \leq |C_w^1|$. Hence, there exists $u \in C_w^1 \setminus A_w$, and by construction

we have $u \in c_1$ and $g_1(u) = w$. Conversely, let $w \in \mathfrak{C} \setminus c$. Then $A_w = C_w^1 = g_1^{-1}(w)$, and since $A_w \cap c_1 = \emptyset$ there exists no $u \in c_1$ such that $g_1(u) = w$.

Next, let $P \subseteq X \cap \mathcal{P}$. We write $\theta_P := (\bigwedge_{p \in P} p) \wedge (\bigwedge_{p \in X \cap \mathcal{P} \setminus P} p)$. Since c_1 is non-degenerated, the pre-cluster $c = g_1[c_1]$ has at least one quasi-predecessor by Proposition 2.12. By Lemma 2.15, it follows that c has at most $2^{|X \setminus P|}$ quasi-predecessors within $\langle \theta_P \rangle_X \cap \hat{\mathfrak{C}}$, and that c_0 has at least $2^{|Y \setminus P|} - |\mathfrak{C}_0|$ predecessors within $\langle \theta_P \rangle_Y$. From $2^{|Y \setminus X|} \geq 2K|\mathfrak{C}|$, we then get

$$2^{|Y \setminus P|} = 2^{|Y \setminus X| + |X \setminus P|} = 2^{|Y \setminus X|} \cdot 2^{|X \setminus P|} \geq (2K|\mathfrak{C}| + 1) \cdot 2^{|X \setminus P|} \geq 2K|\mathfrak{C}| + 2^{|X \setminus P|} \geq |\mathfrak{C}_0| + 2^{|X \setminus P|}$$

and thus $2^{|Y \setminus P|} - |\mathfrak{C}_0| \geq 2^{|X \setminus P|}$. Therefore, the predecessors of c_0 within $\langle \theta_P \rangle_Y$ are more numerous than the quasi-predecessors of c within $\langle \theta_P \rangle_X \cap \hat{\mathfrak{C}}$. We then extend g_0 so that to map the former onto the latter, and by construction g_0 preserves the parameters in X .

Finally, we extend f to a parametric bounded morphism from $\hat{\mathfrak{C}}_1$ to $\hat{\mathfrak{C}}_0$. Consider $c_1 \subseteq \mathfrak{C}_1$ and $u \triangleright c_1$. Let $m_1 := |c_1|$ and $n_1 := |\mathfrak{C}_1 \setminus c_1|$. Let $c_0 := f[c_1]$, $m_0 := |c_0|$ and $n_0 := |\mathfrak{C}_0 \setminus c_0|$. Since f is surjective, we have $(n_0, m_0) \leq (n_1, m_1)$. Thus, c_0 is non-degenerated (as a consequence of Proposition 2.16 again). We also have $g_0[c_0] = g_0 \circ f[c_1] = g_1[c_1]$, and $g_1(u) \triangleright g_1[c_1]$. So by the above construction of g_0 , there exists $v \triangleright c_0$ such that $g_0(v) = g_1(u)$. We then set $f(u) := v$, and we have $g_0 \circ f(u) = g_1(u)$.

It is then clear that f and g_0 are two parametric bounded morphisms and that $g_0 \circ f = g_1$. We also have $g_0[\hat{\mathfrak{C}}_0] \subseteq \hat{\mathfrak{C}}$. We thus obtain a primitive Y -reflection $(\hat{\mathfrak{C}}_0, g_0)$ of $\hat{\mathfrak{C}}$ such that $(\hat{\mathfrak{C}}_0, g_0) \preceq (\hat{\mathfrak{C}}_1, g_1)$, and we are done. \square

We will also need a number of combinatorial properties.

Lemma 3.7. *Let us write $n := |X|$.*

1. *If \mathfrak{C} is a cluster in \mathfrak{F}_X , then $|\mathfrak{C}| \leq 2^n$.*
2. *If \mathfrak{C} is a cluster in \mathfrak{F}_X , then $|\hat{\mathfrak{C}}| \leq 2^{2^{n+1}}$.*
3. *There are at most 2^n isolated points in \mathfrak{F}_X .*
4. *There are at most 2^{2^n} clusters in \mathfrak{F}_X .*
5. *We have $|\mathfrak{F}_X| \leq 2^{2^{n+3}}$.*
6. *Let $Y \subseteq \mathcal{X}$ be finite. Let $m := |Y|$. Then every generated subframe of \mathfrak{F}_X has at most $2^{2^{2^{n+m+4}}}$ primitive Y -reflections.*

We are now ready to prove the announced result. Below, the constant $\alpha_{\mathbf{L}}$ is defined as $\alpha_{\mathbf{L}} := \max \{\log(\log(N_{\mathbf{L}})), \log(\log(M_{\mathbf{L}}))\}$, with the convention that $\log(-1) = \log(0) = 0$.

Proposition 3.8. *Let $X \subseteq \mathcal{X}$ be finite, and \mathfrak{F} be a generated subframe of \mathfrak{F}_X . Let $n := |X|$, and $Y \subseteq \mathcal{X}$ be such that $X \subseteq Y$ and $Y \cap \mathcal{P} = X \cap \mathcal{P}$ and $|Y| = 2 \cdot 2^{2^{n+\alpha_{\mathbf{L}}+5}}$. Then \mathfrak{F} has Y -bounded total reflection.*

Proof. Let $\mathfrak{F}_1, \dots, \mathfrak{F}_k$ be the collection of all primitive closed subframes of \mathfrak{F} . Let $Y' \subseteq \mathcal{X}$ be such that $X \subseteq Y'$ and $Y' \cap \mathcal{P} = X \cap \mathcal{P}$ and $|Y'| = 3n + 2^{\alpha_{\mathbf{L}}} + 2$. We write $m := |Y'|$. Then we have $|Y' \setminus X| = m - n = (n + 2) + (n + 2^{\alpha_{\mathbf{L}}}) \geq n + 2 + \log K$, whence $2^{|Y' \setminus X|} \geq 2^{n+2+\log K} \geq 2K \cdot 2^n + 1$. So, given a cluster \mathfrak{C} in \mathfrak{F}_X , the above inequalities together with Lemma 3.7 ensure that we have $2^{|Y' \setminus X|} \geq 2K|\mathfrak{C}| + 1$. Given $1 \leq i \leq k$, the frame \mathfrak{F}_i is either trivial or a pinned cluster in \mathfrak{F}_X , and by Lemma 3.5 and Lemma 3.6, it follows that \mathfrak{F}_i has bounded primitive Y' -reflection. So the next step is to select $Y \supseteq Y'$ big enough so that \mathfrak{F}_Y contains every primitive Y' -reflection of every \mathfrak{F}_i . To

this end, let \mathcal{R} be the set of pairs of the form (i, \mathfrak{H}', h') where $1 \leq i \leq k$ and (\mathfrak{H}', h') is a primitive Y' -reflection of \mathfrak{F}_i . We introduce a set $Y^* \subseteq \mathcal{X} \setminus (\mathcal{P} \cup Y')$ such that $|Y^*| = \lceil \log |\mathcal{R}| \rceil$. In this case we have $|\mathcal{R}| \leq 2^{|Y^*|}$ and so there exists an injection $\Omega : \mathcal{R} \rightarrow 2^{Y^*}$. By Lemma 3.7, there are at most $2^n + 2^{2^n} \leq 2^{2^{n+1}}$ primitive closed subframes of \mathfrak{F}_X , and each of them has at most $2^{2^{m+n+4}}$ primitive Y' -reflections. Hence, we have $|\mathcal{R}| \leq 2^{2^{n+1}} \cdot 2^{2^{2^{m+n+4}}}$, and thus

$$|Y^*| \leq 2^{n+1} + 2^{2^{m+n+4}} = 2^{n+1} + 2^{2^{(3n+2^{\alpha_L+1})+n+4}} = 2^{n+1} + 2^{2^{4n+2^{\alpha_L+5}}}.$$

As a result we have $|Y' \cup Y^*| \leq 3n + 2^{\alpha_L} + 2 + 2^{n+1} + 2^{2^{4n+2^{\alpha_L+5}}} \leq 2 \cdot 2^{2^{4n+2^{\alpha_L+5}}}$, and so we can select $Y \subseteq \mathcal{X}$ such that $Y' \cup Y^* \subseteq Y$ and $Y \cap \mathcal{P} = X \cap \mathcal{P}$ and $|Y| = 2 \cdot 2^{2^{4n+2^{\alpha_L+5}}}$.

