The Size of Interpolants in Modal Logics

Balder ten Cate
University of Amsterdam

Louwe Kuijer
University of Liverpool

Frank Wolter
University of Liverpool

Abstract

We start a systematic investigation of the size of Craig interpolants, uniform interpolants, and strongest implicates for (quasi-)normal modal logics. Our main upper bound states that for tabular modal logics, the computation of strongest implicates can be reduced in polynomial time to uniform interpolant computation in classical propositional logic. Hence they are of polynomial dag-size iff $\mathrm{NP} \subseteq \mathrm{P}_{/\mathrm{poly}}$. The reduction also holds for Craig interpolants and uniform interpolants if the tabular modal logic has the Craig interpolation property. Our main lower bound shows an unconditional exponential lower bound on the size of Craig interpolants and strongest implicates covering almost all non-tabular standard normal modal logics. For normal modal logics contained in or containing S4 or GL we obtain the following dichotomy: tabular logics have "propositionally sized" interpolants while for non-tabular logics an unconditional exponential lower bound holds.

2012 ACM Subject Classification Theory of computation → Modal and temporal logics

Keywords and phrases Modal logic, strongest implicates, uniform interpolants, Craig interpolants

1 Overview

Craig interpolation and the stronger property of uniform interpolation have been instrumental in understanding and applying modal logics. The Craig interpolation property (CIP) says that whenever an implication $\varphi \to \psi$ is valid there exists a formula χ using only the shared propositional variables of φ and ψ such that both $\varphi \to \chi$ and $\chi \to \psi$ are valid. The formula χ is then called a *Craig interpolant for* $\varphi \to \psi$ and often provides a useful explanation of its validity. Uniform interpolation (UIP) strengthens this. It says that for a signature σ of propositional variables, a Craig interpolant χ can be chosen uniformly for all ψ with $\mathrm{sig}(\psi) \cap \mathrm{sig}(\varphi) \subseteq \sigma$. Then χ is called a *uniform* σ*-interpolant of* φ and can be regarded as the result of eliminating existential quantifiers over, or 'forgetting', the variables $\mathrm{sig}(\varphi) \setminus \sigma$ in φ [28, 36, 23, 35]. If a logic does not have CIP, *strongest* σ*-implicates*, formulas χ that behave like uniform σ-interpolants but in which the condition '$\mathrm{sig}(\psi) \cap \mathrm{sig}(\varphi) \subseteq \sigma$' is replaced by '$\mathrm{sig}(\psi) \subseteq \sigma$' are often a sufficient approximation of uniform interpolants. They coincide with uniform interpolants if the logic has CIP.

While existence results and effective constructions for Craig interpolants, uniform interpolants, and strongest implicates in modal logics have been studied extensively, size bounds often remain unknown, even for standard modal logics. Moreover, in contrast to many other properties of modal logics such as the complexity of reasoning or the size of finite models, no general results have yet been established. In this paper we address the latter problem, and show first general results on the size of interpolants. Although we are mainly interested in normal modal logics, our results can often be more conveniently formulated for the even larger class of quasi-normal modal logics, modal logics determined by classes of (possibly generalised) Kripke frames with a distinguished world. We investigate modal logics with a single modal operator, the extension to polymodal logics is straightforward.

The lack of general results in modal logic is not surprising: even in the case of propositional

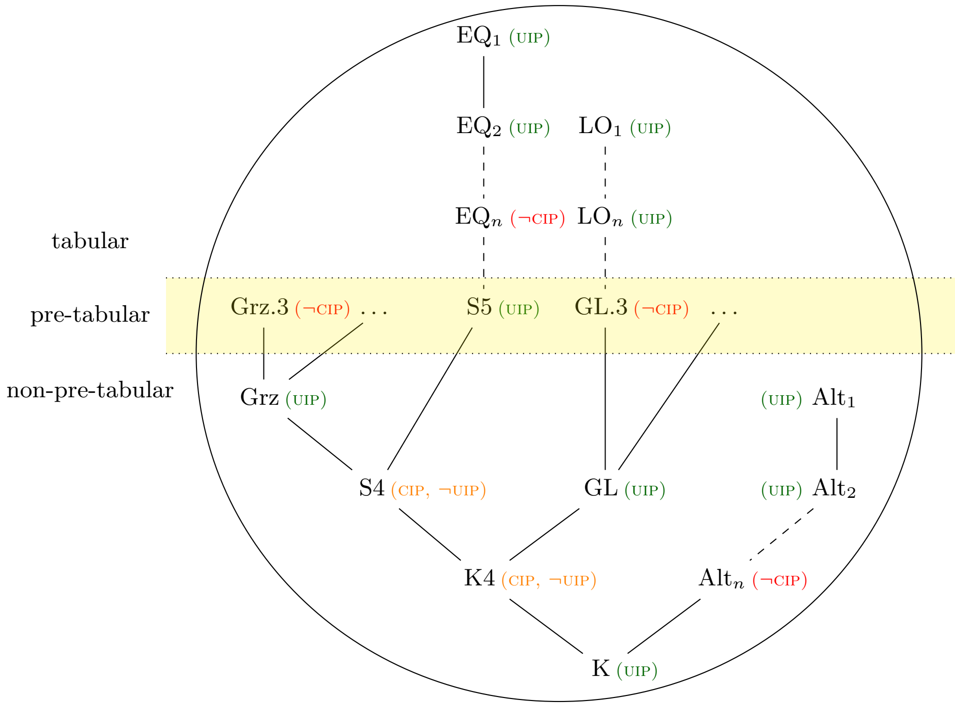

Figure 1 Landscape of non-tabular, pre-tabular, and tabular normal modal logics. To keep the figure simple, not all inclusion relations that hold between the given logics are indicated with edges. The failure of CIP for EQ_n and for Alt_n holds for $n > 2$. Logics are defined formally in Section 2 and pointers to the literature and proofs are supplied in Section 3 and Appendix F.

logic it is a major open problem whether Craig interpolants or uniform interpolants of polynomial size always exist. Although it is conjectured that this is false and both are of superpolynomial size in general, a proof would entail P $\neq$ NP and so is currently out of reach. As a result, the landscape of unconditional known lower bounds is rather sparse: currently, lower bounds have only been established for restricted fragments, most notably the monotone fragment of propositional logic, where one can show exponential lower bounds on the size of interpolants [29, 1]. Clearly, lower bounds obtained for propositional logic automatically transfer to modal logics. Thus, for example, the aforementioned exponential lower bound on the size of Craig interpolants in the monotone fragment of propositional logic immediately yields corresponding exponential lower bounds for monotone fragments of modal logics. An intriguing question, which we will pursue here, is when *upper bounds* for the case of propositional logic transfer to the modal setting.

Not every modal logic has the CIP, let alone the UIP. Indeed, while many standard modal logics have CIP, it is, for instance, known that only at most 37 normal modal logics containing S4 have CIP [10]. We refer the reader to Figure 1 for a small fragment of the landscape of modal logics, containing logics without CIP, logics with CIP and logics with UIP. Here and in what follows, a logic is called tabular if it is complete for a finite set of finite frames and pre-tabular if it is not tabular but all proper extensions are tabular. Note that a tabular logic with CIP also has UIP since one can obtain a uniform σ-interpolant for a formula φ by taking the conjunction over all relevant Craig interpolants (as there are only finitely many non-equivalent formulas over a finite signature of propositional variables).

We next discuss the main contributions of this paper. Unless stated otherwise, the size

$|\varphi|$ of a formula φ is the length of φ if represented as a dag (or equivalently, the number of subformulas of φ). See also Section 7 for further discussion regarding dag-size and tree-size.

First contribution (strongest implicates and Craig interpolants for tabular logics). For every tabular modal logic L, the computation of strongest $L(\sigma)$-implicates can be reduced in poly-time to the computation of uniform interpolants in propositional logic. As a corollary, we obtain that the size of strongest $L(\sigma)$-implicates is bounded polynomially in the size of propositional uniform interpolants.

▶ **Theorem 1.1.** *Let L be a tabular quasi-normal modal logic. We can associate to each signature σ of propositional variables a companion signature $\widehat{\sigma}$ such that there are poly-time translations*

$$tr_L : ML(\sigma) \to PL(\widehat{\sigma}) \quad \text{and} \quad rt_L : PL(\widehat{\sigma}) \to ML(\sigma)$$

such that for every modal formula φ and every uniform $\widehat{\tau}$-interpolant ψ of $tr_L(\varphi)$ in propositional logic, $rt_L(\psi)$ is a strongest $L(\tau)$-implicate for φ.

We recall that quasi-normal modal logics form a generalization of normal modal logic. Hence, the above results (and the results below) apply in particular to normal modal logics.

▶ **Corollary 1.2.** *Let L be a tabular quasi-normal modal logic. Let $f : \mathbb{N} \to \mathbb{N}$ be any function such that every propositional formula φ has a uniform σ-interpolant of size at most $f(|\varphi|)$ for all $\sigma \subseteq sig(\varphi)$. Then every modal formula φ has a strongest $L(\sigma)$-implicate (for $\sigma \subseteq sig(\varphi)$) of size at most $poly(f(poly(|\varphi|)))$.*

Similarly, for tabular modal logics L that have CIP, the computation of Craig interpolants in L can be reduced in poly-time to the computation of Craig interpolants in propositional logic. Consequently, the size of Craig interpolants in L is polynomial in the size of propositional Craig interpolants.

▶ **Theorem 1.3.** *Let L be a tabular quasi-normal modal logic that has CIP. Given a valid modal implication $\varphi \to \psi$ in L, and given a propositional Craig interpolant χ for $tr_L(\varphi) \to tr_L(\psi)$, it holds that $rt_L(\chi)$ is a Craig interpolant for $\varphi \to \psi$ with respect L.*

Here, tr_L and rt_L refer to the poly-time translation from Theorem 1.1.

▶ **Corollary 1.4.** *Let L be a tabular quasi-normal modal logic. Let $f : \mathbb{N} \to \mathbb{N}$ be any function such that every valid propositional implication $\varphi \to \psi$ has a Craig interpolant of size at most $f(|\varphi \to \psi|)$. Then every modal implication $\varphi \to \psi \in L$ has a Craig interpolant in L of size at most $poly(f(poly(|\varphi| + |\psi|)))$.*

The proof uses the standard bisimulation-based characterization of Craig interpolants and strongest implicates. A key ingredient of the proof is a poly-size representation of exponentially many non-bisimilar models using *abstract Φ-models*, for an appropriate poly-size set Φ of propositional formulas. Intuitively, an abstract Φ-model is a partial description of a Kripke model, which does not include a concrete propositional valuation for each world, but provides enough information to infer which worlds have the same propositional valuation.

Second contribution (exponential lower bounds for non-tabular logics) We prove a general exponential lower bound on the size of Craig interpolants and strongest implicates, which applies to many logics not covered by Theorem 1.1.

▶ **Theorem 1.5.** *Let L be any non-tabular normal modal logic that contains or is contained in S4 or Gödel-Löb provability logic GL, or, more generally, has a non-tabular intersection with S4 or GL.*

1. *There exist modal formulas $(\varphi_n)_{n=1,2,\ldots}$ of size polynomial in n and signatures $(\sigma_n)_{n=1,2,\ldots}$ with $\sigma_n \subseteq sig(\varphi_n)$, such that φ_n has a strongest $L(\sigma_n)$-implicate, and every strongest $L(\sigma_n)$-implicate of φ_n has size at least 2^n.*
2. *There are modal formulas $(\varphi_n)_{n=1,2,\ldots}$ and $(\psi_n)_{n=1,2,\ldots}$ of size polynomial in n such that $\varphi \to \psi \in L$ such that a Craig interpolant for $\varphi \to \psi$ exists, and every Craig interpolant for $\varphi \to \psi$ in L has size at least 2^n.*

The proof is based on a complete description of the pre-tabular normal modal logics containing S4 or GL and exploits their *poly-size model property*. In the case of strongest implicates, we actually prove a more general sufficient condition for an exponential lower bound that introduces the *exponential growth property* of a modal logic. It states that there are frames validating the logic in which one can reach exponentially worlds in a polynomial number of steps from a distinguished world. All pre-tabular normal modal logics contained in S4 or GL enjoy that property.

Interestingly, Theorem 1.1 and 1.3 and Theorem 1.5 together provide a dichotomy for normal modal logics L that are contained in or contain S4 or GL:

- If L is tabular, then L admits "propositional sized" strongest $L(\sigma)$-implicates and (if L has CIP) Craig interpolants;
- If L is non-tabular, there are exponential lower bounds on the size of strongest $L(\sigma)$-implicates and Craig interpolants.

Note that tabularity is decidable for normal modal logics containing S4 or GL, so this dichotomy is effective [6, Chapter 17].

We note that from Theorem 1.5 we obtain tight bounds for the size of Craig interpolants for many standard modal logics, including K, K4, S4, GL, Grzegorczyk's logic Grz, and S5.

▶ **Theorem 1.6.** *Let L be either K, K4, S4, Grz, GL, or S5. Then for every $\varphi \to \psi \in L$ there exists a Craig interpolant of at most exponential size. This bound is optimal.*

The exponential upper bounds for K and S5 are shown in [33, 12] (even for uniform interpolants) and they follow from [21, 34] for the remaining logics. The exponential lower bound for GL was independently shown in [27], where exponential lower bounds for K4 and S4 are mentioned as open research problems.