Now we prove that \mathfrak{F} has Y -bounded total reflection. So let (\mathfrak{F}_Z, h) be a total reflection of \mathfrak{F} . Up to applying Corollary 2.27, we can assume that $X \cap \mathcal{P} = Z \cap \mathcal{P}$. Let \mathfrak{H} be a primitive closed subframe of \mathfrak{F}_Z . Then there exists $i \in [1, k]$ such that $h[\mathfrak{H}] \subseteq \mathfrak{F}_i$, and \mathfrak{H} induces a primitive reflection $(\mathfrak{H}, h_{\mathfrak{H}})$ of \mathfrak{F}_i . Since \mathfrak{F}_i has bounded primitive Y' -reflection, there exists a Y' -reflection $(\mathfrak{H}', h'_{\mathfrak{H}})$ of \mathfrak{F}_i such that $(\mathfrak{H}', h'_{\mathfrak{H}}) \preceq (\mathfrak{H}, h_{\mathfrak{H}})$. Then, by Lemma 3.3 there is some primitive Y -reflection $(\mathfrak{G}_{\mathfrak{H}}, g_{\mathfrak{H}})$ of $\mathfrak{F}_{\mathfrak{H}}$ such that $(\mathfrak{G}_{\mathfrak{H}}, g_{\mathfrak{H}}) \preceq (\mathfrak{H}', h'_{\mathfrak{H}})$ and $\mathfrak{G}_{\mathfrak{H}} \subseteq \langle y \rangle_Y$ for all $y \in \Omega(i, \mathfrak{H}', h')$ and $\mathfrak{G}_{\mathfrak{H}} \subseteq \langle \neg y \rangle_Y$ for all $y \in Y^* \setminus \Omega(i, \mathfrak{H}', h')$ – since Ω is injective, this last condition ensures that all the $\mathfrak{G}_{\mathfrak{H}}$'s are pairwise disjoint. The situation is depicted in Figure 5.




Figure 5: Construction of a total reflection

By Lemma 3.2, it follows that $(\mathfrak{G}_{\mathfrak{H}}, g_{\mathfrak{H}}) \preceq (\mathfrak{H}, h_{\mathfrak{H}})$, and so there exists a parametric bounded morphism $f_{\mathfrak{H}} : \mathfrak{H} \rightarrow \mathfrak{G}_{\mathfrak{H}}$ such that $g_{\mathfrak{H}} f_{\mathfrak{H}} = h_{\mathfrak{H}}$. The union of the $f_{\mathfrak{H}}$'s induces a parametric bounded morphism $f : \mathfrak{F}_Z \rightarrow \mathfrak{F}_Y$, and the union of the $g_{\mathfrak{H}}$'s induces a parametric bounded morphism $g' : \mathfrak{G} \rightarrow \mathfrak{F}_X$, where \mathfrak{G} is the union of all the $\mathfrak{G}_{\mathfrak{H}}$'s. We then have $g' f = h$.

There remains to extend g' to a total parametric bounded morphism $g : \mathfrak{F}_Y \rightarrow \mathfrak{F}_X$. Since $Z \cap \mathcal{P} = X \cap \mathcal{P} = Y \cap \mathcal{P}$, we can select an arbitrary parametric substitution $\sigma : \mathcal{L}_Z \rightarrow \mathcal{L}_Y$, so

that we have $\sigma^* : \mathfrak{F}_Y \rightarrow \mathfrak{F}_Z$. Next, the domain of g' is a closed subframe of \mathfrak{F}_Y , and thus so is its complement. So we define $g : \mathfrak{F}_Y \rightarrow \mathfrak{F}_X$ by $g(w) := g'(w)$ if w belongs to the domain of g' , and $g(w) := h(\sigma^*(w))$ otherwise. It is then clear that g is a parametric bounded morphism and that $\text{Im } g \subseteq \mathfrak{F}$. Hence, (\mathfrak{F}_Y, g) is a total reflection of \mathfrak{F} . It is also immediate that $gf = h$. Therefore $(\mathfrak{F}_X, g) \preceq (\mathfrak{F}_Y, h)$, and we are done. \square

4 Computing a most general unifier via filtering

Now that we have obtained the dual properties that we are looking for, we lift them to the syntactic level, and meet our final goal by applying results from [AABM23]. In this paper, we find a technique allowing to ‘merge’ two unifiers into a more general third unifier – a property known as *filtering*. In particular, if one is given a finite non-empty complete set Σ of unifiers, then one can obtain a MGU by iterating this operation on the elements of Σ , and this simple procedure will give birth to our algorithm. Let us set these ideas formally. We fix $\varphi \in \mathcal{L}$ and $X := \text{Var}(\varphi)$.

Definition 4.1. Let $\sigma : \mathcal{L}_X \rightarrow \mathcal{L}_Y$ and $\tau : \mathcal{L}_X \rightarrow \mathcal{L}_Z$ be two substitutions, and t be a fresh variable. The *filtering* of σ and τ is the substitution $\sigma \sqcap \tau : \mathcal{L}_X \rightarrow \mathcal{L}_{Y \cup Z \cup \{t\}}$ defined by

$$(\sigma \sqcap \tau)(x) := ((\Box \Box t \wedge (t \vee \Diamond \top)) \wedge \sigma(x)) \vee ((\Diamond \Diamond \neg t \vee (\neg t \wedge \Box \perp)) \wedge \tau(x))$$

for all $x \in X$.

Proposition 4.2. Let $\sigma : \mathcal{L}_X \rightarrow \mathcal{L}_Y$ and $\tau : \mathcal{L}_X \rightarrow \mathcal{L}_Z$ be two substitutions.

1. If σ and τ are parametric, then so is $\sigma \sqcap \tau$.
2. We have $\sigma \sqcap \tau \preceq \sigma$ and $\sigma \sqcap \tau \preceq \tau$.
3. If σ and τ are two unifiers of φ , then so is $\sigma \sqcap \tau$.

Proof. For 2 and 3, we refer to [AABM23, Prop. 20]. For 1, let us write $\psi := \Box \Box t \wedge (t \vee \Diamond \top)$. If $p \in \mathcal{P} \cap X$ then by assumption we have $(\sigma \sqcap \tau)(p) \equiv_{\mathbf{L}} (\psi \wedge p) \vee (\neg \psi \wedge p) \equiv_{\mathbf{L}} p$, as desired. \square

The following is a consequence of Lemma 2.19 and Proposition 4.2.

Corollary 4.3. Let $\Sigma = \{\sigma_1, \dots, \sigma_n\}$ be a finite non-empty complete set of unifiers of φ . Then $(\sigma_1 \sqcap \sigma_2) \sqcap \dots \sqcap \sigma_n$ is a most general unifier of φ .

Before moving to the main result, we will need an upper bound on the size of modal formulas, up to equivalence in \mathbf{L} .

Lemma 4.4. Let $X \subseteq \mathcal{X}$ be finite and $n := |X|$.

1. For all $w \in \mathfrak{F}_X$, there exists $\theta_w \in \mathcal{L}_X$ such that $|\theta_w| = O(2^n \cdot n)$ and $\theta_w \in u$ iff $w = u$, for all $u \in \mathfrak{F}_X$.
2. For all $\psi \in \mathcal{L}_X$ there exists $\chi \in \mathcal{L}_X$ such that $\psi \equiv_{\mathbf{L}} \chi$ and $|\chi| = O(2^{2^{n+3}} \cdot 2^n \cdot n)$.

Theorem 4.5. If φ is unifiable, then φ admits a most general unifier, which can be effectively computed in space 4-exponential in the size of φ .

Proof. First, since $\mathbf{K5} \subseteq \mathbf{L}$, the validity problem for \mathbf{L} is in NP [HR07], and thus also in **PSPACE**. Hence, there is a procedure $\text{valid}_{\mathbf{L}}(\cdot)$ which decides whether $\vdash_{\mathbf{L}} \psi$ in space polynomial in $|\psi|$, for every $\psi \in \mathcal{L}$. Let \mathfrak{F} be the subframe of \mathfrak{F}_X generated by $\langle \varphi \rangle_X^\infty$. We write $n := |X|$. Let $Y \subseteq \mathcal{X}$ be

such that $X \subseteq Y$ and $Y \cap \mathcal{P} = X \cap \mathcal{P}$ and $|Y| = 2 \cdot 2^{4n + \alpha_{\mathbf{L}} + 5}$. Let Σ_0 be the set of all parametric substitutions from \mathcal{L}_X to \mathcal{L}_Y , up to equivalence. Then since \mathbf{L} is locally tabular, the set Σ_0 is finite, and we consider $\Sigma := \{\sigma \in \Sigma_0 \mid \sigma \text{ is a unifier of } \varphi\}$. By Proposition 3.8, \mathfrak{F} has Y -bounded total reflection, and by duality (Corollary 2.26) it follows that Σ is a complete set of unifiers of φ . If φ is unifiable, then Σ is non-empty, and Corollary 4.3 provides a MGU of φ . We describe this procedure precisely in Algorithm 1, of which we now discuss the complexity. We see that the

```

Input:  $\varphi$ 
 $X \leftarrow \text{Var}(\varphi)$ ;
 $n \leftarrow |X|$ ;
select  $Y \subseteq \mathcal{X}$  such that  $X \subseteq Y$  and  $|Y| = 2 \cdot 2^{4n + \alpha_{\mathbf{L}} + 5}$ ;
 $\sigma \leftarrow \clubsuit$ ;
foreach parametric substitution  $\tau : \mathcal{L}_X \rightarrow \mathcal{L}_Y$  do
  if  $\text{valid}_{\mathbf{L}}(\tau(\varphi))$  then
    if  $\sigma = \clubsuit$  then
       $\sigma \leftarrow \tau$ ;
    else
       $\sigma \leftarrow \sigma \sqcap \tau$ ;
    end
  end
end
Output:  $\sigma$ 

```

Algorithm 1: returns a MGU of φ if there is one, and returns \clubsuit otherwise

program requires three quantities of space: one to store τ , one to store σ , and one used by the execution of $\text{valid}_{\mathbf{L}}$ on $\tau(\varphi)$.

1. The parametric substitution τ can be encoded as a map from X to \mathcal{L}_X . By Lemma 4.4, a formula in \mathcal{L}_Y can be written, up to equivalence, in size $O(2^{2^{|Y|+3}} \cdot 2^{|Y|} \cdot |Y|)$. In addition, since there are arbitrarily many variables to be handled by the program, they cannot be represented by primitive symbols. Instead, they need to be encoded by (say) numbers, and each variable will use up to $\log |Y|$ digits. Hence, τ has size $O(n \cdot 2^{2^{|Y|+3}} \cdot 2^{|Y|} \cdot |Y| \cdot \log |Y|)$.
2. The parametric substitution σ is, in the worst case, the iteration of the operation \sqcap over all of Σ_0 . A formula in \mathcal{L}_X can be identified to a subset of \mathfrak{F}_X . By Lemma 3.7, we have $|\mathfrak{F}_X| \leq 2^{2^{n+3}}$, and so there at most $2^{2^{2^{n+3}}}$ formulas in \mathcal{L}_X , up to equivalence. It follows that $|\Sigma_0| \leq \left(2^{2^{2^{n+3}}}\right)^n$. Further, the size of $\mu \sqcap \mu'$ is polynomial in the size of μ and μ' . Hence, the size of σ is polynomial in

$$2^{n \cdot 2^{2^{n+3}}} \cdot n \cdot 2^{2^{|Y|+3}} \cdot 2^{|Y|} \cdot |Y| \cdot \log |Y|.$$

3. Given a parametric substitution $\tau : \mathcal{L}_X \rightarrow \mathcal{L}_Y$, the size of the formula $\tau(\varphi)$ is polynomial in $|\varphi| \cdot \max\{|\tau(x)| : x \in X\}$. As mentioned in 1, $\tau(x)$ is written in space $O(2^{2^{|Y|+3}} \cdot 2^{|Y|} \cdot |Y| \cdot \log |Y|)$ for all $x \in X$. Hence, the amount of space used by $\text{valid}_{\mathbf{L}}$ on $\tau(\varphi)$ will be polynomial in $|\varphi| \cdot 2^{2^{|Y|+3}} \cdot 2^{|Y|} \cdot |Y| \cdot \log |Y|$.

In total, the space used by Algorithm 1 is polynomial in $(n + 2^{n \cdot 2^{2^{n+3}}} \cdot n + |\varphi|) \cdot 2^{2^{|Y|+3}} \cdot 2^{|Y|} \cdot |Y| \cdot \log |Y|$ and thus 4-exponential in the size of φ . \square

Combining this theorem with Proposition 2.22, we also obtain a complexity result for the problems of unification and admissibility.

Corollary 4.6. *Unification and admissibility in \mathbf{L} are decidable in 4-exponential space.*

Though the complexity of our algorithm is huge, extra assumptions on \mathbf{L} may enable simplifications. Some logics feature what we call *concise unification*, which basically allows us to always take $Y = X$, thus lowering the space complexity down to 3-exponential. In the next section we investigate this condition.

5 Concise, transparent and projective unification

In this section we investigate special types of unary unification. Let $\varphi \in \mathbf{L}$ and $X := \text{Var}(\varphi)$. A unifier $\sigma : \mathcal{L}_X \rightarrow \mathcal{L}_X$ of φ is said to be *transparent* if we have $\tau\sigma \simeq \tau$ for all unifiers τ of φ [Wro95]. It is said to be *projective* if we have $\varphi \vdash_{\mathbf{L}} \sigma(x) \leftrightarrow x$ for all $x \in X$ [Ghi97]. The formula φ is *projective* (resp. *transparent*) if it admits a projective unifier.

The logic \mathbf{L} is said to be *transparent* (resp. *projective*) every unifiable formula is projective (resp. transparent). It is said to be *parameter-free projective* if every unifiable parameter-free formula is projective. Finally, the logic \mathbf{L} has *concise unification* if every unifiable formula φ admits a MGU $\sigma : \mathcal{L}_X \rightarrow \mathcal{L}_X$, where $X := \text{Var}(\varphi)$ – that is, a MGU that does not use more variables than those already present in φ .

Proposition 5.1. *Let $\varphi \in \mathbf{L}$ and $X := \text{Var}(\varphi)$. Let $\sigma : \mathcal{L}_X \rightarrow \mathcal{L}_X$ be a unifier of φ .*

1. *If σ is transparent then σ is a most general unifier of φ .*
2. *If σ is projective then σ is transparent.*

Proof. Item 1 is immediate. For Item 2, let τ be a unifier of φ and let $x \in X$. From $\varphi \vdash_{\mathbf{L}} \sigma(x) \leftrightarrow x$, it follows that $\tau(\varphi) \vdash_{\mathbf{L}} \tau(\sigma(x) \leftrightarrow x)$, whence $\tau(\varphi) \vdash_{\mathbf{L}} \tau(\sigma(x)) \leftrightarrow \tau(x)$. Then since $\vdash_{\mathbf{L}} \tau(\varphi)$, we obtain $\vdash_{\mathbf{L}} \tau(\sigma(x)) \leftrightarrow \tau(x)$. Therefore $\tau\sigma \simeq \tau$. \square

As a result, it is clear that all projective logics are parameter-free projective and transparent, and that all transparent logics are unitary and have concise unification. Soon we will see that some of these implications are proper. First, we exhibit the dual properties that correspond to transparency and projectivity.

Definition 5.2. Let $\varphi \in \mathbf{L}$ and $X := \text{Var}(\varphi)$. A dual unifier $g : \mathfrak{F}_Y \rightarrow \mathfrak{F}_X$ of φ is said to be:

1. *transparent* if for all dual unifiers $f : \mathfrak{F}_Y \rightarrow \mathfrak{F}_X$ of φ , we have $gf = f$;
2. *projective* if for all $w \in \langle \varphi \rangle_Y^\infty$, we have $\sigma^*(w) = w$.

Proposition 5.3. *Let $\varphi \in \mathbf{L}$ and $X := \text{Var}(\varphi)$. Let $\sigma : \mathcal{L}_X \rightarrow \mathcal{L}_X$ be a unifier of φ . Then σ is transparent iff σ^* is transparent, and σ is projective iff σ^* is projective.*

Proof. The first equivalence is a simple consequence of Proposition 2.25. For the second equivalence, we refer to [BG22, Th. 4.5]. \square

The projective extensions of $\mathbf{K5}$ were characterized in [BG22], but in the parameter-free setting only. Below we provide a similar result for projective parametric unification.

Theorem 5.4 ([BG22, Th. 5.5]). *The parameter-free projective extensions of **K5** are exactly the extensions of **K45**. More precisely, if **K45** $\not\subseteq$ **L**, then the formula $\diamond x \rightarrow \diamond \diamond x$ is not projective.*

Theorem 5.5. *The projective extensions of **K5** are exactly the extensions of **Kt5**. More precisely, if **Kt5** $\not\subseteq$ **L**, then the formula $(p \wedge \diamond \top) \rightarrow \diamond(p \vee x)$ with parameter p is not projective.*