Finally, we refute the natural conjecture that the above dichotomy for S4 and GL generalizes to all normal modal logics by showing that the logic Alt_1 determined by frames in which every world has at most one successor admits a poly-time reduction of Craig and uniform interpolants to propositional Craig and, respectively, uniform interpolants.

The paper is structured as follows. In Section 2 we introduce the relevant notation and results in modal logic. Next, in Section 3, we discuss the main notions investigated in this paper and also review related work in propositional and modal logic. Section 4 is the core of this paper where we present the reduction of strongest implicates in tabular modal logics to uniform interpolants in propositional logic. The results and techniques presented in Section 4 are then also applied to give the reduction for Craig interpolants in tabular modal logics with CIP to propositional Craig interpolants in Section 5. Exponential lower bounds for interpolants are presented in Section 6. Missing definitions for normal and quasi-normal modal logics and detailed proofs are given in the appendix. In Section C, we also provide illustrating examples for the main reduction given in Section 4.

Log($\mathcal{F}$)	Class $\mathcal{F}$ of frames $\mathfrak{F} = (W, R)$
K	all frames
Alt$_n$	frames in which each node has at most n successors
S4	transitive and reflexive frames
Grz	frames validating S4 and no infinite strict chains $w_0 R w_1 R w_2 \cdots$
GL	transitive and no infinite chains $w_0 R w_1 R w_2 \cdots$
S5	equivalence relations
Grz.3/GL.3	GL/Grz frames satisfying $\forall w \forall v (wRv \vee vRw \vee w = v)$
LO$_n$	GL.3 frames of length at most n
EQ$_n$	S5 frames of size at most n

Table 1 List of relevant normal modal logics

2 Preliminaries

We start by introducing relevant notions from modal and propositional logic [6, 5]. Fix a countable infinite set *Prop* of *atoms* or *propositional variables*. A *signature* σ is a finite subset of *Prop*. The language PL of *propositional formulas* consists of all formulas constructed from atoms using the propositional connectives $\top$, $\neg$, and $\wedge$ in the usual way. Similarly, the language ML of *modal formulas* consists of formulas constructed from atoms using the propositional connectives $\top$, $\neg$, and $\wedge$ as well as the modal operator $\Box$. We set $\Diamond\varphi = \neg\Box\neg\varphi$, define $\Box^{\leq n}\varphi = \varphi \wedge \Box\varphi \wedge \cdots \wedge \Box^n\varphi$, and let $\Diamond^{\leq n}\varphi = \neg\Box^{\leq n}\neg\varphi$, for $n \geq 0$. By sig(φ) we denote the set of atoms in the (propositional or modal) formula φ. For a given signature σ, PL(σ) consists of those formulas $\varphi \in$ PL with sig(φ) $\subseteq \sigma$, and ML(σ) again consists of those formulas $\varphi \in$ ML with sig(φ) $\subseteq \sigma$.

ML(σ) is interpreted in *frames* $\mathfrak{F} = (W, R)$ with W a nonempty set of worlds and $R \subseteq W \times W$ an accessibility relation. A *pointed frame* is a pair $(\mathfrak{F}, w)$ with $w \in W$. A world $w \in W$ is called a *root* of $\mathfrak{F}$ if for every $v \in W$ there exists an R-path from w to v. A pointed frame $(\mathfrak{F}, w)$ with w a root of $\mathfrak{F}$ is called *rooted*. The *size* $|\mathfrak{F}|$ of a frame $\mathfrak{F}$ is the number of worlds in $\mathfrak{F}$.

A *model* based on $\mathfrak{F}$ is a pair $\mathfrak{M} = (\mathfrak{F}, V)$ with $V : W \to 2^{Prop}$ a *valuation*. The satisfaction-relation $\models$ between *pointed models* $(\mathfrak{M}, w)$ with $w \in W$ and modal formulas φ is defined as usual and denoted $\mathfrak{M}, w \models \varphi$. We write $\mathfrak{F}, w \models \varphi$ and call φ *valid in* $(\mathfrak{F}, w)$ if $\mathfrak{M}, w \models \varphi$ for all models $\mathfrak{M}$ based on $\mathfrak{F}$. For $w \in W$ denote by $\mathfrak{F}_w$ the restriction of $\mathfrak{F}$ to the set of worlds $v \in W$ such that there is an R-path from w to v. Then w is a root of $\mathfrak{F}_w$ and one can easily show that $\mathfrak{F}, w \models \varphi$ iff $\mathfrak{F}_w, w \models \varphi$.

Let $\mathcal{F}$ be a set of pointed frames. Then Log($\mathcal{F}$) $= \{\varphi \in \text{ML} \mid \text{ for all } (\mathfrak{F}, w) \in \mathcal{F} : \mathfrak{F}, w \models \varphi\}$ is the *logic determined by* $\mathcal{F}$. Log($\mathcal{F}$) is always a *quasi-normal modal logic*. Observe that Log($\mathcal{F}$) = Log($\{(\mathfrak{F}_w, w) \mid (\mathfrak{F}, w) \in \mathcal{F}\}$), and so we can always work with rooted frames. If L = Log($\mathcal{F}$) for some class $\mathcal{F}$ such that for every $(\mathfrak{F}, w) \in \mathcal{F}$ the set $\mathcal{F}$ also contains all $(\mathfrak{F}_v, v)$ with v in $\mathfrak{F}$, then Log($\mathcal{F}$) is a *normal modal logic*. For normal modal logics L we often drop the distinguished worlds when defining L as Log($\mathcal{F}$). Table 1 defines relevant normal modal logics:

We refer the reader to [6] for a detailed discussion of (quasi-)normal modal logics. A quasi-normal modal logic L is *tabular* if there exists a finite set $\mathcal{F}$ of finite pointed frames such that L = Log($\mathcal{F}$).

Let σ be a signature and let $\mathfrak{M}_1 = (\mathfrak{F}_1, V_1)$ and $\mathfrak{M}_2 = (\mathfrak{F}_2, V_2)$ with $\mathfrak{F}_i = (W_i, R_i)$ for $i =$

$1,2$ be models. $\mathfrak{M}_1, w_1$ and $\mathfrak{M}_2, w_2$ are *ML*(σ)*-indistinguishable*, in symbols $\mathfrak{M}_1, w_1 \equiv_{\mathrm{ML}(\sigma)} \mathfrak{M}_2, w_2$, if $\mathfrak{M}_1, w_1$ and $\mathfrak{M}_2, w_2$ satisfy the same formulas in ML(σ). ML(σ)-indistinguishability is characterized using σ-bisimilations. A *σ-bisimulation* between $\mathfrak{M}_1$ and $\mathfrak{M}_2$ is any relation $Z \subseteq W_1 \times W_2$ such that for all $w_1 \in W_1$ and $w_2 \in W_2$:

Atoms if $w_1 Z w_2$, then $\mathfrak{M}_1, w_1 \models p$ iff $\mathfrak{M}_2, w_2 \models p$ for all $p \in \sigma$,

Forth if $w_1 Z w_2$ and $(w_1, v_1) \in R_1$, then there is a $v_2 \in W_2$ such that $(w_2, v_2) \in R_2$ and $v_1 Z v_2$,

Back if $w_1 Z w_2$ and $(w_2, v_2) \in R_2$, then there is a $v_1 \in W_1$ such that $(w_1, v_1) \in R_1$ and $v_1 Z v_2$.

If Z is a σ-bisimulation between $\mathfrak{M}_1$ and $\mathfrak{M}_2$ such that $w_1 Z w_2$, we write $Z : \mathfrak{M}_1, w_1 \sim_\sigma \mathfrak{M}_2, w_2$. If there is any Z such that $Z : \mathfrak{M}_1, w_1 \sim_\sigma \mathfrak{M}_2, w_2$ we write $\mathfrak{M}_1, w_1 \sim_\sigma \mathfrak{M}_2, w_2$ and call $\mathfrak{M}_1, w_1$ and $\mathfrak{M}_2, w_2$ *σ-bisimilar*. It is known that $\sim_\sigma$ and $\equiv_\sigma$ coincide in finite pointed models [5].

3 Main Notions

In this paper, our main concern is the size and computation of strongest σ-implicates, uniform interpolants, and Craig interpolants. We next introduce these notions in detail in both modal and propositional logic.

Consider a quasi-normal logic L, a modal formula φ and a signature $\sigma \subseteq \mathrm{sig}(\varphi)$. A modal formula χ is called a *strongest $L(\sigma)$-implicate of* φ if $\mathrm{sig}(\chi) = \sigma$ and

- $\varphi \to \chi \in L$;
- if $\varphi \to \psi \in L$ and $\mathrm{sig}(\psi) \subseteq \sigma$, then $\chi \to \psi \in L$.

By definition, any two strongest $L(\sigma)$-implicates of φ are logically equivalent in L, so we speak about *the* strongest $L(\sigma)$-implicate if it exists. In this paper, we are interested in the computation and size of strongest $L(\sigma)$-implicates. There is a close relationship between strongest implicates and uniform interpolants. Recall that a formula χ is called a *Craig interpolant in* L for an implication $\varphi \to \psi \in L$ if $\varphi \to \chi \in L$, $\chi \to \psi \in L$, and $\mathrm{sig}(\chi) \subseteq \mathrm{sig}(\varphi) \cap \mathrm{sig}(\psi)$. A logic L has the *Craig interpolation property (CIP)* if for any $\varphi \to \psi \in L$ a Craig interpolant exists. A formula χ is called a *uniform $L(\sigma)$-interpolant* for φ if it satisfies the conditions of a strongest $L(\sigma)$-implicate of φ even if Condition 2 is strengthened by replacing the condition '$\mathrm{sig}(\psi) \subseteq \sigma$' by the condition $\mathrm{sig}(\psi) \cap \mathrm{sig}(\varphi) \subseteq \sigma$. It follows that a uniform $L(\sigma)$-interpolant for φ is a Craig interpolant for all $\varphi \to \psi \in L$ with $\mathrm{sig}(\psi) \cap \mathrm{sig}(\varphi) \subseteq \sigma$. The following is straightforward.

▶ **Lemma 3.1.** *If L has CIP, then strongest $L(\sigma)$-implicates and uniform $L(\sigma)$-interpolants coincide: if χ is a strongest $L(\sigma)$-implicate of φ, then it is a uniform $L(\sigma)$-interpolant for φ.*

▶ **Example 3.2.** Let $\mathfrak{F}_{1;3}$ be frame shown below, and $L_{1;3} = \mathrm{Log}(\{(\mathfrak{F}_{1;3}, 0)\})$.

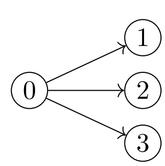

The logic $L_{1;3}$ does not have the CIP. For example, the strongest $L_{1;3}(\{p\})$-implicate of $\Diamond(p \wedge q) \wedge \Diamond(p \wedge \neg q)$ is $\Diamond p$. Yet $\Diamond p$ is not a uniform $L_{1;3}(\{p\})$-interpolant, since $(\Diamond(p \wedge q) \wedge \Diamond(p \wedge \neg q)) \to (\Diamond(\neg p \wedge r) \to \Box(\neg p \to r)) \in L_{1;3}$ while $\Diamond p \to (\Diamond(\neg p \wedge r) \to \Box(\neg p \to r)) \notin L_{1;3}$.

The following result characterizes strongest $L(\sigma)$-implicates using bisimulations: a strongest $L(\sigma)$-implicate of φ expresses existentially quantifying $\mathrm{sig}(\varphi) \setminus \sigma$ *modulo* a σ-bisimulation. Similar characterisations have been shown for many modal logics [28, 36, 7, 24].

▶ **Theorem 3.3.** *Let $L = Log(\mathcal{F}_L)$ with $\mathcal{F}_L$ a finite sets of finite pointed frames. Let φ be a modal formula and $\sigma \subseteq sig(\varphi)$. Then the following conditions are equivalent for all modal formulas χ with $sig(\chi) \subseteq \sigma$:*

1. *χ is the strongest $L(\sigma)$-implicate of φ;*
2. *for all models $(\mathfrak{M}, w)$ based on some $(\mathfrak{F}, w) \in \mathcal{F}_L$: $\mathfrak{M}, w \models \chi$ iff there exists a model $(\mathfrak{M}', w')$ based on some $(\mathfrak{F}', w') \in \mathcal{F}_L$ with $\mathfrak{M}', w' \sim_\sigma \mathfrak{M}, w$ and $\mathfrak{M}', w' \models \varphi$.*

The definitions of strongest σ-implicates, uniform interpolants, and Craig interpolants in propositional logic are derived from the definitions for modal logic in the obvious way. As propositional logic has CIP, strongest propositional σ-implicates are uniform propositional σ-interpolants.

Define the *dag-size* $|\varphi|$ of a propositional or modal formula φ as the number of its distinct subformulas. This corresponds to representing formulas as directed acyclic graphs and is the standard measure used in circuit complexity theory. If we take into account the number of different occurrences of subformulas, we obtain the *tree-size* $s(\varphi)$ of φ. Clearly, $s(\varphi) \geq |\varphi|$ and it is conjectured that in propositional logic there is a superpolynomial (even exponential) gap between dag-size and tree-size. However, this conjecture remains open [30].

It is also conjectured that both propositional Craig interpolants and uniform interpolants are (in the worst case) necessarily of exponential dag-size. In more detail, the following result shows how the size of propositional interpolants is related to open problems in circuit complexity [20], see also [26, 31].