Proof. From right to left, suppose that **Kt5** \subseteq **L**. Let $\varphi \in \mathbf{L}$ and $X := \text{Var}(\varphi)$, and assume that φ is unifiable. Then there exists a unifier $\sigma : \mathcal{L}_X \rightarrow \mathcal{L}_Y$ of φ , and we can assume that $X \subseteq Y$. By Lemma 2.18, we also assume that $X \cap \mathcal{P} = Y \cap \mathcal{P}$. This allows selecting an arbitrary parametric substitution $\tau : \mathcal{L}_Y \rightarrow \mathcal{L}_X$. Dually, we thus have a dual unifier $\sigma^* : \mathfrak{F}_Y \rightarrow \mathfrak{F}_X$ of φ , and a parametric bounded morphism $\tau^* : \mathfrak{F}_X \rightarrow \mathfrak{F}_Y$. Since **Kt5** \subseteq **L**, it is easy to see that every point in \mathfrak{F}_X is either isolated, or belongs to a cluster. In this case, $\langle \varphi \rangle_X^\infty$ is then a closed subframe of \mathfrak{F}_X . We now define a projective dual unifier $f : \mathfrak{F}_X \rightarrow \mathfrak{F}_X$ of φ by

$$f(w) := \begin{cases} w & \text{if } w \in \langle \varphi \rangle_X^\infty, \\ \sigma^* \tau^*(w) & \text{otherwise,} \end{cases}$$

for all $w \in \mathfrak{F}_X$. We first check that f is a bounded morphism. Suppose that $w R_X u$. If $w \in \langle \varphi \rangle_X^\infty$, then $u \in \langle \varphi \rangle_X^\infty$ as well, so $f(w) R_X f(u)$ is immediate. Conversely, assume that $w \notin \langle \varphi \rangle_X^\infty$. Then $u \notin \langle \varphi \rangle_X^\infty$ as well. Since τ^* and σ^* are bounded morphisms, it follows that $\sigma^* \tau^*(w) R_X \sigma^* \tau^*(u)$, whence $f(w) R_X f(u)$. Now suppose that $f(w) R_X u'$. If $w \in \langle \varphi \rangle_X^\infty$ then $w = f(w) R_X u'$ and so $u \in \langle \varphi \rangle_X^\infty$ too, whence $w R_X u'$ with $f(u') = u'$. Otherwise $f(w) = \sigma^* \tau^*(w)$, and since τ^* and σ^* are bounded morphisms, there exists $u \in \mathfrak{F}_X$ such that $w R_X u$ and $\sigma^* \tau^*(u) = u'$. Then $u \notin \langle \varphi \rangle_X^\infty$, whence $f(u) = \sigma^* \tau^*(u) = u'$. It is also clear that f is parametric, that $\text{Im } f \subseteq \langle \varphi \rangle_X^\infty$, and that $f(w) = w$ for all $w \in \langle \varphi \rangle_X^\infty$. Therefore, f is a projective dual unifier of φ . By Proposition 5.3, it follows that φ is projective.

Conversely, assume that **Kt5** $\not\subseteq$ **L**. Let $\varphi := (p \wedge \diamond \top) \rightarrow \diamond(p \vee x)$, with $p \in \mathcal{P}$ and $x \notin \mathcal{P}$. We write $X := \{x, p\}$. Then we have $\not\vdash_{\mathbf{L}} \varphi$, for otherwise we would obtain $\vdash_{\mathbf{L}} p \wedge \diamond \top \rightarrow \diamond p$ after substituting p for x , contradicting **Kt5** $\not\subseteq$ **L**. Further, we observe that φ is unifiable, as it suffices to substitute \top for x . Toward a contradiction, assume that φ is projective. Then by Proposition 5.3 there exists a dual unifier $f : \mathfrak{F}_X \rightarrow \mathfrak{F}_X$ of φ such that $f(w) = w$ for all $w \in \langle \varphi \rangle_X^\infty$. Since $\not\vdash_{\mathbf{L}} \varphi$, there exists $w \in \langle p \wedge \diamond \top \wedge \neg \diamond(p \vee x) \rangle_X$. Since $w \in \langle p \rangle_X$ and f is parametric, we then have $f(w) \in \langle p \rangle_X$. Since $f(w) \in \langle \varphi \rangle_X^\infty$ by construction, it follows that $f(w) \in \langle \diamond(p \vee x) \rangle_X$. Hence, there exists $u \in \langle p \rangle_X$ such that $f(w) R_X u$. Then by the backward condition, there is $v \in \mathfrak{F}_X$ such that $f(v) = u$ and $w R_X v$. Then v belongs to a cluster, and so it is clear that $v \in \langle \varphi \rangle_X^\infty$, whence $f(v) = v$. Then $u = v$ and $w R_X u$, contradicting $w \in \langle \neg \diamond p \rangle_X$. We conclude that **L** is not projective. \square

Next, we exhibit infinitely many extensions of **K5**, including **K5** and **KD5** themselves, that are transparent but not (parameter-free) projective. Given $m \geq 1$, we introduce **K5W** $_m := \mathbf{K5} + \text{Ax}_{-1,m}$ and **KD5W** $_m := \mathbf{KD5} + \text{Ax}_{-1,m}$. The set Δ associated to these two logics is thus $\{(-1, m)\}$, which means that every cluster in their canonical frame has size $< m$ – hence the ‘**W**’, which stands for ‘bounded width’. We also write **K5W** $_\infty := \mathbf{K5}$ and **KD5W** $_\infty := \mathbf{KD5}$.

Proposition 5.6. *If $m \geq 1$ (including $m = \infty$), then **K5W** $_m$ and **KD5W** $_m$ are transparent.*

Proof. Let $\mathbf{L} := \mathbf{K5W}_m$. Let $\varphi \in \mathcal{L}$ and assume that φ is unifiable. Let $X := \text{Var}(\varphi)$. Given $w \in \mathfrak{F}_X$, we write $P_w := \{p \in \mathcal{P} \cap X : w \in \langle p \rangle_X\}$. A cluster \mathfrak{C} in \mathfrak{F}_X is said to be *realizable* if there

exists a dual unifier $h : \mathfrak{F}_Y \rightarrow \mathfrak{F}_X$ of φ such that $\mathfrak{C} \subseteq h[\mathfrak{F}_Y]$.⁴ In this case, we claim that for all non-degenerated pre-clusters $c \subseteq \mathfrak{C}$ and $P \subseteq \mathcal{P}$, there exists $u \in \langle \varphi \rangle_X^\infty$ such that $u \triangleright c$ and $P_u = P$. Indeed, let c' be a pre-cluster in \mathfrak{F}_Y such that $h[c'] \subseteq c$ and $|c'| = |c|$. Then c' is non-degenerated (by construction of \mathbf{L}), and by Lemma 2.15 there exists $w \triangleright c'$ such that $P_w = P$. Since h is a parametric bounded morphism, it follows that $h(w) \triangleright c$ and $P_{h(w)} = P$. Since h is a dual unifier of φ , we also have $h(w) \in \langle \varphi \rangle_X^\infty$, as desired.

Next, since φ is unifiable, it admits a dual unifier $g : \mathfrak{F}_Y \rightarrow \mathfrak{F}_X$, and we can assume that $X \subseteq Y$. By Corollary 2.27, we can also assume that $X \cap \mathcal{P} = Y \cap \mathcal{P}$. Hence, we can select an arbitrary parametric substitution $\tau : \mathcal{L}_Y \rightarrow \mathcal{L}_X$. We construct a transparent dual unifier $f : \mathfrak{F}_X \rightarrow \mathfrak{F}_X$ of φ as follows.

1. First, let \mathfrak{C} be a realizable cluster in \mathfrak{F}_X . For every $w \in \hat{\mathfrak{C}} \cap \langle \varphi \rangle_X^\infty$, we set $f(w) := w$. Note that since \mathfrak{C} is realizable, we have $\mathfrak{C} \subseteq \langle \varphi \rangle_X^\infty$, and so this case already covers all the elements of \mathfrak{C} . If instead $w \in \hat{\mathfrak{C}} \setminus \langle \varphi \rangle_X^\infty$, we have $w \triangleright c$ for some $c \subseteq \mathfrak{C}$. As shown earlier, there exists $u \in \langle \varphi \rangle_X^\infty$ such that $u \triangleright c$ and $P_u = P_w$, and we set $f(w) := u$.
2. Next, let \mathfrak{C} be a non-realizable cluster in \mathfrak{F}_X . Then for all $w \in \hat{\mathfrak{C}}$, we set $f(w) := g(\tau^*(w))$.
3. Finally, let w be isolated in \mathfrak{F}_X . We set $f(w) := w$ in case $w \in \langle \varphi \rangle_X^\infty$, and $f(w) := g(\tau^*(w))$ otherwise.

First, it is easy to check that f is a parametric bounded morphism. Now, consider $w \in \mathfrak{F}_X$. If $w \in \hat{\mathfrak{C}}$ for some realizable cluster \mathfrak{C} , then we already know that $f(w) \in \langle \varphi \rangle_X^\infty$. If $w \in \hat{\mathfrak{C}}$ for some non-realizable cluster \mathfrak{C} , then we have $f(w) = g(\tau^*(w)) \in \langle \varphi \rangle_X^\infty$ by assumption on g . If w is isolated, the reasoning is the same. This proves that f is a dual unifier of φ , and there remains to check that f is transparent. Consider a dual unifier $h : \mathfrak{F}_Z \rightarrow \mathfrak{F}_X$ of φ . Given $w \in \mathfrak{F}_X$, the element $h(w)$ belongs to $\langle \varphi \rangle_X^\infty$ and is either initial, or a member of some realizable cluster. By construction, it follows that $f(h(w)) = h(w)$, as desired. Finally, if $\mathbf{L} = \mathbf{KD5W}_m$ then the proof is exactly the same, except that we skip step 3 above. \square

Corollary 5.7. *There are infinitely many transparent extensions of $\mathbf{K5}$ that are not parameter-free projective.*

Proof. We know that the logics $\mathbf{K5W}_m$ and $\mathbf{KD5W}_m$, where $m \geq 3$ (including $m = \infty$), are transparent by Proposition 5.6. Let \mathbf{L} be one of them, and let $X \subseteq \mathcal{X}$ be non-empty. Then by Proposition 2.16, there exists a cluster $\mathfrak{C} = \{w, u\}$ of size 2 in \mathfrak{F}_X . There also exists $v \triangleright \{w\}$, and we have $vRuRu$ but not vRu . This proves that \mathbf{L} is not an extension of $\mathbf{K45}$, and thus not parameter-free projective by Theorem 5.4. Finally, it is clear that all of these logics are pairwise distinct, and thus in infinite number. \square

These results show the abundance of concise logics, but it would also be interesting to make sure that *not* all logics are concise – otherwise, the work of section 3 would be plain useless. To prove this, we identify a necessary condition on Δ for concise unification.

Proposition 5.8. *Suppose that \mathbf{L} has concise unification. Then either $\mathbf{K45} \subseteq \mathbf{L}$, or there is no $n \geq -1$ such that $(n, 1) \in \uparrow\Delta$ and $(-1, n + 1) \notin \uparrow\Delta$.*

Proof. Toward a contradiction, assume that $\mathbf{K45} \not\subseteq \mathbf{L}$ and that $(n, 1) \in \uparrow\Delta$ and $(-1, n + 1) \notin \uparrow\Delta$ for some $n \geq -1$. From $\mathbf{K45} \not\subseteq \mathbf{L}$ it follows that $(1, 1) \notin \uparrow\Delta$ (see Table 2). Hence, we also have

⁴Note: this amounts to say that $\hat{\mathfrak{C}}$ admits a reflection.

$n \geq 2$. Now let us fix $X \subseteq \mathcal{X} \setminus \mathcal{P}$ such that $|X| = 1$. We have $(-1, 2) \leq (-1, 3) \leq (-1, n+1)$, so $(-1, 2)$ and $(-1, 3)$ do not belong to $\uparrow\Delta$. Hence, by Proposition 2.16, there exist a one-element cluster $\{v_1\}$ and a two-element cluster $\{w, u\}$ in \mathfrak{F}_X . In addition, since $(1, 1) \notin \uparrow\Delta$, the pre-cluster $\{w\}$ is non-degenerated, so there exists $w_0 \triangleright \{w\}$.

If $\mathsf{D}^? = \diamond\top$, we select a formula $\varphi \in \mathcal{L}_X$ such that $\langle \varphi \rangle_X = \{v_1, w_0, w, u\}$. Otherwise, there exists an isolated point v_0 in \mathfrak{F}_X , and we choose $\langle \varphi \rangle_X = \{v_0, v_1, w_0, w, u\}$ instead. In both cases we observe that $\langle \varphi \rangle_X$ is upward closed, in the sense that $t \in \langle \varphi \rangle_X$ and $tR_X t'$ implies $t' \in \langle \varphi \rangle_X$. In this case, it is easy to check that we have $\langle \varphi \rangle_X^\infty = \langle \varphi \rangle_X$. We can then define a dual unifier $f : \mathfrak{F}_X \rightarrow \mathfrak{F}_X$ of φ by setting $f(t) := v_0$ in case t is isolated, and $f(t) := v_1$ for any other t . So φ is unifiable, and by assumption it admits a MGU $\sigma : \mathcal{L}_X \rightarrow \mathcal{L}_X$. We write $g := \sigma^*$.

Now let us select $Y \subseteq \mathcal{X} \setminus \mathcal{P}$ finite such that $2^{|Y|} \geq n+1$. Then since $(-1, n+1) \notin \uparrow\Delta$, there exists by Proposition 2.16 a cluster \mathfrak{C} in \mathfrak{F}_Y such that $|\mathfrak{C}| = n+1$. Further, let $v \in \mathfrak{C}$. Since $(n, 1) \in \uparrow\Delta$, the pre-cluster $\{v\}$ is degenerated by Proposition 2.16. We define $h : \mathfrak{F}_Y \rightarrow \mathfrak{F}_X$ as follows: we set $h(v) := u$, $h(t) := w_0$ whenever $t \triangleright c \subseteq \mathfrak{C} \setminus \{v\}$, and $h(t) := w$ for every other $t \in \hat{\mathfrak{C}}$; then we set $h(t) := v_0$ for all t isolated, and $h(t) := v_1$ for every other t . The construction is depicted in Figure 6. The map h is readily checked to be a dual unifier of φ , and thus there exists a bounded morphism $f : \mathfrak{F}_Y \rightarrow \mathfrak{F}_X$ such that $gf = h$. Then there is a cluster \mathfrak{C}' in \mathfrak{F}_X such that $f[\mathfrak{C}] = \mathfrak{C}'$. It follows that $g[\mathfrak{C}'] = gf[\mathfrak{C}] = h[\mathfrak{C}] = \{w, u\}$. Since $|X| = 1$, we have also $|\mathfrak{C}'| \leq 2$, so the only possibility for \mathfrak{C}' is to have size 2. We write $\mathfrak{C}' = \{w', u'\}$ with $g(w') = w$ and $g(u') = u$. Again, there exists $u'_0 \triangleright \{u'\}$, and by Proposition 2.12 we must have $g(u'_0) \triangleright \{u\}$. But then $g(u'_0) \notin \langle \varphi \rangle_X^\infty$, a contradiction. \square

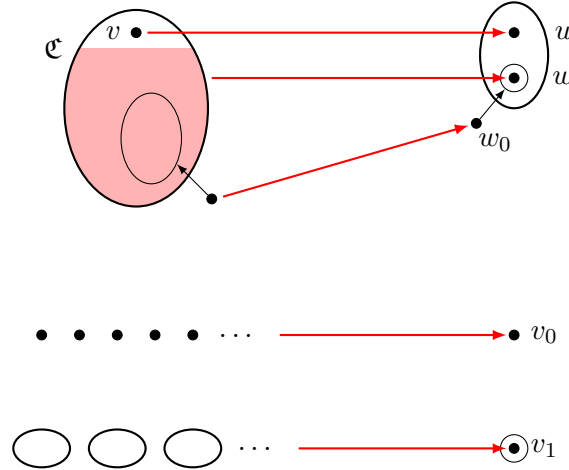


Figure 6: The bounded morphism h (in bold arrows) from \mathfrak{F}_Y to $\langle \varphi \rangle_X^\infty$. The colored area is mapped to w . All isolated points are mapped to v_0 , and all clusters other than \mathfrak{C} are mapped to v_1 .

Corollary 5.9. *There are infinitely many extensions of $\mathbf{K5}$ that do not have concise unification.*

Proof. For all $n \geq 2$, we consider the logic \mathbf{L} defined by $\Delta := \{(n, 1)\}$ and $\mathsf{D}^? = \top$. Then $\mathbf{K45} \not\subseteq \mathbf{L}$, since $(1, 1) \notin \uparrow\Delta$ (see Table 2), and we also have $(n, 1) \in \uparrow\Delta$ and $(-1, n+1) \notin \uparrow\Delta$. So \mathbf{L} falsifies the condition of Proposition 5.8, and thus lacks concise unification. It is also clear that all of these logics are pairwise distinct. \square

6 Conclusion

By thoroughly exploiting the dual properties of Euclidean logics, we proved not only that they have unary unification, but also that most general unifiers can be computed effectively. The weakness of our algorithm is obviously its complexity, which is so gigantic that it prevents any hope of actually implementing it. However, one should keep in mind that our upper bounds are extremely coarse and non-specific, in the sense that we have designed a general method for any extension of **K5**, and that we have not looked at how it refines to special cases. As we have seen, logics with concise unification are better behaved, and even in less lucky cases, there is much room for improvement. Indeed, our complexity estimates rely on the upper bounds of Lemma 3.7, which can be refined depending on the properties of the canonical frame. For instance, if Δ contains the axiom $(-1, m)$ for some $m \geq 1$, then all clusters in \mathfrak{F}_X have size bounded by $m - 1$, instead of $2^{|X|}$. So anyone interested in applying this work should carefully check how every upper bound can be optimized.

We have also provided a characterization of projective logics, along with counter-examples showing that the transparent and projective types are distinct, as are the concise and unary types. It is also interesting to note that the logic **K45** is parameter-free projective but not projective, since it is not an extension of **Kt5** – thus revealing a discordance between the parametric and non-parametric settings. In spite of this advancement, the landscape of Euclidean logics is yet to be fully described: it could be interesting to characterize exactly the transparent/concise extensions of **K5**, and to determine whether the concise and transparent types coincide.