▶ **Theorem 3.4.**

(1) Propositional logic has poly-dag-size uniform interpolants iff $\mathbf{NP} \subseteq \mathbf{P}_{/\mathbf{poly}}$.

(2) If propositional logic has poly-dag-size Craig interpolants, then $\mathbf{NP} \cap \mathbf{coNP} \subseteq \mathbf{P}_{/\mathbf{poly}}$.

We remind the reader that the question whether $\mathbf{NP} \subseteq \mathbf{P}_{/\mathbf{poly}}$ is a long-standing open problem. It is unlikely that it holds since, by the Karp-Lipton Theorem, if $\mathbf{NP} \subseteq \mathbf{P}_{/\mathbf{poly}}$, then the polynomial hierarchy collapses at the second level [19, 18, 2]. Conversely, if we could prove $\mathbf{NP} \not\subseteq \mathbf{P}_{/\mathbf{poly}}$, then $\mathbf{NP} \neq \mathbf{P}$ would follow since $\mathbf{P} \subseteq \mathbf{P}_{/\mathbf{poly}}$. It is also regarded as unlikely that $\mathbf{NP} \cap \mathbf{coNP} \subseteq \mathbf{P}_{/\mathbf{poly}}$. However, if we could prove $\mathbf{NP} \cap \mathbf{coNP} \not\subseteq \mathbf{P}_{/\mathbf{poly}}$, then again $\mathbf{NP} \neq \mathbf{P}$ would follow.

We next review relevant results on interpolants in modal logic. Among the main modal logics considered in this paper, K, S5, GL, and Grz have UIP [36] while S4 and K4 do not [13, 4]. S4 and K4 have CIP, however, while GL.3 and Grz.3 do not even enjoy CIP [11]. Regarding the size of Craig/uniform interpolants for the logics with CIP and UIP considered in this paper, for modal logics K and S5 it is known that uniform interpolants of at most exponential size always exist [33, 12]. In contrast, for GL and Grz, to date, only non-elementary methods are known for constructing uniform interpolants. Interestingly, however, there are tools available for computing uniform interpolants in these logics [8]. For Craig interpolants, for K and S5, the exponential upper bound is inherited from the exponential upper bound for uniform interpolants. For K4, S4, Grz, and GL one can obtain exponential upper bounds by using the reductions to Craig interpolants in K given in [21, Section 3.7] and applying the tableau developed in [34]. While this is done explicitly only for K4, it is straightforward to extend the proof to S4, Grz, and GL.

It is also worth noting that it is undecidable whether a finitely axiomatized normal modal logic has CIP [21, Section 9.6]. Using the same technique, one can show the undecidability of UIP. This does not necessarily imply that we cannot show effective dichotomies for the size of interpolants since even for finitely axiomatized normal modal logics containing GL the CIP is undecidable [6, Theorem 17.19], but we have nevertheless shown a (conditional) dichotomy result for size.

4 Constructing Strongest Implicates via Propositional Logic

We show how one can construct, for tabular quasi-normal logics, in polynomial time strongest implicates from uniform interpolants in propositional logic.

4.1 Abstraction

The number of different valuations V of atoms in a signature σ in a fixed finite frame is exponential in $|\sigma|$, even for frames with a single world. To avoid working with all valuations, in this subsection we introduce abstract models, of which we only require polynomially many to supply sufficient information about the underpinning models for our purposes. Unless stated otherwise, in this section, we only consider finite rooted pointed frames $(\mathfrak{F}, w)$. We typically denote the root w by 0.

Given a signature σ, we call a finite set $\Phi \subseteq \mathrm{PL}(\sigma)$ an *exhaustive and mutually exclusive set over* σ, σ-EME for short, if (i) $\models \bigvee_{\varphi\in\Phi}\varphi$ and (ii) $\models \neg(\varphi \wedge \psi)$ for all $\varphi \neq \psi$ with $\varphi, \psi \in \Phi$. If $\mathfrak{M} = (\mathfrak{F}, V)$ is a model and Φ a σ-EME, we say that Φ is a σ*-cover* of $\mathfrak{M}$ if for every $\varphi \in \Phi$ and every $w_1, w_2 \in W$ if $\mathfrak{M}, w_1 \models \varphi$ and $\mathfrak{M}, w_2 \models \varphi$ then $\mathfrak{M}, w_1 \equiv_{\mathrm{PL}(\sigma)} \mathfrak{M}, w_2$.

If Φ is a σ-EME, then every world of $\mathfrak{M}$ satisfies exactly one formula $\varphi \in \Phi$. So Φ is a σ-cover of $\mathfrak{M}$ if, and only if, any two worlds of $\mathfrak{M}$ that are distinguishable in $\mathrm{PL}(\sigma)$ are distinguished by Φ.

▶ **Lemma 4.1.** *For every σ-EME Φ and for every rooted model $\mathfrak{M}, 0$ of size at most N, we have Φ is a σ-cover of $\mathfrak{M}$ if, and only if, $\mathfrak{M}, 0 \models \mathit{cover}(\sigma, \Phi)$ where*

$$\mathit{cover}(\sigma, \Phi) = \bigwedge_{\varphi\in\Phi}\bigwedge_{p\in\sigma}\left(\Diamond^{\leq N}(\varphi \wedge p) \to \Box^{\leq N}(\varphi \to p)\right)$$

The proof is straightforward, so we omit it here.

A key observation that we will make use of is that, even though the number of possible σ-valuations over a given frame is exponential in $|\sigma|$, the number of σ-EMEs needed in order to cover all of them is only polynomial in $|\sigma|$.

▶ **Lemma 4.2.** *Fix a finite frame $\mathfrak{F} = (W, R)$. For every signature σ, we can compute in polynomial time a set $\Gamma_{\mathfrak{F}}$ of σ-EMEs such that every model based on $\mathfrak{F}$ has a σ-cover in $\Gamma_{\mathfrak{F}}$. Moreover, each $\Phi \in \Gamma_{\mathfrak{F}}$ consists of at most $|W|$ formulas, each of size $O(|W|)$.*

The proof (given in the appendix) is along the following lines: in any model based on $\mathfrak{F}$, there are at most $|W|$ distinguishable worlds. Using a partition refinement procedure we can find, for each indistinguishability class, an identifying formula that is a conjunction of literals over σ of size at most $O(|W|)$. Hence every model has a σ-cover Φ such that $|\Phi| \leq |W|$ and $|\varphi| \leq O(|W|)$ for each $\varphi \in \Phi$. Because the bounds on $|\Phi|$ and $|\varphi|$ do not depend on σ, each such cover contains only a constant number of atoms. There are polynomially many ways to choose a constant number of atoms from σ, so $\Gamma_{\mathfrak{F}}$ can be computed in polynomial time (and $|\Gamma_{\mathfrak{F}}|$ is polynomial).

▶ **Example 4.3.** Assume $\sigma = \{p, q\}$. If $\mathfrak{F}$ has a single world, then we can take $\Gamma_{\mathfrak{F}} = \{\{\top\}\}$ and if $\mathfrak{F}$ has two worlds, $\Gamma_{\mathfrak{F}} = \{\{p, \neg p\}, \{q, \neg q\}, \{\top\}\}$ is as required.

If Φ is a σ-EME, then an *abstract Φ-model* is a pair $\mathfrak{A} = (\mathfrak{F}, f)$, where $\mathfrak{F} = (W, R)$ is a frame and $f : W \to \Phi$ is a function. One can regard abstract Φ-models as partially specified models: rather than stating for every world w and atom p whether p is true at w, they only state which element of Φ is true in w. If Φ is a σ-cover of $\mathfrak{M}$, then the *Φ-abstraction of* $\mathfrak{M}$ is the unique abstract Φ-model $\mathfrak{A} = (\mathfrak{F}, f)$ such that $\mathfrak{M}, w \models f(w)$ for all $w \in W$.

We also lift bisimulations from the level of models to the level of abstract Φ-models. Let Φ be a σ-EME, and $\mathfrak{A}_1 = (\mathfrak{F}_1, f_1)$ and $\mathfrak{A}_2 = (\mathfrak{F}_2, f_2)$ abstract Φ-models. A relation $Z \subseteq W_1 \times W_2$ is an *abstract Φ-bisimulation* between $\mathfrak{A}_1$ and $\mathfrak{A}_2$ if conditions **Forth** and **Back** introduced in the definition of σ-bisimulations hold and

Abstract atoms for all $w_1 \in W_1$ and $w_2 \in W_2$, if $w_1 Z w_2$, then $f_1(w_1) = f_2(w_2)$.

We write $Z : \mathfrak{A}_1, w_1 \sim_\Phi \mathfrak{A}_2, w_2$ if Z is an abstract Φ-bisimulation such that $w_1 Z w_2$. If there is a Z such that $Z : \mathfrak{A}_1, w_1 \sim_\Phi \mathfrak{A}_2, w_2$ we write $\mathfrak{A}_1, w_1 \sim_\Phi \mathfrak{A}_2, w_2$ and say that $\mathfrak{A}_1, w_1$ and $\mathfrak{A}_2, w_2$ are *abstractly Φ-bisimilar*.

▶ **Proposition 4.4.** *Let $\mathfrak{M}, 0$ be a rooted model, Φ a σ-cover of $\mathfrak{M}$ and $\mathfrak{A}$ the Φ-abstraction of $\mathfrak{M}$. Then for every rooted frame $(\mathfrak{F}', 0)$ with $\mathfrak{F}' = (W', R')$ and for every $Z \subseteq W \times W'$ the following are equivalent:*
1. *$Z : (\mathfrak{M}, 0) \sim_\sigma (\mathfrak{M}', 0)$ for some model $\mathfrak{M}'$ based on $\mathfrak{F}'$,*
2. *$Z : (\mathfrak{A}, 0) \sim_\Phi (\mathfrak{A}', 0)$ for some abstract Φ-model $\mathfrak{A}'$ based on $\mathfrak{F}'$.*

Furthermore, if such $\mathfrak{A}'$ and $\mathfrak{M}'$ exist then $\mathfrak{A}'$ is the Φ-abstraction of $\mathfrak{M}'$.

For a σ-EME Φ, we often treat the propositional formulas in Φ as if they are atoms. Specifically, we will denote by $\mathrm{ML}(\Phi)$ the set of modal formulas constructed from the formulas in Φ (treated as atoms) in the usual way. In abstract Φ-models $\mathfrak{A}$, formulas $\psi \in \mathrm{ML}(\Phi)$ are evaluated in the expected way by setting $\mathfrak{A}, w \models \varphi$ iff $f(w) = \varphi$ for all $\varphi \in \Phi$ and then using the standard truth conditions for modal formulas. Observe that if $\psi \in \mathrm{ML}(\Phi)$ and $\mathfrak{A}$ is the Φ-abstraction of $\mathfrak{M}$, then $\mathfrak{M}, w \models \psi$ if, and only if, $\mathfrak{A}, w \models \psi$.

▶ **Proposition 4.5.** *Fix a natural number $N > 0$. Given a σ-EME Φ (with σ a signature) and a rooted abstract Φ-model $(\mathfrak{A}, 0)$ of size at most N, we can compute in polynomial time a formula $\delta_\mathfrak{A} \in ML(\Phi)$ such that, for all rooted abstract Φ-models $(\mathfrak{A}', 0)$ of size at most N, $\mathfrak{A}', 0 \models \delta_\mathfrak{A}$ iff $\mathfrak{A}', 0 \sim_\Phi \mathfrak{A}, 0$. Moreover, $|\delta_\mathfrak{A}| \leq O(2^N \cdot \max_{\varphi \in \Phi} |\varphi|)$.*

The proof (given in the appendix) is along the following lines: let $\mathfrak{A} = (\mathfrak{F}, f)$ be given. For each $w \in W$, let $[w] = \{w' \mid \mathfrak{A}, w \sim_\Phi \mathfrak{A}, w'\}$. We first use a partition refinement procedure to compute, in polynomial time, a formula $\varphi_{[w]}$ such that $\mathfrak{A}, w' \models \varphi_w$ iff $w' \in [w]$. We define

$$\delta_\mathfrak{A} = \varphi_{[0]} \wedge \Big(\Box^{\leq N} \bigvee_{w \in W} \varphi_{[w]}\Big) \wedge \Big(\bigwedge_{w \in W} \bigwedge_{w' \in W} \Box^{\leq N} (\varphi_{[w]} \to \pm\Diamond\varphi_{[w']})\Big)$$

where, $\varphi_{[w]} \to \pm\Diamond\varphi_{[w']}$ is equal to $\varphi_{[w]} \to \Diamond\varphi_{[w']}$ if there are $j \in [w]$ and $j' \in [w']$ such that $(j, j') \in R$, and equal to $\varphi_{[w]} \to \neg\Diamond\varphi_{[w']}$ otherwise. We then show that $\delta_\mathfrak{A}$ has the desired properties. See the proof in the appendix for more details.

Fix now a tabular quasi-normal modal logic L. Since L is tabular, there is a finite set of rooted finite frames $\mathcal{F}_L = \{(\mathfrak{F}_1, 0), \ldots, (\mathfrak{F}_n, 0)\}$ with $L = \mathrm{Log}(\mathcal{F}_L)$. Fix such a set. Note that we assume without loss of generality that the root of each $\mathfrak{F}_i$ is 0. We will often be sloppy and write $\mathfrak{F} \in \mathcal{F}$ instead of $(\mathfrak{F}, 0) \in \mathcal{F}$. Let $N = \max\{|\mathfrak{F}| \mid \mathfrak{F} \in \mathcal{F}_L\}$. By an *abstract Φ-model for L* we will mean an abstract Φ-model whose frame belongs to $\mathcal{F}_L$. If $\mathfrak{A}$ is an abstract Φ-model for L, its *Φ-bisimulation class* with respect to L, denoted $[\mathfrak{A}]_L$, is the set of abstract Φ-models $\mathfrak{A}'$ for L such that $\mathfrak{A}, 0 \sim_\Phi \mathfrak{A}', 0$. An *abstract class identifier* for a Φ-bisimulation class $[\mathfrak{A}]_L$ is a formula $\delta \in \mathrm{ML}(\Phi)$ such that for every abstract Φ-model $\mathfrak{A}'$ for L, we have $\mathfrak{A}', 0 \models \delta$ if, and only if, $\mathfrak{A}' \in [\mathfrak{A}]_L$. If δ is an abstract class identifier for $[\mathfrak{A}]_L$ we also write $[\delta]_L$ for $[\mathfrak{A}]_L$.

Note that the formula $\delta_{\mathfrak{A}}$ from Proposition 4.5 is an abstract class identifier for $[\mathfrak{A}]$.

▶ **Definition 4.6.** *A* σ-encoding *for L is a pair $(\Gamma, \{\Delta_\Phi\}_{\Phi\in\Gamma})$ where Γ is a set of σ-EMEs and Δ_Φ is a set of Φ-abstract class identifiers for L such that for every model $\mathfrak{M}$ based on a frame in $\mathcal{F}_L$, there are a $\Phi \in \Gamma$ and $\delta \in \Delta_\Phi$ such that Φ is a σ-cover of $\mathfrak{M}$ and $\mathfrak{M}, 0 \models \delta$.*

The following proposition summarizes the relevant results proved above in a more compact way:

▶ **Proposition 4.7.** *Fix a tabular quasi-normal modal logic L. Given a signature σ, we can compute in polynomial time a σ-encoding for L.*

Proof. We first use Lemma 4.2 to construct, for each $\mathfrak{F} \in \mathcal{F}_L$, in time polynomial in $|\sigma|$, a set $\Gamma_{\mathfrak{F}}$ of σ-EMEs such that every model based on $\mathfrak{F}$ has a σ-cover in $\Gamma_{\mathfrak{F}}$. Let $\Gamma = \bigcup_{\mathfrak{F}\in\mathcal{F}_L} \Gamma_{\mathfrak{F}}$. Next, we apply Proposition 4.5 to obtain, for each $\Phi \in \Gamma$ and for each abstract Φ-model $\mathfrak{A}$ for L, an abstract class identifier $\delta_{\mathfrak{A}}$. Note that there are at most $\Sigma_{\mathfrak{F}\in\mathcal{F}_L} |\Phi|^{|\mathfrak{F}|}$ many abstract Φ-models for L, which is polynomial in $|\Phi|$ and hence in $|\sigma|$. Let Δ be the set of all these abstract class identifiers. It follows from the construction that (Γ, Δ) is a σ-encoding for L. ◀

4.2 From Modal Logic to Propositional Logic and Back

In this section, we use propositional formulas to speak about pointed models. Given $\mathfrak{F} = (W, R)$ with root 0 and σ, we form the set of atoms $\sigma_{\mathfrak{F}} = \sigma \times W$. For notational convenience we denote $(p, w) \in \sigma_{\mathfrak{F}}$ as p_w. Then $\mathfrak{M} = (\mathfrak{F}, V)$ defines a propositional valuation $v_{\mathfrak{M}} : \sigma \times W \to \{0, 1\}$ for the language $\mathrm{PL}(\sigma_{\mathfrak{F}})$ by setting $v_{\mathfrak{M}}(p_w) = 1$ if $\mathfrak{M}, w \models p$ and $v_{\mathfrak{M}}(p_w) = 0$ otherwise. Observe that, conversely, every propositional valuation v for $\mathrm{PL}(\sigma_{\mathfrak{F}})$ defines a modal σ-model $\mathfrak{M}_v = (\mathfrak{F}, V_v)$ by setting $\mathfrak{M}_v, w \models p$ iff $v(p_w) = 1$. For every modal formula in $\mathrm{ML}(\sigma)$ we can easily find an 'equivalent' propositional formula in $\mathrm{PL}(\sigma_{\mathfrak{F}})$ which we define next.

▶ **Definition 4.8.** *The translation functions $tr_{\mathfrak{F},w} : ML(\sigma) \to PL(\sigma_{\mathfrak{F}})$ with $w \in W$ are defined by mutual induction:*

$$\begin{aligned} tr_{\mathfrak{F},w}(p) &= p_w \\ tr_{\mathfrak{F},w}(\neg\varphi) &= \neg tr_{\mathfrak{F},w}(\varphi) \\ tr_{\mathfrak{F},w}(\varphi \wedge \psi) &= tr_{\mathfrak{F},w}(\varphi) \wedge tr_{\mathfrak{F},w}(\psi) \\ tr_{\mathfrak{F},w}(\Diamond\varphi) &= \bigvee_{(w,w')\in R} tr_{\mathfrak{F},w'}(\varphi) \end{aligned}$$

We summarize the main properties of $tr_{\mathfrak{F},w}(\varphi)$.

▶ **Proposition 4.9.** *Let $\mathfrak{F} = (W, R)$ be a finite frame with root 0, $w \in W$, and $\mathfrak{M}$ a model based on $\mathfrak{F}$.*

1. *for every $\varphi \in ML(\sigma)$, we have $\mathfrak{M}, w \models \varphi$ iff $v_{\mathfrak{M}} \models tr_{\mathfrak{F},w}(\varphi)$.*
2. *for every $\varphi \in ML(\sigma)$, $(\mathfrak{F}, 0) \models \varphi$ iff $tr_{\mathfrak{F},0}(\varphi)$ is a tautology.*

Note that not every propositional formula over $\sigma \times W$ is equivalent to the translation of a modal formula.

▶ **Example 4.10.** Consider the following pointed models $(\mathfrak{M}_1, 0)$ and $(\mathfrak{M}_2, 0)$:

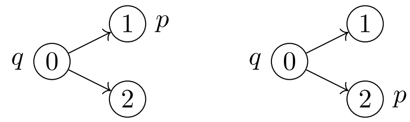

Since these models are σ-bisimilar for $\sigma = \{p, q\}$, no formula in $\mathrm{ML}(\sigma)$ can distinguish between them. It follows that no translation of a formula in $\mathrm{ML}(\sigma)$ can distinguish between the corresponding propositional models $v_{\mathfrak{M}_1}$ and $v_{\mathfrak{M}_2}$. This implies that there is no translation of a formula in $\mathrm{ML}(\sigma)$ that is equivalent to the propositional formula p_1.

This concludes the presentation of the "forward" translation tr, i.e., from modal formulas to propositional formulas. We also need a reverse translation rt, in order to translate the propositional uniform interpolant back to a modal formula. For this, we make use of the techniques we developed in the previous subsection. In particular, we make use of Proposition 4.7 and Lemma 4.1.

Fix a tabular quasi-normal modal logic L, and let $\mathcal{F}_L$ be a finite set of finite rooted frames such that $L = \mathrm{Log}(\mathcal{F}_L)$. Let $N = \max\{|\mathfrak{F}'| \mid \mathfrak{F}' \in \mathcal{F}_L\}$. Recall that we denote the root of each $\mathfrak{F} \in \mathcal{F}_L$ by 0.

▶ **Definition 4.11.** *Fix a tabular quasi-normal modal logic L, and let $\mathfrak{F} \in \mathcal{F}_L$. Then $rt_{\mathfrak{F},L} : PL(\sigma_{\mathfrak{F}}) \to ML(\sigma)$ is given by*

$$rt_{\mathfrak{F},L}(\xi) = \bigwedge_{\Phi\in\Gamma} \bigwedge_{\delta\in\Delta_\Phi} \left((cover(\sigma, \Phi) \wedge \delta) \to \bigvee_{(\mathfrak{F},f)\in[\delta]_L} \xi^f \right)$$

where $(\Gamma, \{\Delta_\Phi\}_{\Phi\in\Gamma})$ is a σ-encoding for L as given by Proposition 4.7, and ξ^f is the result of replacing, in ξ, each indexed atom $p_w \in \sigma_{\mathfrak{F}}$ by $\Diamond^{\leq N}(f(w) \wedge p)$.

By the results in the previous subsection, all computations involved in this translation can be performed in polynomial time (assuming L is fixed, and hence $N = \max_{\mathfrak{F}\in\mathcal{F}_L} |\mathfrak{F}|$ can be treated as a constant). In particular, note that the number of $(\mathfrak{F}, f) \in [\delta]_L$ is at most $|\Phi|^{|\mathfrak{F}|} \leq |\Phi|^N$, hence is also polynomial.

▶ **Proposition 4.12.** *Fix a tabular quasi-normal modal logic L. Let $\mathfrak{M}, 0$ be any pointed model based on a pointed frame in $\mathcal{F}_L$. Furthermore, let $\mathfrak{F} \in \mathcal{F}_L$ and let $\xi \in PL(\sigma_{\mathfrak{F}})$. Then the following are equivalent:*

1. $\mathfrak{M}, 0 \models rt_{\mathfrak{F},L}(\xi)$,
2. *there is a model $\mathfrak{M}'$ based on the frame $\mathfrak{F}$ such that $\mathfrak{M}, 0 \sim_\sigma \mathfrak{M}', 0$ and $v_{\mathfrak{M}'} \models \xi$.*

The proof (given in the appendix) is along the following lines. Let $\Phi \in \Gamma$ be a σ-cover of $\mathfrak{M}$, and $\mathfrak{A} = (\mathfrak{F}, f)$ the Φ-abstraction of $\mathfrak{M}$. Furthermore, let $\delta \in \Delta_\Phi$ be such that $\mathfrak{M}, 0 \models cover(\sigma, \Phi) \wedge \delta$. Then $\mathfrak{A} \in [\delta]_L$.

It follows that the models bisimilar to $\mathfrak{M}$ are exactly those $\mathfrak{M}'$ that have a Φ-abstraction $\mathfrak{A}' = (\mathfrak{F}, f) \in [\delta]_L$. For any such $\mathfrak{M}'$ and any $\varphi \in \Phi$, if $\mathfrak{M}', w \models \varphi$ then $v_{\mathfrak{M}'} \models p_w$ if, and only if, $\mathfrak{M}, 0 \models \Diamond^{\leq N}(f(w) \wedge p)$. This implies that $\mathfrak{M}, 0 \models \xi^f$ if, and only if, $v_{\mathfrak{M}'} \models \xi$.

That, in turn, implies that $\mathfrak{M}, 0 \models \bigvee_{(\mathfrak{F},f)\in[\delta]_L} \xi^f$ if, and only if, there is some $\mathfrak{M}', 0$ bisimilar to $\mathfrak{M}, 0$ such that $v_{\mathfrak{M}'} \models \xi$. This is true for any Φ, δ such that $\mathfrak{M}, 0 \models cover(\sigma, \Phi) \wedge \delta$, and so the proposition follows.

The following result implies Theorem 1.1 for $L = \mathrm{Log}(\mathcal{F})$ for a singleton $\mathcal{F} = \{(\mathfrak{F}, 0)\}$.

▶ **Theorem 4.13.** *Fix a tabular quasi-normal modal logic L with $L = Log(\mathcal{F})$ for a singleton $\mathcal{F} = \{(\mathfrak{F}, 0)\}$. Let $\varphi \in ML$ and let $\sigma \subseteq sig(\varphi)$. If $\xi \in PL(\sigma_{\mathfrak{F}})$ is a propositional uniform $\sigma_{\mathfrak{F}}$-interpolant for $tr_{\mathfrak{F},0}(\varphi)$, then $rt_{\mathfrak{F},L}(\xi)$ is a strongest $L(\sigma)$-implicate of φ.*

Proof. (1) We show that $\varphi \to rt_{\mathfrak{F},L}(\xi) \in L$. Assume $\mathfrak{M}, 0 \models \varphi$ with $\mathfrak{M}$ based on $\mathfrak{F}$. Then $v_{\mathfrak{M}} \models tr_{\mathfrak{F},0}(\varphi)$. Hence $v_{\mathfrak{M}} \models \xi$. By the backward direction of Proposition 4.12 (using the identity bisimulation), $\mathfrak{M}, 0 \models rt_{\mathfrak{F},L}(\xi)$.