References

- [AABM23] Majid Alizadeh, Mohammad Ardeshir, Philippe Balbiani, and Mojtaba Mojtahedi. Unification types in Euclidean modal logics. *Logic Journal of the IGPL*, 31(3):422–440, 2023.
- [Bal19] Philippe Balbiani. Remarks about the unification type of several non-symmetric non-transitive modal logics. *Logic Journal of the IGPL*, 27(5):639–658, 2019.
- [BdRV01] Patrick Blackburn, Maarten de Rijke, and Yde Venema. *Modal Logic*. Cambridge Tracts in Theoretical Computer Science. Cambridge University Press, Cambridge, 2001.
- [BG11] Franz Baader and Silvio Ghilardi. Unification in modal and description logics. *Logic Journal of IGPL*, 19(6):705–730, 2011.
- [BG17] Philippe Balbiani and Çiğdem Gencer. **KD** is nullary. *Journal of Applied Non-Classical Logics*, 27(3-4):196–205, 2017.
- [BG20] Philippe Balbiani and Çiğdem Gencer. About the unification type of modal logics between **KB** and **KTB**. *Studia Logica*, 108(5):941–966, 2020.
- [BG22] Philippe Balbiani and Quentin Gougeon. Projective unification through duality. In *Advances in Modal Logic (AiML 2022)*, 2022.
- [BS01] Franz Baader and Wayne Snyder. Unification theory. In *Handbook of Automated Reasoning*, pages 439–526. Elsevier, 2001.
- [CZ97] Alexander Chagrov and Michael Zakharyashev. *Modal Logic*. Oxford Logic Guides. Clarendon Press. Consecutive, Oxford, 1997.

- [Ghi97] Silvio Ghilardi. Unification through projectivity. *Journal of Logic and Computation*, 7(6):733–752, 1997.
- [Ghi99] Silvio Ghilardi. Unification in intuitionistic logic. *The Journal of Symbolic Logic*, 64(2):859–880, 1999.
- [Ghi00] Silvio Ghilardi. Best solving modal equations. *Annals of Pure and Applied Logic*, 102(3):183–198, 2000.
- [Ghi04] Silvio Ghilardi. Unification, finite duality and projectivity in varieties of Heyting algebras. *Annals of Pure and Applied Logic*, 127(1-3):99–115, 2004.
- [GS04] Silvio Ghilardi and Lorenzo Sacchetti. Filtering unification and most general unifiers in modal logic. *The Journal of Symbolic Logic*, 69(3):879–906, 2004.
- [HR07] Joseph Y. Halpern and Leandro Chaves Rêgo. Characterizing the NP-PSPACE gap in the satisfiability problem for modal logic. *Journal of Logic and Computation*, 17(4):795–806, 2007.
- [Jeř15] Emil Jeřábek. Blending margins: the modal logic **K** has nullary unification type. *Journal of Logic and Computation*, 25(5):1231–1240, 2015.
- [Kos18] Sławomir Kost. Projective unification in transitive modal logics. *Logic Journal of the IGPL*, 26(5):548–566, 2018.
- [NT85] Michael C. Nagle and Steven K. Thomason. The extensions of the modal logic **K5**. *The Journal of symbolic logic*, 50(1):102–109, 1985.
- [Wro95] Andrzej Wronski. Transparent unification problem. *Reports on Mathematical Logic*, 1995.

A Proofs for Section 2 (Background)

Proof of Proposition 2.12. Suppose that f is a bounded morphism. First, 1 follows immediately from the backward condition. For 2, let C be a cluster in \mathfrak{F} . Let $w', u' \in f[C]$. Then $w' = f(w)$ and $u' = f(u)$ for some $w, u \in C$. Since C is a cluster we have wRu , whence $f(w)R'f(u)$ by the forward condition. Now let $w' \in f[C]$ and assume that $w'R'u'$. Then $w' = f(w)$ for some $w \in C$, and so by the backward condition there exists $u \in W$ such that wRu and $f(u) = u'$. Since C is a cluster we have $u \in C$, whence $f(u) \in f[C]$. This proves that $f[C]$ is a cluster.

Next, suppose that c is a pre-cluster, and that $w \triangleright c$. We prove that for all $u' \in W'$ we have $f(w)R'u'$ iff $u' \in f[c]$. First suppose that $u' \in f[c]$. Then $u' = f(u)$ for some $u \in c$. Then wRu and thus $f(w)R'f(u)$ by the forward condition, whence $f(w)R'u'$. Conversely, suppose that $f(w)R'u'$. Then by the backward condition there exists u such that wRu and $f(u) = u'$. Since $w \triangleright c$ we have $u \in c$ and thus $u' \in f[c]$, as desired. Hence we have either $f(w) \triangleright f[c]$, in case $f(w)$ is initial, or $f(w) \in f[c]$ otherwise. This proves 3.

Now consider the case where $f[c]$ is a proper pre-cluster. First we have $f[c] \subset C'$ for some cluster C' . Then there exists $u' \in C' \setminus f[c]$. Suppose that $f(w)$ is not initial, that is, $f(w) \in C'$. Then $f(w)R'u'$, and so by the backward condition there exists $u \in W$ such that $f(u) = u'$ and wRu . Then $u \in c$, and so $u' = f(u) \in f[c]$, a contradiction. Therefore $f(w)$ is initial in \mathfrak{F}' , and by 3 it follows that $f(w) \triangleright f[c]$.

Conversely, suppose that f satisfies 1, 2, 3 and 4. Assume that wRu . If w is not initial, then w and u belong to a cluster C . By 2, $f[C]$ is a cluster, and thus $f(w)R'f(u)$. Otherwise we have $w \triangleright c$ for some pre-cluster c . By definition we then have $u \in c$. By 3 it follows that $f(w) \in f[c]$ or $f(w) \triangleright f[c]$, and in both cases we obtain $f(w)R'f(u)$.

Now assume that $f(w)R'u'$. By 1, w cannot be isolated. First suppose that w belongs to some cluster C . Then $f(w) \in f[C]$, and by 2 we obtain $u' \in f[C]$. Thus, there exists $u \in C$ such that $f(u) = u'$. We then have wRu . Now suppose that w is initial. Then $w \triangleright c$ for some pre-cluster c . By 3 we then have $f(w) \in f[c]$ or $f(w) \triangleright f[c]$. If $f(w) \triangleright f[c]$, then in particular $u' \in f[c]$. Otherwise we have $f(w) \in f[c]$, and so $f[c]$ must be a cluster according to 4. Thus $u' \in f[c]$ as well. Therefore, there exists $u \in c$ such that $u' = f(u)$, and since $w \triangleright c$ we have wRu , as desired. \square

Proof of Lemma 2.14. Given $w \in \mathfrak{C}$, we will denote by $[w]$ the equivalence class of w modulo \sim . Since $|\mathfrak{C}_{/\sim}| \leq 2^{|Y \setminus P|}$, there exists an injective map $\Omega : \mathfrak{C}_{/\sim} \rightarrow 2^{Y \setminus P}$. Given $w \in \mathfrak{C}$, we then define the Boolean formula $\theta_w \in \mathcal{L}_Y$ as the conjunction of the following literals:

- p where $p \in P$ and $w \in \langle p \rangle_Z$,
- $\neg p$ where $p \in P$ and $w \notin \langle p \rangle_Z$,
- y where $y \in \Omega([w])$,
- $\neg y$ where $y \in Y \setminus (P \cup \Omega([w]))$.

We then introduce a valuation $V : \mathfrak{C} \rightarrow 2^Y$ so that for all $w \in \mathfrak{C}$, we have $\mathfrak{C}, V, w \models \theta_w$. Then, let $\varphi := (\bigwedge_{w \in \mathfrak{C}} \diamond \theta_w) \wedge \square (\bigvee_{w \in \mathfrak{C}} \theta_w)$ and $\Gamma_0 := \{\square \psi \rightarrow \psi : \psi \in \mathcal{L}_Y\}$. We can check that for all $w \in \mathfrak{C}$, we have $\mathfrak{C}, V, w \models \{\varphi\} \cup \Gamma_0$. Hence, $\{\varphi\} \cup \Gamma_0$ is consistent, and so there exists a maximal consistent set $\Gamma \in \mathfrak{F}_Y$ such that $\{\varphi\} \cup \Gamma_0 \subseteq \Gamma$. Since $\Gamma_0 \subseteq \Gamma$, it follows that Γ is reflexive. Therefore, Γ belongs to a cluster \mathfrak{C}' . By construction, there exists, for every $w \in \mathfrak{C}$, some $f(w) \in \mathfrak{C}' \cap \langle \theta_w \rangle_Y$. This defines a map $f : \mathfrak{C} \rightarrow \mathfrak{C}'$, which we prove to be surjective. For consider $w' \in \mathfrak{C}'$. Then there exists $w \in \mathfrak{C}$ such that $w' \in \langle \theta_w \rangle_Y$. Since $f(w) \in \langle \theta_w \rangle_Y$ too, it follows that $w' \sim_Y f(w)$. Since w' and $f(w)$ belong to the same cluster and satisfy the same literals, they satisfy the same formulas in \mathcal{L}_Y , and thus $w' = f(w)$. It remains to prove that f satisfies the required properties.

1. Suppose that $w \sim u$. Then by assumption, we have $w \sim_P u$. In addition, $\Omega([w]) = \Omega([u])$, and so we see that $\theta_w = \theta_u$. Since $f(w) \in \langle \theta_w \rangle_Y$ and $f(u) \in \langle \theta_u \rangle_Y$, it follows that $f(w) \sim_Y f(u)$, and reasoning as above we obtain $f(w) = f(u)$. Conversely, suppose that $f(w) = f(u)$. Then $f(w) \in \langle \theta_u \rangle_Y$, and thus $\Omega([w]) = \Omega([u])$. Since Ω is injective, it follows that $[w] = [u]$, that is, $w \sim u$.
2. The equivalence $w \sim_P f(w)$ stems from the construction of θ_w and $f(w)$. □

Proof of Lemma 2.15. For all $w \in \mathfrak{F}_X$, we will write $Y_w := \{x \in X : w \in \langle x \rangle_X\}$. Note that for all $w, u \in \mathfrak{C}$, we have $Y_w = Y_u \implies w = u$, as pointed out earlier. Given $Y \subseteq X$ and $c \subseteq \mathfrak{C}$, we define

$$\theta_Y^c := \left(\bigwedge_{y \in Y} y \right) \wedge \left(\bigwedge_{y \in Y} \neg y \right) \wedge \left(\bigwedge_{u \in c \text{ and } \varphi \in \mathcal{L}_X \text{ and } u \in \langle \varphi \rangle_X} \diamond \varphi \right) \wedge \square \left(\bigvee_{u \in c \text{ and } \varphi \in \mathcal{L}_X \text{ and } u \in \langle \varphi \rangle_X} \varphi \right).$$

1. Let w be a predecessor of c . Let $Y \subseteq X$ and assume that $Y \neq Y_u$ for all $u \in \mathfrak{C}$. Let V be a valuation on \mathfrak{F}_X such that $w \in V(x) \iff x \in Y$ holds for all $x \in X$, and V coincides with V_X everywhere else. Then $\mathfrak{F}_X, V, w \models \theta_Y^c$, and so $\{\theta_Y^c\}$ is consistent. Thus, there exists $w_Y \in \langle \theta_Y^c \rangle_X$. Then $w_Y \triangleright \mathfrak{C}$. If $w_Y \in \mathfrak{C}$ then by construction $Y \neq Y_{w_Y} = Y$, a contradiction. Hence $w_Y \triangleright \mathfrak{C}$. Conversely, if $u \triangleright \mathfrak{C}$, then $u \in \langle \theta_{Y_u}^c \rangle_X$ and so $u = w_{Y_u}$. Thus, the predecessors of \mathfrak{C} are exactly the w_Y 's, which can be checked to be pairwise distinct. We conclude that there are exactly $2^{|X|} - |\mathfrak{C}|$ predecessors of \mathfrak{C} .
2. Let $c \subset \mathfrak{C}$ and $w \triangleright c$. Given $Y \subseteq X$, we reason as above and obtain the existence of $w_Y \in \langle \theta_Y^c \rangle_X$. Then w_Y is initial, because there exists $u \in \mathfrak{C} \setminus c$, and thus some $\varphi \in \mathcal{L}$ such that $w_Y \in \langle \diamond \varphi \wedge \neg \varphi \rangle_X$. Therefore $w_Y \triangleright \mathfrak{C}$. Conversely, if $u \triangleright c$, then $u \in \langle \theta_{Y_u}^c \rangle_X$ and so $u = w_{Y_u}$. Hence, c has exactly $2^{|X|}$ predecessors. □

Proof of Proposition 2.16. 1. Suppose that there is a cluster \mathfrak{C} in \mathfrak{F}_X with $|\mathfrak{C}| = m'$. As pointed out earlier, the relation \sim_X coincides with the equality on \mathfrak{C} . So the elements of \mathfrak{C} correspond one-to-one with the equivalence classes of \sim_X , which can themselves be identified to subsets of X . Hence $m' = |\mathfrak{C}| \leq 2^{|X|}$. Further, \mathfrak{C} is isomorphic to the frame $\mathfrak{F}_{-1, m'}$, and so $\mathfrak{F}_{-1, m'} \models \mathbf{L}$. So for all $(n, m) \in \Delta$ we have $\mathfrak{F}_{-1, m'} \models \mathbf{Ax}_{n, m}$, whence $(n, m) \not\leq (-1, m')$. Conversely, suppose that $2^{|X|} \geq m'$ and $(-1, m') \notin \uparrow \Delta$. Then for all $(n, m) \in \Delta$ we have $(n, m) \not\leq (-1, m')$, whence $\mathfrak{F}_{-1, m'} \models \mathbf{L}$. Since $2^{|X|} \geq m'$, there exists an injection $\Omega : \mathfrak{F}_{-1, m'} \rightarrow 2^X$. For all $w \in \mathfrak{F}_{-1, m'}$, let us write $\theta_w := \bigwedge_{x \in \Omega(w)} x \wedge \bigwedge_{x \in X \setminus \Omega(w)} \neg x$. Given some arbitrary $w \in \mathfrak{F}_{-1, m'}$, we then introduce $\varphi := \theta_w \wedge (\bigwedge_{u \in \mathfrak{F}_{-1, m'}} \diamond \theta_u) \wedge \square (\bigvee_{u \in \mathfrak{F}_{-1, m'}} \theta_u)$ and $\Gamma_0 := \{\square \psi \rightarrow \psi : \psi \in \mathcal{L}_Z\}$. Clearly $\Gamma_0 \cup \{\varphi\}$ is satisfiable on $\mathfrak{F}_{-1, m'}$, and thus consistent with \mathbf{L} . Hence, there exists $w' \in \mathfrak{F}_X$ such that $\Gamma_0 \cup \{\varphi\} \subseteq w'$, and it is then clear that w' belongs to a cluster of size m .

2. Suppose that c is non-degenerated. Then c admits a predecessor u , and the subframe of \mathfrak{F}_X generated by u is isomorphic to $\mathfrak{F}_{n', m'}$. Then $\mathfrak{F}_{n', m'} \models \mathbf{L}$. Hence, for every $(n, m) \in \Delta$ we have $\mathfrak{F}_{n', m'} \models \mathbf{Ax}_{n, m}$, and thus $(n, m) \not\leq (n', m')$. Conversely, suppose that for all $(n, m) \in \Delta$ we have $(n, m) \not\leq (n', m')$. Then we have $\mathfrak{F}_{n', m'} \models \mathbf{L}$. We have $\mathfrak{F}_{n', m'} = (W, R)$ with $W = \{w\} \cup A \cup B$ defined as above, and $|B| = |\mathfrak{C} \setminus c| = n'$ and $|A| = |c| = m'$. Thus, there exists a bijection $f : \mathfrak{C} \rightarrow A \cup B$ such that $f[c] = A$ and

$f[\mathfrak{C} \setminus c] = B$. Now let V be a valuation on $\mathfrak{F}_{n',m'}$ defined by $V(x) := f[V_X(x) \cap \mathfrak{C}]$ for all $x \in X$. Since $\mathfrak{F}_{n',m'} \models \mathbf{L}$, the set $\Gamma := \{\varphi \in \mathcal{L}_X : \mathfrak{F}_{n',m'}, V, w \models \varphi\}$ is maximal consistent, and so belongs to W_X . Given $u \in c$ and $\varphi \in w$, we have $\mathfrak{F}_{n',m'}, V, f(u) \models \varphi$ by construction of V . Since $f(u) \in A$ we get $\mathfrak{F}_{n',m'}, V, w \models \diamond\varphi$, and thus $\diamond\varphi \in \Gamma$. Hence $\Gamma R_X u$. Conversely, consider $u \notin \mathfrak{C} \setminus c$. Then for all $v \in c$ there exists $\varphi_v \in v$ such that $\varphi_v \notin u$. By setting $\varphi := \bigvee_{v \in c} \varphi_v$, we obtain $\varphi \notin u$ and $c \subseteq \langle \varphi \rangle_X$. As a result we have $A = f[c] \subseteq \langle \varphi \rangle_X$, and thus $\mathfrak{F}_{n',m'}, V, w \models \Box\varphi$. Hence $\Box\varphi \in \Gamma$. Since $\varphi \notin u$, it follows that $\text{not } \Gamma R_X u$. This proves that $\Gamma \triangleright c$, as desired. \square