(2) Suppose that $\varphi \to \chi \in L$ for some $\chi \in \mathrm{ML}(\sigma)$. We must show that $rt_{\mathfrak{F},L}(\xi) \to \chi \in L$. Assume $\mathfrak{M}, 0 \models rt_{\mathfrak{F},L}(\xi)$. By Proposition 4.12, there is a σ-bisimilar $\mathfrak{M}', 0$ based on $\mathfrak{F}$ with $v_{\mathfrak{M}'} \models \xi$. Next note that by Proposition 4.9, $tr_{\mathfrak{F},0}(\varphi) \to tr_{\mathfrak{F},0}(\chi)$ is a tautology. Hence, as ξ is a uniform $\sigma_{\mathfrak{F}}$-interpolant for $tr_{\mathfrak{F},0}(\varphi)$, $\xi \to tr_{\mathfrak{F},0}(\chi)$ is a tautology. It follows that $v_{\mathfrak{M}'} \models tr_{\mathfrak{F},0}(\chi)$ and so $\mathfrak{M}', 0 \models \chi$. Therefore, by invariance under σ-bisimulations, $\mathfrak{M}, 0 \models \chi$. ◀

4.3 From Single Frames to Multiple Frames

We now lift Theorem 4.13 to tabular quasi-normal logics defined by a *finite set* of finite rooted frames. Let as before $\mathcal{F}$ be a finite set of finite pointed frames and $L = \mathrm{Log}(\mathcal{F})$. We first show that the strongest $L(\sigma)$-implicate of a modal formula φ can be obtained in polynomial time from the propositional uniform $\sigma_{\mathfrak{F}}$-interpolants of $tr_{\mathfrak{F}}(\varphi)$, $\mathfrak{F} \in \mathcal{F}$. To strengthen this result to a reduction of L-implicate computation to a *single* propositional uniform interpolant, we then generalize the translation $tr_{\mathfrak{F}}$ to sets $\mathcal{F}$ of frames and show how uniform interpolants for $tr_{\mathcal{F}}(\varphi)$ are obtained from uniform interpolants for $tr_{\mathfrak{F}}(\varphi)$, $\mathfrak{F} \in \mathcal{F}$.

▶ **Theorem 4.14.** *Fix a tabular quasi-normal modal logic L with $L = Log(\mathcal{F})$. Let $\varphi \in ML$ and $\sigma \subseteq sig(\varphi)$. Let $\xi_{\mathfrak{F}} \in PL(\sigma_{\mathfrak{F}})$ be propositional uniform $\sigma_{\mathfrak{F}}$-interpolants for $tr_{\mathfrak{F}}(\varphi)$, for $\mathfrak{F} \in \mathcal{F}$. Then $\bigvee_{\mathfrak{F} \in \mathcal{F}} rt_{\mathfrak{F},L}(\xi_{\mathfrak{F}})$ (cf. Proposition 4.12) is a strongest $L(\sigma)$-implicate of φ.*

Proof. (1) We show that $\varphi \to \bigvee_{\mathfrak{F} \in \mathcal{F}} rt_{\mathfrak{F},L}(\xi_{\mathfrak{F}}) \in L$. Assume $\mathfrak{M}, 0 \models \varphi$ with $\mathfrak{M}$ based on $\mathfrak{F} \in \mathcal{F}$. Then $v_{\mathfrak{M}} \models tr_{\mathfrak{F}}(\varphi)$. Hence $v_{\mathfrak{M}} \models \xi_{\mathfrak{F}}$. By Proposition 4.12, $\mathfrak{M}, 0 \models rt_{\mathfrak{F},L}(\xi_{\mathfrak{F}})$, as required.

(2) We show that $\bigvee_{\mathfrak{F} \in \mathcal{F}} rt_{\mathfrak{F},L}(\xi_{\mathfrak{F}}) \to \chi \in L$ for all $\mathrm{ML}(\sigma)$-formulas χ with $\varphi \to \chi \in L$. To this end, it suffices to show that if $\mathfrak{M}, 0 \models \bigvee_{\mathfrak{F} \in \mathcal{F}} rt_{\mathfrak{F},L}(\xi_{\mathfrak{F}})$ for some $\mathfrak{M}$ based on $\mathfrak{F} \in \mathcal{F}$, then there exists a σ-bisimilar $\mathfrak{M}'$ with $\mathfrak{M}', 0 \models \varphi$ based on some $\mathfrak{F}' \in \mathcal{F}$. So assume $\mathfrak{M}, 0 \models \bigvee_{\mathfrak{F} \in \mathcal{F}} rt_{\mathfrak{F},L}(\xi_{\mathfrak{F}})$ with $\mathfrak{M}$ based on $\mathfrak{F}$. Then we can take $\mathfrak{F}' \in \mathcal{F}$ with $\mathfrak{M}, 0 \models rt_{\mathfrak{F}',L}(\xi_{\mathfrak{F}'})$. By Proposition 4.12, there is a σ-bisimilar $\mathfrak{M}'$ based on $\mathfrak{F}'$ with $v_{\mathfrak{M}'} \models \xi_{\mathfrak{F}'}$. As $\xi_{\mathfrak{F}'}$ is a uniform $\sigma_{\mathfrak{F}'}$-interpolant for $tr_{\mathfrak{F}'}(\varphi)$, we can modify the values of non-$\sigma_{\mathfrak{F}'}$ atoms in $v_{\mathfrak{M}'}$ to obtain a model v with $v \models tr_{\mathfrak{F}'}(\varphi)$. But then $\mathfrak{M}'', 0 \sim_\sigma \mathfrak{M}', 0 \sim_\sigma \mathfrak{M}, 0$ and $\mathfrak{M}'', 0 \models \varphi$ for the model $\mathfrak{M}''$ based on $\mathfrak{F}'$ with $v_{\mathfrak{M}''} = v$. ◀

The above theorem can be interpreted as a Turing reduction: it provides an efficient algorithm for computing a strongest $L(\sigma)$-implicate for a given modal formula, where the algorithm is allows to ask one or more queries to an oracle that computes propositional uniform interpolants. In particular, it follows that if propositional uniform interpolants can be computed in polynomial time, then strongest $L(\sigma)$-implicates can be computed in polynomial time. Theorem 1.1 from the introduction (restated below) is slightly stronger, as it provides an algorithm that asks only one oracle query. Our next aim is therefore to strengthen Theorem 4.14 to bridge this gap.

▶ **Theorem 1.1.** *Let L be a tabular quasi-normal modal logic. We can associate to each signature σ of propositional variables a companion signature $\widehat{\sigma}$ such that there are poly-time* *translations*

$$tr_L : ML(\sigma) \to PL(\widehat{\sigma}) \quad \text{and} \quad rt_L : PL(\widehat{\sigma}) \to ML(\sigma)$$

such that for every modal formula φ and every uniform $\widehat{\tau}$-interpolant ψ of $tr_L(\varphi)$ in propositional logic, $rt_L(\psi)$ is a strongest $L(\tau)$-implicate for φ.

Proof. We only sketch the proof here, see the Appendix for a detailed proof. First we extend translation $tr_{\mathfrak{F}}$ to a translation tr_L for any tabular quasi-normal logics L. Fix $L = \mathrm{Log}(\mathcal{F}_L)$ for a finite set $\mathcal{F}_L$ of finite rooted frames. We assume that the worlds of $\mathfrak{F}$ in $(\mathfrak{F}, w) \in \mathcal{F}_L$ are mutually disjoint (and so do not use 0 to denote the roots of frames in $\mathcal{F}_L$). Take atoms $r_{\mathfrak{F},w}$, for $(\mathfrak{F}, w) \in \mathcal{F}_L$. They are used to identify the frame $(\mathfrak{F}, w) \in \mathcal{F}_L$ in which we evaluate a modal formula. For any signature σ, let $\sigma_L = \bigcup_{(\mathfrak{F},w)\in\mathcal{F}_L} \sigma_{\mathfrak{F},w}$. The signatures $\sigma_{\mathfrak{F},w}$ are assumed to be mutually disjoint and also disjoint from $\{r_{\mathfrak{F},w} \mid (\mathfrak{F}, w) \in \mathcal{F}_L\}$. Let $\hat{\sigma} = \sigma_L \cup \{r_{\mathfrak{F},w} \mid (\mathfrak{F}, w) \in \mathcal{F}_L\}$. Next let for any modal formula φ,

$$tr_L(\varphi) = \mathrm{Unique}(\mathcal{F}_L) \wedge \bigwedge_{(\mathfrak{F},w)\in\mathcal{F}_L} (r_{\mathfrak{F},w} \to tr_{\mathfrak{F},w}(\varphi)).$$

where

$$\mathrm{Unique}(\mathcal{F}_L) = (\bigvee_{(\mathfrak{F},w)\in\mathcal{F}_L} (r_{\mathfrak{F},w} \wedge \bigwedge_{(\mathfrak{G},v)\in\mathcal{F}_L\setminus\{(\mathfrak{F},w)\}} \neg r_{\mathfrak{G},v}))$$

Hence $tr_L(\varphi) \in \mathrm{PL}(\hat{\sigma})$ uses $\mathrm{Unique}(\mathcal{F}_L)$ to pick $(\mathfrak{F}, w) \in \mathcal{F}_L$ and then states that $tr_{\mathfrak{F},w}(\varphi)$ holds.

We next define the translation rt_L (which is meaningful only if applied to propositional uniform interpolants). To this end, we first construct from a propositional uniform $\hat{\sigma}$-interpolant of $tr_L(\varphi)$ propositional uniform $\sigma_{\mathfrak{F},w}$-interpolants for $tr_{\mathfrak{F},w}(\varphi)$, for every $(\mathfrak{F}, w) \in \mathcal{F}_L$. Assume φ is given and $\sigma \subseteq \mathrm{sig}(\varphi)$. Let ξ be any propositional uniform $\hat{\sigma}$-interpolant of $tr_L(\varphi)$. Then obtain for $(\mathfrak{F}, w) \in \mathcal{F}$teh formula $\xi^{\uparrow\mathfrak{F},w}$ from ξ by

- replacing $r_{\mathfrak{F},w}$ by $\top$,
- replacing all $r_{\mathfrak{G},v}$ with $(\mathfrak{G}, v) \in \mathcal{F}_L \setminus \{(\mathfrak{F}, w)\}$ by $\bot$, and
- replacing all $p_{\mathfrak{G},v}$ with v not in $(\mathfrak{F}, w)$ by $\bot$.

It is easy to see that $\xi^{\uparrow\mathfrak{F},w}$ is a propositional uniform $\sigma_{\mathfrak{F},w}$-interpolant for $tr_{\mathfrak{F},w}(\varphi)$, for every $(\mathfrak{F}, w) \in \mathcal{F}_L$. Now we simply define for any $\xi \in \mathrm{PL}(\hat{\sigma})$,

$$rt_L(\xi) = \bigvee_{(\mathfrak{F},w)\in\mathcal{F}_L} rt_{(\mathfrak{F},w),L}(\xi^{\uparrow\mathfrak{F},w})$$

and show using Theorem 4.14 that rt_L is as required. ◀

▶ **Corollary 4.15.** *Let L be a quasi-normal tabular logic, $\sigma \subseteq \tau$ signatures and $\rho : \mathbb{N} \to \mathbb{N}$ a monotone function. If every $\chi \in PL(\widehat{\tau})$ has a uniform $\widehat{\sigma}$-interpolant of size at most $\rho(|\chi|)$, then every $\psi \in ML(\tau)$ has a strongest σ-implicate of size at most $poly(\rho(poly(|\psi|)))$.*

5 Constructing Craig Interpolants via Propositional Logic

We show that if a tabular quasi-normal logic $L = \mathrm{Log}(\mathcal{F}_L)$ has CIP, then one can compute Craig interpolants for an implication $\varphi \to \psi \in L$ from propositional Craig interpolants for $tr_L(\varphi) \to tr_L(\psi)$ in polynomial time. To this end, we first give a criterion for CIP for tabular quasi-normal logics which shows that L has CIP iff no non-trivial bisimulations between models based on frames in $\mathcal{F}_L$ exist. We then show that the backward translation $\bigvee_{\mathfrak{F}\in\mathcal{F}_L} rt_{\mathfrak{F},L}(\xi_{\mathfrak{F}})$ is a Craig interpolant for $\varphi \rightarrow \psi$ if the $\xi_{\mathfrak{F}}$ are Craig interpolants for $tr_{\mathfrak{F}}(\varphi) \rightarrow tr_{\mathfrak{F}}(\psi)$ in $\mathrm{Log}(\mathfrak{F})$ for $\mathfrak{F} \in \mathcal{F}_L$.

We close with a discussion of Craig interpolants for modal logics that do not have CIP.

Let $L = \mathrm{Log}(\mathcal{F}_L)$ with $\mathcal{F}_L = \{(\mathfrak{F}_1, 0), \dots, (\mathfrak{F}_n, 0)\}$. Clearly we may assume that $\mathcal{F}_L$ is *reduced* in the sense that $\mathrm{Log}(\mathfrak{F}_i, 0) \not\subseteq \mathrm{Log}(\mathfrak{F}_j, 0)$ for $i \neq j$. We start by giving a criterion for CIP of L. The *σ-reduct* of a model $\mathfrak{M}$ is the restriction of $\mathfrak{M}$ to atoms in σ. An *isomorphism* between the σ-reducts of pointed models $\mathfrak{M}_1, w_1$ and $\mathfrak{M}_2, w_2$ is an isomorphism f between the underlying frames with $f(w_1) = w_2$ such that $\mathfrak{M}_1, w \models p$ iff $\mathfrak{M}_2, f(w) \models p$, for all $p \in \sigma$.