B Proofs for Section 3 (Reflections)

Lemma B.1. *Let $\sigma : \mathcal{L}_X \rightarrow \mathcal{L}_Y$ be a substitution.*

- *If σ is injective, then σ^* is surjective.*
- *If σ is surjective, then σ^* is injective.*

Proof. This is a direct consequence of [BdRV01, Prop. 5.52], together with the correspondence between substitutions and modal algebra homomorphisms discussed e.g. in [BG22, Sect. 3]. \square

Proof of Lemma 3.3. We write $Y^* := Y \setminus (Y_0 \cup Y_1)$. Since $|Z| = |Y^*|$, there exists $j : Y^* \rightarrow Z$ bijective, and since $Y^* \cap \mathcal{P} \subseteq Y \cap \mathcal{P} \subseteq Z$ we can assume that $j(p) = p$ for all $p \in Y^* \cap \mathcal{P}$. We then define a parametric substitution $\sigma : \mathcal{L}_Y \rightarrow \mathcal{L}_Z$ by

$$\sigma(y) := \begin{cases} \top & \text{if } y \in Y_1, \\ \perp & \text{if } y \in Y_0, \\ j(y) & \text{if } y \in Y^*. \end{cases}$$

That σ is parametric follows from the fact that $Y \cap \mathcal{P} = X \cap \mathcal{P} \subseteq Y^* \cap \mathcal{P}$. We also have a substitution $\tau : \mathcal{L}_Z \rightarrow \mathcal{L}_Y$ defined by $\tau(z) := j^{-1}(z)$. It is then clear that σ is surjective, τ is injective, and $\sigma\tau = \text{id}_{\mathcal{L}_Z}$. By Lemma B.1, it follows that $a := \sigma^* : \mathfrak{F}_Z \rightarrow \mathfrak{F}_Y$ is an injective parametric bounded morphism, $b := \tau^* : \mathfrak{F}_Y \rightarrow \mathfrak{F}_Z$ is a surjective bounded morphism, and we have $ba = \text{id}_{\mathfrak{F}_Z}$.

If \mathfrak{H} is trivial, then by Proposition 2.12, $\mathfrak{G} := a[\mathfrak{H}]$ is trivial as well, and b maps \mathfrak{G} to \mathfrak{H} . Otherwise, \mathfrak{H} is a pinned cluster $\hat{\mathfrak{C}}$, and then by Proposition 2.12, $\mathfrak{G} := a[\hat{\mathfrak{C}}]$ is a pinned cluster $\hat{\mathfrak{C}}'$. Then a induces a bijection from \mathfrak{C} to \mathfrak{C}' , and so b maps \mathfrak{C}' to \mathfrak{C} . Therefore, b maps $\hat{\mathfrak{C}}'$ to $\hat{\mathfrak{C}}$.

In both cases, we have constructed a closed primitive subframe \mathfrak{G} of \mathfrak{F}_Y such that a maps \mathfrak{H} to \mathfrak{G} , and b maps \mathfrak{G} to \mathfrak{H} . By construction of a , we also have $\mathfrak{G} \subseteq \langle y \rangle_Y$ for all $y \in Y_1$, and $\mathfrak{G} \subseteq \langle \neg y \rangle_Y$ for all $y \in Y_0$. We introduce the bounded morphism $g := hb|_{\mathfrak{G}} : \mathfrak{G} \rightarrow \mathfrak{F}$, and we claim that g is parametric. Indeed, if $p \in X \cap \mathcal{P}$ and $w \in \mathfrak{G}$, we have

$$g(w) \in \langle p \rangle_Y \iff h(\tau^*(w)) \in \langle p \rangle_Y \iff \tau^*(w) \in \langle p \rangle_Z \iff w \in \langle \tau(p) \rangle_Z = \langle j^{-1}(p) \rangle_Z = \langle p \rangle_Z$$

since h is parametric and $p \in X \cap \mathcal{P} \subseteq Y^* \cap \mathcal{P}$. In addition, $g[\mathfrak{G}] \subseteq h[\mathfrak{H}] \subseteq \mathfrak{F}$. We thus obtain a Y -reflection (\mathfrak{G}, g) of \mathfrak{F} satisfying $(\mathfrak{G}, g) \preceq (\mathfrak{H}, h)$. \square

Proof of Lemma 3.7. 1. See Proposition 2.16.

2. By 1, we have $|\mathfrak{C}| \leq 2^n$ and so \mathfrak{C} contains at most 2^{2^n} pre-clusters. By Lemma 2.15, every pre-cluster has at most 2^n predecessors. Therefore, we have

$$|\hat{\mathfrak{C}}| = 2^n + 2^{2^n} \cdot 2^n \leq 2^{2^n+n} + 2^{2^n+n} = 2^{2^n+n+1} \leq 2^{2^n+2^n} = 2^{2^{n+1}}.$$

3. Given w and u isolated in \mathfrak{F}_X , we have $w \sim_X u$ iff $w = u$. So there are at most 2^n isolated points.
4. As explained in the proof of Proposition 2.16, the clusters in \mathfrak{F}_X can be identified to sets of subsets of X . Hence, there are at most 2^{2^n} of them.
5. By the previous results, \mathfrak{F}_X contains at most 2^n isolated points of size 1, plus at most 2^{2^n} pinned clusters of size $\leq 2^{2^{n+1}}$. It follows that

$$|\mathfrak{F}_X| \leq 2^n + 2^{2^n} \cdot 2^{2^{n+1}} \leq 2^n + 2^{2^{n+1}+2^{n+1}} = 2^n + 2^{2^{n+2}} \leq 2 \cdot 2^{2^{n+2}} = 2^{2^{n+2}+1} \leq 2^{2^{n+3}}.$$

6. Let \mathfrak{F} be a generated subframe of \mathfrak{F}_X . A primitive Y -reflection of \mathfrak{F} is given by a primitive subframe \mathfrak{G} of \mathfrak{F}_Y and a map g from \mathfrak{G} to \mathfrak{F}_X . First, by 3 and 4, there are at most $2^m + 2^{2^m} \leq 2^{2^{m+1}}$ possible choices for \mathfrak{G} . Then, by 2 and 5, there are at most $(2^{2^{n+3}})^{2^{2^{m+1}}}$ possible choices for g . In total, the number of possibilities amounts to

$$\begin{aligned} 2^{2^{m+1}} \cdot (2^{2^{n+3}})^{2^{2^{m+1}}} &= 2^{2^{m+1}+2^{n+3} \cdot 2^{2^{m+1}}} \\ &\leq 2^{2^{m+1}+2^{2^{m+1}+2^{n+3}}} \leq 2^{2^{2^{m+1}+2^{n+3}}} \leq 2^{2^3 \cdot 2^{n+m+2}} \leq 2^{2^{n+m+4}}. \end{aligned}$$

□

C Proofs for Section 4 (Computing a most general unifier via filtering)

Proof of Lemma 4.4.

1. Given $w \in \mathfrak{F}_X$, we introduce $\rho_w := \left(\bigwedge_{x \in X \text{ and } w \in \langle x \rangle_Y} x \right) \wedge \left(\bigwedge_{x \in X \text{ and } w \notin \langle x \rangle_Y} \neg x \right)$ and then

$$\theta_w := \rho_w \wedge \left(\bigwedge_{w R_X u} \diamond \rho_u \right) \wedge \left(\bigwedge_{w R_X^2 u} \diamond \diamond \rho_u \right) \wedge \square \left(\bigvee_{w R_X u} \rho_u \right) \wedge \square \square \left(\bigvee_{w R_X^2 u} \rho_u \right).$$

By Proposition 2.10 it is then clear that θ_w has the desired property. In addition, the big conjunctions and disjunctions in θ_w quantify over the worlds reachable from w . In the worst case, w belongs to a pinned cluster, and so there are at most 2^n of them by Lemma 3.7. Since every ρ_u has size linear in n , it follows that $|\theta_w| = O(2^n \cdot n)$.

2. Let $\chi := \bigvee_{w \in \langle \psi \rangle_X} \theta_w$. We prove that $\langle \psi \rangle_X = \langle \chi \rangle_X$. Indeed, if $w \in \langle \chi \rangle_X$ then $w \in \langle \theta_u \rangle_X$ for some $u \in \langle \psi \rangle_X$. Then $w = u$ by construction of θ_u , and it follows that $w \in \langle \psi \rangle_X$. Conversely, if $w \in \langle \psi \rangle_X$, then $w \in \langle \theta_w \rangle_X$ and so $w \in \langle \chi \rangle_X$. Therefore $\psi \equiv_{\mathbf{L}} \chi$. By Lemma 3.7 we also have $|\langle \psi \rangle_X| \leq 2^{2^{n+3}}$, and thus $|\chi| = O(2^{2^{n+3}} \cdot 2^n \cdot n)$.

□