▶ **Theorem 5.1.** *Let $L = Log(\mathcal{F}_L)$ be a tabular quasi-normal logic with $\mathcal{F}_L$ a reduced set of finite rooted frames. Then the following conditions are equivalent:*

1. *L has CIP;*
2. *for all signatures σ: if $\mathfrak{M}, 0 \sim_\sigma \mathfrak{N}, 0$ with $\mathfrak{M}$ a model based on $\mathfrak{F}_i$ and $\mathfrak{N}$ a model based on $\mathfrak{F}_j$, then the σ-reducts of $\mathfrak{M}, 0$ and $\mathfrak{N}, 0$ are isomorphic (in particular, $i = j$).*

The proof is given in the appendix. The following example illustrates Theorem 5.1.

▶ **Example 5.2.** (1) Let $\mathfrak{F}_{1;3}$ and $L_{1;3}$ be from Example 3.2. Observe that for $\sigma = \{p\}$ and $\mathfrak{M}_1$ and $\mathfrak{M}_2$ based of $\mathfrak{F}_{1;3}$ with p true in 1 in $\mathfrak{M}_1$ and p true in $1, 2$ in $\mathfrak{M}_2$ we have $\mathfrak{M}_1, 0 \sim_\sigma \mathfrak{M}_2, 0$ but $\mathfrak{M}_1, 0$ and $\mathfrak{M}_2, 0$ are not isomorphic. Hence $L_{1;3}$ does not have CIP.

(2) Observe that any two models $\mathfrak{M}_1, 0$ and $\mathfrak{M}_2, 0$ based on reflexive frames are σ-bisimilar, for $\sigma = \emptyset$. Hence, if $L = Log(\mathcal{F})$ for a finite set of finite reflexive pointed frames and L has CIP, then $L = Log(\mathfrak{F})$ for a single finite pointed frame $\mathfrak{F}$.

▶ **Theorem 5.3.** *Fix a tabular quasi-normal modal logic L with CIP and with $L = Log(\mathcal{F}_L)$. Let $\varphi \rightarrow \psi \in L$. If $\xi_{\mathfrak{F}} \in PL(\sigma_{\mathfrak{F}})$ are propositional Craig interpolants for $tr_{\mathfrak{F}}(\varphi) \rightarrow tr_{\mathfrak{F}}(\psi)$, $\mathfrak{F} \in \mathcal{F}_L$, then $\bigvee_{\mathfrak{F}\in\mathcal{F}_L} rt_{\mathfrak{F},L}(\xi_{\mathfrak{F}})$ is a Craig interpolant for $\varphi \rightarrow \psi$ in L.*

Proof. Assume $\varphi \rightarrow \psi \in L$ is given. We show that $\varphi \rightarrow \bigvee_{\mathfrak{F}\in\mathcal{F}_L} rt_{\mathfrak{F},L}(\xi_{\mathfrak{F}}) \in L$ and $\bigvee_{\mathfrak{F}\in\mathcal{F}_L} rt_{\mathfrak{F},L}(\xi_{\mathfrak{F}}) \rightarrow \psi \in L$. The claim $\varphi \rightarrow \bigvee_{\mathfrak{F}\in\mathcal{F}_L} rt_{\mathfrak{F},L}(\xi_{\mathfrak{F}}) \in L$ can be shown in exactly the same way as in the proof of Theorem 4.14. Now assume $\mathfrak{M}, 0 \models \bigvee_{\mathfrak{F}\in\mathcal{F}_L} rt_{\mathfrak{F},L}(\xi_{\mathfrak{F}})$ with $\mathfrak{M}$ based on $\mathfrak{F}_i$. We show that $\mathfrak{M}, 0 \models \psi$. Take $\mathfrak{F}_j$ with $\mathfrak{M}, 0 \models rt_{\mathfrak{F}_j,L}(\xi_{\mathfrak{F}_j})$. By Proposition 4.12, there is a σ-bisimilar $\mathfrak{M}', 0$ based on $\mathfrak{F}_j$ with $v_{\mathfrak{M}'} \models \xi_{\mathfrak{F}_j}$. By Theorem 5.1, $i = j$ and we have an isomorphism f between the σ-reducts of $(\mathfrak{M}, 0)$ and $(\mathfrak{M}', 0)$. We modify $\mathfrak{M}'$ in the obvious way to obtain a model $\mathfrak{M}''$ such that f is an isomorphism from $(\mathfrak{M}, 0)$ onto $(\mathfrak{M}'', 0)$. Then $v_{\mathfrak{M}''} \models \xi_{\mathfrak{F}_i}$ since $v_{\mathfrak{M}'}$ and $v_{\mathfrak{M}''}$ coincide on $\sigma_{\mathfrak{F}_i}$. We obtain $v_{\mathfrak{M}''} \models tr_{\mathfrak{F}}(\psi)$ since $\xi_{\mathfrak{F}_i}$ is a Craig interpolant for $tr_{\mathfrak{F}}(\varphi) \rightarrow tr_{\mathfrak{F}}(\psi)$. Hence $\mathfrak{M}'', 0 \models \psi$. But then $\mathfrak{M}, 0 \models \psi$ since $(\mathfrak{M}, 0)$ and $(\mathfrak{M}'', 0)$ are isomorphic. ◀

▶ **Theorem 1.3.** *Let L be a tabular quasi-normal modal logic that has CIP. Given a valid modal implication $\varphi \rightarrow \psi$ in L, and given a propositional Craig interpolant χ for $tr_L(\varphi) \rightarrow tr_L(\psi)$, it holds that $rt_L(\chi)$ is a Craig interpolant for $\varphi \rightarrow \psi$ with respect L.*

The proof is similar to the proof of Theorem 1.1 and uses Theorem 5.3.

We briefly discuss what happens for logics L without CIP. Clearly, in this case there must be $\varphi \rightarrow \psi \in L$ such that for any propositional Craig interpolant ξ for $tr_L(\varphi) \rightarrow tr_L(\psi)$ the modal formula $rt_L(\xi)$ is not a Craig interpolant for φ, ψ in L. For instance, consider $L_{1;3} = \mathrm{Log}(\{(\mathfrak{F}_{1;3}, 0)\})$ defined in Example 3.2. Take the formulas

$$\varphi = (\Diamond(p \land q) \land \Diamond(p \land \neg q)), \quad \psi = (\Diamond(\neg p \land r) \rightarrow \Box(\neg p \rightarrow r))$$

and let $\sigma = \{p\}$. Then the propositional formula ξ stating that "at least two successors satisfy p" is (up to logical equivalence) the propositional uniform $\sigma_{\mathfrak{F}_{1;3}}$-interpolant for $tr_{\mathfrak{F}_{1;3},0}(\varphi)$ and also the unique (up to logical equivalence) Craig interpolant for $tr_{\mathfrak{F}_{1;3},0}(\varphi) \to tr_{\mathfrak{F}_{1;3},0}(\psi)$. The backward translation $rt_{\mathfrak{F}_{1;3},0}(\xi)$ is equivalent to $\Diamond p$. It is a strongest $L_{1;3}(\sigma)$-implicate of φ but neither a uniform $L_{1;3}(\sigma)$-interpolant for φ nor a Craig interpolant for $\varphi \to \psi$ in $L_{1;3}$.

There has recently been work on the computation of Craig interpolants for logics without CIP [3, 22]. Even for a tabular logic without CIP, it is of interest to compute, given $\varphi \to \psi \in L$ such that there exists a Craig interpolant for φ, ψ in L, such a Craig interpolant from any Craig interpolant for $tr_L(\varphi) \to tr_L(\psi)$ in polynomial time. We next show that this is impossible if there are valid implications in propositional logic for which all Craig interpolants are of superpolynomial size. Assume $\varphi_n \to \psi_n$, $n \geq 0$, is a family of propositional formulas such that all Craig interpolants for $\varphi_n \to \psi_n$ are of superpolynomial size. Now consider again the tabular logic $L_{1;3}$. Let σ_n be the signature of φ_n, ψ_n. Now consider

$$\varphi'_n = \Diamond(p \wedge q) \wedge \Diamond(p \wedge \neg q) \wedge \Box\varphi_n$$

and

$$\psi'_n = (\Diamond(\neg p \wedge r) \to \Box(\neg p \to r)) \vee \Box\psi_n$$

where we assume that $\{p, q, r\}$ is disjoint from σ_n. Then we have (i) polynomial size propositional Craig interpolants for $tr_{L_{1;3}}(\varphi'_n) \to tr_{L_{1;3}}(\psi'_n)$, (ii) there are Craig interpolants for $\varphi'_n \to \psi'_n$ in $L_{1;3}$, and (iii) there do not exist any Craig interpolants for $\varphi'_n \to \psi'_n$ in $L_{1;3}$ of polynomial size. Of course, this result does not refute the possibility that a different translation can be used to polynomially reduce Craig interpolants for modal logics without CIP to propositional Craig interpolants.

6 Lower Bounds

Our aim in this section is to prove exponential lower bounds on the size of strongest implicates and Craig interpolants, for large classes of non-tabular modal logics. Our main result is as follows.

▶ **Theorem 1.5.** *Let L be any non-tabular normal modal logic that contains or is contained in S4 or Gödel-Löb provability logic GL, or, more generally, has a non-tabular intersection with S4 or GL.*

1. *There exist modal formulas $(\varphi_n)_{n=1,2,\ldots}$ of size polynomial in n and signatures $(\sigma_n)_{n=1,2,\ldots}$ with $\sigma_n \subseteq sig(\varphi_n)$, such that φ_n has a strongest $L(\sigma_n)$-implicate, and every strongest $L(\sigma_n)$-implicate of φ_n has size at least 2^n.*
2. *There are modal formulas $(\varphi_n)_{n=1,2,\ldots}$ and $(\psi_n)_{n=1,2,\ldots}$ of size polynomial in n such that $\varphi \to \psi \in L$ such that a Craig interpolant for $\varphi \to \psi$ exists, and every Craig interpolant for $\varphi \to \psi$ in L has size at least 2^n.*

While for Craig interpolants this theorem fully captures our results, for strongest implicates Part 1 follows from a significantly stronger result based on sufficient model-theoretic conditions for exponential size implicates. These conditions are defined next. Recall that a quasi-normal modal logic L has the *poly-size model property* if for every formula $\varphi \notin L$ there exists a rooted frame $(\mathfrak{F}, w)$ validating L and refuting φ which is of polynomial size in $|\varphi|$. We say that L has the *exponential growth property* if there is a polynomial function f such that for every $n > 0$ there exists a rooted frame $(\mathfrak{F}, w)$ validating L such that at least 2^n worlds are reachable from w by an R-path of length at most $f(n)$ starting at w. A few standard modal logics such as S5, GL.3, and Grz.3 have both the poly-size model property and exponential growth property. We are in a position now to formulate a sufficient condition for exponential size strongest implicates.

▶ **Theorem 6.1.** *Let L be a quasi-normal modal logic contained in a quasi-normal modal logic L' with the poly-size model property and the exponential growth property. Then there exist modal formulas φ_n and signatures $\sigma_n \subseteq sig(\varphi_n)$ with $|\varphi_n|$ polynomial in n such that the strongest $L(\sigma_n)$-implicates of φ_n exist and are of size $\geq 2^n$.*

Proof. By the poly-size model property of L', we can take $k > 0$ such that satisfiable formulas of size m are satisfied in models of size at most m^k, for sufficiently large m. We also take a polynomial function f witnessing the exponential growth property of L'. Consider atoms $\sigma_n = \{p_0, \ldots, p_{kn-1}\}$ and let $q_0, \ldots, q_{kn-1}$ be a set of atoms disjoint from σ_n. Let

$$\varphi_n = \Box^{\leq f(kn)} \bigvee_{i=0}^{kn-1} \neg(p_i \leftrightarrow q_i) \wedge \Box^{\leq f(kn)-1}(\bigwedge_{i=0}^{kn-1} ((q_i \rightarrow \Box q_i) \wedge (\neg q_i \rightarrow \Box\neg q_i))$$

The formula $\chi_n = \bigvee_{t \in T_{kn}} \Box^{\leq f(n)} \neg t$ is then easily seen to be a strongest $L(\sigma_n)$-implicate of φ_n, where T_n is the set of types $\bigwedge_{i=0}^{kn-1} l_i$ with $l_i \in \{p_i, \neg p_i\}$ for $i = 0, \ldots, kn-1$. Assume now for a proof by contradiction that χ_n is equivalent on frames validating L to a formula ψ_n of size $< 2^n$, for sufficiently large n. Then this is also the case for $L' \supseteq L$. Take $\neg\psi_n$ which states that all 2^{kn} types in T_{kn} are satisfied in worlds reachable in $f(kn)$ steps. By the exponential growth property of L', $\neg\psi_n$ is satisfiable in an L'-frame. We derive a contradiction since then $\neg\psi_n$ is satisfied in a model of size $< 2^{kn}$, for sufficiently large n. ◀

We next spell out the consequences of Theorem 6.1 and show that the first part of Theorem 1.5 follows. Observe first that every non-tabular quasi-normal modal logic containing K4 has the exponential growth property since every world reachable from a world is reachable in one step already. Hence, above K4, it suffices to analyze the poly-size modal property. Call a normal modal logic L *pre-tabular* if every normal modal logic properly containing L is tabular.[1] Pre-tabular logics have been investigated extensively in modal logic as they separate tabular from non-tabular logics. For the following result and further details on pre-tabular logics, we refer the reader to [6, Chapter 12].

▶ **Proposition 6.2.** *(1) Every non-tabular normal modal logic is contained in a pre-tabular normal modal logic.*

(2) All normal modal logics containing a tabular normal modal logic are tabular.

(3) There are five pre-tabular normal modal logics containing S4. These are $L(\mathcal{F})$ for $\mathcal{F}$ depicted in Figure 2.

(4) There are countably many pre-tabular normal modal logics containing GL. These are $L(\mathcal{F})$ for $\mathcal{F}$ depicted in Figure 3.

Observe that Grz.3 $= L(\mathcal{F}_{\text{Grz.3}})$, S5 $= L(\mathcal{F}_{\text{S5}})$ are pre-tabular logics containing S4 and GL.3 $= L(\mathcal{F}_{\text{GL.3}})$ is a pre-tabular logic containing GL. Now the first part of Theorem 1.5 follows from Theorem 6.1 and the following result.

▶ **Proposition 6.3.** *All pre-tabular normal modal logics containing S4 or GL have both the poly-size model property and, trivially, the exponential growth property.*

[1] We use the term pre-tabular logic only in the class of *normal* modal logics. Pre-tabular logics have also been studied in the class of quasi-normal modal logics [6]. The two notions are indeed different. We focus on normal ones for simplicity here.

Figure 2 The finite rooted frames for the five pre-tabular normal logics containing S4. A boundary around a set of worlds indicates a *cluster* of worlds with wRw' for all w, w' in the cluster. Reflexive and transitive edges are omitted for clarity.

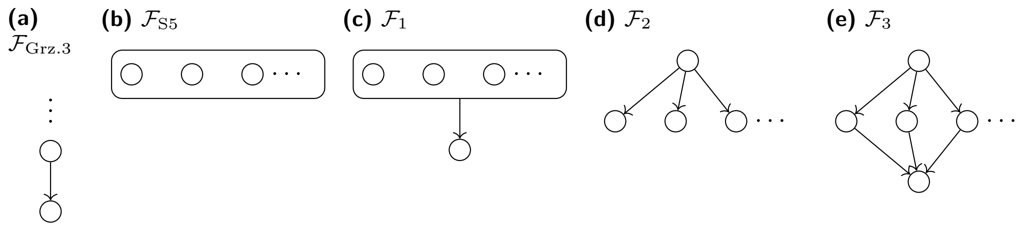

Figure 3 The finite rooted frames for the pre-tabular normal logics containing GL. Transitive edges are omitted for clarity.

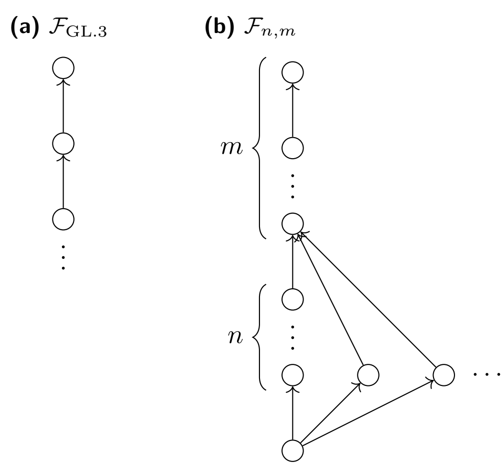

The proof of Proposition 6.3 is straightforward. It also follows from our proof of the second part of Theorem 1.5.

The proof of the second half of Theorem 1.5 is direct and rather lengthy. We define modal formulas $(\varphi_n)_{n=1,2,\dots}$ and $(\psi_n)_{n=1,2,\dots}$ of size polynomial in n and show for each pre-tabular normal modal logic L' containing S4 or GL separately that for every normal modal logic L contained in L' Craig interpolants for $\varphi_n \to \psi_n$ exist but are of exponential size. The proof is given in the appendix.

The logic Alt_1 can be used to show that there are non-tabular normal modal logics for which an exponential lower bound for the size of uniform interpolants would imply $\mathrm{P} \neq \mathrm{NP}$. The proof is again by reduction to uniform interpolants for propositional logic. Observe that Alt_1 has the poly-size model property but does not have the exponential growth property.

▶ **Theorem 6.4.** *Let σ be a signature, and $\widehat{\sigma} = \sigma \times \mathbb{N}$. There are poly-time translations*

$$tr_{alt} : ML(\sigma) \to PL(\widehat{\sigma}) \quad \text{and} \quad rt_{alt} : PL(\widehat{\sigma}) \to ML(\sigma)$$

such that for every modal formula φ, every $\tau \subseteq \sigma$ and every uniform $\tau \times \mathbb{N}$-interpolant ψ for $tr_{alt}(\varphi)$ in propositional logic, $rt_{alt}(\psi)$ is a uniform $Alt_1(\tau)$-interpolant for φ.

Here, if $\varphi \in ML(\sigma)$ has modal depth n, then $tr_{alt}(\varphi) \in PL(\sigma \times \{0, \dots, n\})$.

We also note that it is open whether the dichotomy above can be extended to K4 since the pre-tabular normal logics containing K4 are significantly more complex than the pre-tabular normal modal logics containing S4 or GL. For instance, there are pre-tabular normal modal logics containing K4 without the poly-size model property [6, Chapter 12].

7 Conclusion

We have investigated size bounds for Craig interpolants, uniform interpolants, and strongest implicates, for various families of normal modal logics, and, more generally, quasi-normal modal logics.

For *tabular* quasi-normal modal logics, our results show that size upper bounds for strongest implicates are essentially the same as in the case of propositional logic (modulo polynomial translations). This is, in some sense, the best one can hope for: a polynomial upper bound would imply that $\mathrm{NP} \subseteq \mathrm{P}/_{\mathrm{poly}}$, a longstanding open problem in complexity theory, whereas a super-polynomial lower bound would imply that $\mathrm{NP} \subsetneq \mathrm{P}/_{\mathrm{poly}}$. We obtained analogous results regarding Craig interpolants and uniform interpolants (for tabular quasi-normal modal logics that have CIP or, equivalently, UIP).

Complementing these results for tabular logics, we established exponential lower bounds on the size of strongest implicates, uniform interpolants, and Craig interpolants, which apply to all *non-tabular* normal logics that are contained in or contain S4 or GL (or, more generally, that have a non-tabular intersection with S4 or GL). This covers most non-tabular normal modal logics in practice, and matches known single-exponential upper bounds on the size of Craig interpolants for common modal logics such as K, K4, GL, S5, and Grz, as well as known single-exponential upper bounds on the size of uniform interpolants for S5 and K. The problem of finding tight bound on the size of uniform interpolants for GL and Grz remains open.

The main problem left open is whether one can establish a dichotomy of the following form for all (quasi-)normal modal logics: either computing strongest implicates is poly-time reducible to computing uniform interpolants in classical propositional logic or strongest implicates are of at least exponential size, and similarly for Craig interpolants for modal logics having CIP. To show this, one has to replace tabularity by some weaker condition. For normal modal logics logics containing K4, however, we conjecture that tabularity can still be used to show a dichotomy. Another intriguing problem is the size of Craig interpolants (when they exist) for tabular modal logics without CIP. We have seen that our current reduction does not work in this case, but there could still be a completely different approach. On the other hand, it is known that Craig interpolants often become much harder to compute for logics without CIP (and uniform interpolants for logics without UIP) [17, 16, 15].

As discussed earlier, these results all pertain to formula-size as measured based on a dag representation. Interestingly, it is not difficult to show that our reduction of strongest implicates in tabular quasi-normal modal logics to uniform interpolants in propositional logic (and, similarly, of Craig interpolants for tabular logics with CIP) provides 'propositional sized' implicates also if a standard tree-representation is used. In fact, while the tree-size of the 'forward-translation' $\mathrm{tr}_L(\varphi)$ of a formula φ can be exponential, the 'backward-translation' $rt_L(\varphi)$ is always of polynomial tree-size. The latter suffices to obtain strongest implicates of polynomial tree-size in the size of the propositional uniform interpolant and $\mathrm{tr}_L(\varphi)$ can always be represented, using additional propositional variables, as a formula of polynomial tree-size. This observation regarding rt_L can be used to lift further properties from propositional logic to tabular modal logics. For instance, assume that every propositional formula of dag-size n is equivalent to a propositional formula ψ of tree-size $f(n)$. Then, in tabular modal logics, every modal formula of dag-size n is equivalent to a modal formula of tree-size $poly(f(poly(n)))$.

Note that, for the modal logic K (which has UIP), it is known that the tree-size of uniform interpolants (and hence also strongest σ-implicates and Craig interpolants) can be bounded by an exponential function in the tree-size of input formula(s). On the other hand, it remains an open problem whether the exponential upper bounds for Craig interpolants in K4, S4, GL, and Grz still hold if one uses a tree representation of interpolants.

References

1. Noga Alon and Ravi B. Boppana. The monotone circuit complexity of boolean functions. *Comb.*, 7(1):1–22, 1987. doi:10.1007/BF02579196.
2. Sanjeev Arora and Boaz Barak. *Computational Complexity - A Modern Approach*. Cambridge University Press, 2009. URL: http://www.cambridge.org/catalogue/catalogue.asp?isbn=9780521424264.
3. Alessandro Artale, Jean Christoph Jung, Andrea Mazzullo, Ana Ozaki, and Frank Wolter. Living without Beth and Craig: Definitions and interpolants in description and modal logics with nominals and role inclusions. *ACM Trans. Comput. Log.*, 24(4):34:1–34:51, 2023. doi:10.1145/3597301.
4. Marta Bílková. Uniform interpolation and propositional quantifiers in modal logics. *Stud Logica*, 85(1):1–31, 2007. URL: https://doi.org/10.1007/s11225-007-9021-5, doi:10.1007/S11225-007-9021-5.
5. Patrick Blackburn, Maarten de Rijke, and Yde Venema. *Modal Logic*. Cambridge Tracts in Theoretical Computer Science. Cambridge University Press, 2001.
6. Alexander V. Chagrov and Michael Zakharyaschev. *Modal Logic*, volume 35 of *Oxford logic guides*. Oxford University Press, 1997.
7. Giovanna D'Agostino and Marco Hollenberg. Logical questions concerning the mu-calculus: Interpolation, lyndon and los-tarski. *J. Symb. Log.*, 65(1):310–332, 2000. doi:10.2307/2586539.
8. Hugo Férée, Iris van der Giessen, Sam van Gool, and Ian Shillito. Mechanised uniform interpolation for modal logics k, gl, and isl. In Christoph Benzmüller, Marijn J. H. Heule, and Renate A. Schmidt, editors, *Automated Reasoning - 12th International Joint Conference, IJCAR 2024, Nancy, France, July 3-6, 2024, Proceedings, Part II*, volume 14740 of *Lecture*

Notes in Computer Science, pages 43–60. Springer, 2024. doi:10.1007/978-3-031-63501-4_3.

9 Dov M Gabbay. Craig's interpolation theorem for modal logics. In *Conference in mathematical logic—London'70*, pages 111–127. Springer, 1972.

10 Dov M. Gabbay and Larisa Maksimova. *Interpolation and Definability: Modal and Intuitionistic Logics.* Oxford University Press, 2005.

11 Dov M. Gabbay and Larisa Maksimova. *Interpolation and Definability: Modal and Intuitionistic Logics.* Oxford University Press, 05 2005. doi:10.1093/acprof:oso/9780198511748.001.0001.

12 Silvio Ghilardi, Carsten Lutz, Frank Wolter, and Michael Zakharyaschev. Conservative extensions in modal logic. In Guido Governatori, Ian M. Hodkinson, and Yde Venema, editors, *Advances in Modal Logic 6, papers from the sixth conference on "Advances in Modal Logic," held in Noosa, Queensland, Australia, on 25-28 September 2006*, pages 187–207. College Publications, 2006. URL: http://www.aiml.net/volumes/volume6/Ghilardi-Lutz-Wolter-Zakharyaschev.ps.

13 Silvio Ghilardi and Marek W. Zawadowski. Undefinability of propositional quantifiers in the modal system S4. *Stud Logica*, 55(2):259–271, 1995.

14 Jean Christoph Jung and Jedrzej Kolodziejski. Modal separation of fixpoint formulae. In Olaf Beyersdorff, Michal Pilipczuk, Elaine Pimentel, and Kim Thang Nguyen, editors, *42nd International Symposium on Theoretical Aspects of Computer Science, STACS 2025, March 4-7, 2025, Jena, Germany*, volume 327 of *LIPIcs*, pages 55:1–55:20. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2025. doi:10.4230/LIPICS.STACS.2025.55.

15 Jean Christoph Jung, Jędrzej Kołodziejski, and Frank Wolter. Computation of interpolants for description logic concepts in hard cases. *arXiv preprint arXiv:2507.15689*, 2025.

16 Jean Christoph Jung, Carsten Lutz, Hadrien Pulcini, and Frank Wolter. Separating data examples by description logic concepts with restricted signatures. In Meghyn Bienvenu, Gerhard Lakemeyer, and Esra Erdem, editors, *Proceedings of the 18th International Conference on Principles of Knowledge Representation and Reasoning, KR 2021, Online event, November 3-12, 2021*, pages 390–399, 2021. URL: https://doi.org/10.24963/kr.2021/37, doi:10.24963/KR.2021/37.

17 Jean Christoph Jung and Frank Wolter. Living without beth and craig: Definitions and interpolants in the guarded and two-variable fragments. In *36th Annual ACM/IEEE Symposium on Logic in Computer Science, LICS 2021, Rome, Italy, June 29 - July 2, 2021*, pages 1–14. IEEE, 2021. doi:10.1109/LICS52264.2021.9470585.

18 Richard Karp. Turing machines that take advice. *Enseign. Math.*, 28:191–209, 1982.

19 Richard M Karp and Richard J Lipton. Some connections between nonuniform and uniform complexity classes. In *Proceedings of the twelfth annual ACM symposium on Theory of computing*, pages 302–309, 1980.

20 Patrick Koopmann, Christoph Wernhard, and Frank Wolter. Interpolation in classical propositional logic, 2025. URL: https://arxiv.org/abs/2508.11449, arXiv:2508.11449.

21 Marcus Kracht. *Tools and Techniques in Modal Logic.* Elsevier, 1999.

22 Agi Kurucz, Frank Wolter, and Michael Zakharyaschev. A non-uniform view of Craig interpolation in modal logics with linear frames. *CoRR*, abs/2312.05929, 2023.

23 Paolo Liberatore. The ghosts of forgotten things: A study on size after forgetting. *Ann. Pure Appl. Log.*, 175(8):103456, 2024. URL: https://doi.org/10.1016/j.apal.2024.103456, doi:10.1016/J.APAL.2024.103456.

24 Carsten Lutz and Frank Wolter. Foundations for uniform interpolation and forgetting in expressive description logics. In *IJCAI 2011, Proceedings of the 22nd International Joint Conference on Artificial Intelligence, Barcelona, Catalonia, Spain, July 16-22, 2011*, pages 989–995. IJCAI/AAAI, 2011. doi:10.5591/978-1-57735-516-8/IJCAI11-170.

25 Maarten Marx. Interpolation in modal logic. In *Proceedings of the 7th International Conference on Algebraic Methodology and Software Technology, AMAST 1998*, pages 154–163. Springer, 1998.

26 Daniele Mundici. A lower bound for the complexity of Craig's interpolants in sentential logic. *Arch. Math. Log.*, 23(1):27–36, 1983. `doi:10.1007/BF02023010`.

27 Konstantinos Papafilippou and David Fernández-Duque. Exponential lower bounds on definable fixed points. In Jörg Endrullis and Sylvain Schmitz, editors, *33rd EACSL Annual Conference on Computer Science Logic, CSL 2025, February 10-14, 2025, Amsterdam, Netherlands*, volume 326 of *LIPIcs*, pages 25:1–25:19. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2025. URL: `https://doi.org/10.4230/LIPIcs.CSL.2025.25`, `doi:10.4230/LIPICS.CSL.2025.25`.

28 Andrew M. Pitts. On an interpretation of second order quantification in first order intuitionistic propositional logic. *J. Symb. Log.*, 57(1):33–52, 1992. `doi:10.2307/2275175`.

29 Alexander Alexandrovich Razborov. Lower bounds on the monotone complexity of some boolean functions. In *Doklady Akademii Nauk*, volume 281, pages 798–801. Russian Academy of Sciences, 1985.

30 Benjamin Rossman. Formulas versus circuits for small distance connectivity. *SIAM J. Comput.*, 47(5):1986–2028, 2018. `doi:10.1137/15M1027310`.

31 Uwe Schöning and Jacobo Torán. A note on the size of Craig interpolants. In Thomas Schwentick, Denis Thérien, and Heribert Vollmer, editors, *Circuits, Logic, and Games, 08.11. - 10.11.2006*, volume 06451 of *Dagstuhl Seminar Proceedings*. Internationales Begegnungs- und Forschungszentrum fuer Informatik (IBFI), Schloss Dagstuhl, Germany, 2006. URL: `http://drops.dagstuhl.de/opus/volltexte/2007/973`.

32 Fatemeh Seifan, Lutz Schröder, and Dirk Pattinson. Uniform interpolation in coalgebraic modal logic. *CoRR*, abs/2205.00448, 2022. URL: `https://doi.org/10.48550/arXiv.2205.00448`, `arXiv:2205.00448`, `doi:10.48550/ARXIV.2205.00448`.

33 Balder ten Cate, Willem Conradie, Maarten Marx, and Yde Venema. Definitorially complete description logics. In *Proc. of KR*, pages 79–89, 2006.

34 Balder ten Cate, Enrico Franconi, and Inanç Seylan. Beth definability in expressive description logics. *J. Artif. Intell. Res.*, 48:347–414, 2013.

35 Iris van der Giessen, Raheleh Jalali, and Roman Kuznets. Uniform interpolation via nested sequents and hypersequents. *J. Log. Comput.*, 35(6), 2025. URL: `https://doi.org/10.1093/logcom/exae053`, `doi:10.1093/LOGCOM/EXAE053`.

36 Albert Visser et al. Uniform interpolation and layered bisimulation. In *Gödel'96: Logical foundations of mathematics, computer science and physics—Kurt Gödel's legacy*, pages 139–164. Association for Symbolic Logic, 1996.

37 Frank Wolter and Michael Zakharyaschev. Modal decision problems. In Patrick Blackburn, J. F. A. K. van Benthem, and Frank Wolter, editors, *Handbook of Modal Logic*, volume 3 of *Studies in logic and practical reasoning*, pages 427–489. North-Holland, 2007. URL: `https://doi.org/10.1016/s1570-2464(07)80010-3`, `doi:10.1016/S1570-2464(07)80010-3`.

A Further Details for Section 2

We give the inductive definition of the truth-condition between models $\mathfrak{M} = (\mathfrak{F}, V)$ with $\mathfrak{F} = (W, R)$ and formulas φ:

- $\mathfrak{M}, w \models p$ if $w \in V(p)$;
- $\mathfrak{M}, w \models \neg\varphi$ if $\mathfrak{M}, w \not\models \varphi$;
- $\mathfrak{M}, w \models \varphi \wedge \psi$ if $\mathfrak{M}, w \models \varphi$ and $\mathfrak{M}, w \models \psi$;
- $\mathfrak{M}, w \models \Box\varphi$ if for all v with wRv it follows that $\mathfrak{M}, v \models \varphi$.

We have defined quasi-normal and normal modal logics L of the form $L = \text{Log}(\mathcal{F})$ with $\mathcal{F}$ a class of pointed frames. Such (quasi-)normal modal logics are called *Kripke-complete*. By definition, all tabular (quasi-)normal modal logics are Kripke-complete. However, there are (quasi-)normal modal logics which are not Kripke-complete [6]. They do not play an important role in this paper, but since some of our results about exponential lower bounds on the size of implicates and interpolants speak about all (quasi-)normal modal logics, we give a quick definition. Let $K = \text{Log}(\mathcal{F}_K)$, where $\mathcal{F}_K$ is the class of all pointed frames. A *quasi-normal modal logic* is a set L of modal formulas containing K which is closed under modus ponens and uniform substitutions of atoms by modal formulas. It is easy so see that every $L = \text{Log}(\mathcal{F})$ with $\mathcal{F}$ a class of pointed frames is indeed a quasi-normal modal logic. A *normal modal logic* is a quasi-normal modal logic that is also closed under the rule: if $\varphi \in L$, then $\Box\varphi \in L$. Again, it is easy to see that every $L = \text{Log}(\mathcal{F})$ with $\mathcal{F}$ containing any $(\mathfrak{F}_v, v)$ if $(\mathfrak{F}, w) \in \mathcal{F}$ and v in $\mathfrak{F}$, is indeed a normal modal logic according to this definition.

B Proofs for Section 4

▶ **Lemma 4.2.** *Fix a finite frame $\mathfrak{F} = (W, R)$. For every signature σ, we can compute in polynomial time a set $\Gamma_{\mathfrak{F}}$ of σ-EMEs such that every model based on $\mathfrak{F}$ has a σ-cover in $\Gamma_{\mathfrak{F}}$. Moreover, each $\Phi \in \Gamma_{\mathfrak{F}}$ consists of at most $|W|$ formulas, each of size $O(|W|)$.*

Proof. Let $\mathfrak{M} = (\mathfrak{F}, V)$ be a model based on $\mathfrak{F}$. We can obtain a σ-cover for $\mathfrak{M}$ by using a refinement procedure. We begin with $\Psi_0 = \{\top\}$, which is a σ-EME. Then, we obtain Ψ_{m+1} from Ψ_m as follows. If, for some $\psi \in \Psi_m$ and $p \in \sigma$, both $\psi \wedge p$ and $\psi \wedge \neg p$ are satisfied in $\mathfrak{M}$, we let $\Psi_{m+1} = \Psi_m \setminus \{\psi\} \cup \{\psi \wedge p, \psi \wedge \neg p\}$. Otherwise, we let $\Psi_{m+1} = \Psi_m$. This process will terminate after at most $|W| - 1$ refinements, call the result Φ. This Φ is by construction a σ-EME and any two elements of $\mathfrak{M}$ that satisfy the same ψ are propositionally indistinguishable over σ, since otherwise further refinement would have been possible. Hence Φ is a σ-cover of $\mathfrak{M}$.

Note that not every σ-cover is of the form generated by this procedure, but every $\mathfrak{M}$ has at least one such σ-cover. Furthermore, every $\psi \in \Phi$ is a conjunction of at most $|W|$ literals, and hence $|\psi| \leq O(|W|)$.

To find a bound on the number of σ-covers we need, it is convenient to represent them as labeled trees $(\mathit{vert}, \mathit{edge}, \mathit{label})$ where $\mathit{edge} \subseteq \mathit{vert} \times \mathit{vert}$ and label is a function that assigns each vertex $x \in \mathit{vert}$ a formula $\mathit{label}(x)$.

Corresponding with Ψ_0, we take $T_0 = (\mathit{vert}_0, \mathit{edge}_0, \mathit{label}_0)$ where $\mathit{vert} = \{x_0\}$, $\mathit{edge}_0 = \emptyset$ and $\mathit{label}_0(x_0) = \top$. Then, if from Ψ_m to Ψ_{m+1} we split ψ into $\psi \wedge p$ and $\psi \wedge \neg p$, then we take

- $\mathit{vert}_{m+1} = \mathit{vert}_m \cup \{x_m, y_m\}$,
- $\mathit{edge}_{m+1} = \mathit{edge}_m \cup \{(x, x_m), (x, y_m)\}$, where x is the leaf such that $\mathit{label}_m(x) = \psi$, and

- $label_{m+1}(x) = \begin{cases} \psi \wedge p & \text{if } x = x_m \\ \psi \wedge \neg p & \text{if } x = y_m \\ label_m(x) & \text{otherwise} \end{cases}$

and if $\Psi_m = \Psi_{m+1}$ then we take $T_m = T_{m+1}$. Call the final tree $T = (\mathit{vert}, \mathit{edges}, \mathit{label})$.

The unlabeled tree $\overline{T} = (\mathit{vert}, \mathit{edges})$ is a sub-tree of the binary tree of depth $|W|$, and has at most $|W|$ leaves. Each such unlabeled tree can therefore be identified by choosing at most N leaves from among the $2^{|W|}$ nodes of the binary tree. So there are at most $\binom{2^{|W|}}{|W|} \leq \left(2^{|W|}\right)^{|W|} = 2^{|W|^2}$ unlabeled trees. Any labeling is determined by the atoms that were used, in each non-leaf vertex, to differentiate between the two successors of that vertex. There are at most $|W| - 1$ such non-leaf vertices, so there are at most $|\sigma|^{|W|-1}$ labelings for a given unlabeled tree. There are, therefore, at most $2^{|W|^2} \cdot |\sigma|^{|W|-1}$ different σ-covers of this type. Since every model has such a σ-cover, the same number bounds the amount of σ-covers that we need.

The time bound follows because we can construct $\Gamma_{\mathfrak{F}}$ by iterating over all labeled trees. ◀

▶ **Proposition 4.4.** *Let $\mathfrak{M}, 0$ be a rooted model, Φ a σ-cover of $\mathfrak{M}$ and $\mathfrak{A}$ the Φ-abstraction of $\mathfrak{M}$. Then for every rooted frame $(\mathfrak{F}', 0)$ with $\mathfrak{F}' = (W', R')$ and for every $Z \subseteq W \times W'$ the following are equivalent:*

1. *$Z : (\mathfrak{M}, 0) \sim_\sigma (\mathfrak{M}', 0)$ for some model $\mathfrak{M}'$ based on $\mathfrak{F}'$,*
2. *$Z : (\mathfrak{A}, 0) \sim_\Phi (\mathfrak{A}', 0)$ for some abstract Φ-model $\mathfrak{A}'$ based on $\mathfrak{F}'$.*

Furthermore, if such $\mathfrak{A}'$ and $\mathfrak{M}'$ exist then $\mathfrak{A}'$ is the Φ-abstraction of $\mathfrak{M}'$.

Proof. For (1) $\Rightarrow$ (2), suppose that $Z : (\mathfrak{M}, 0) \sim_\sigma (\mathfrak{M}', 0')$. As Z is a bisimulation between the roots of $\mathfrak{M}$ and $\mathfrak{M}'$, it must be a total relation. Hence for every $w' \in W$ there is some $w \in W$ such that wZw'. Because Z is a bisimulation, this implies that $V(w) = V'(w')$. This implies that Φ is a σ-cover of $\mathfrak{M}'$ as well as of $\mathfrak{M}$, so $\mathfrak{M}'$ has a Φ-abstraction. Let $\mathfrak{A}' = (\mathfrak{F}', f')$ be the Φ-abstraction of $\mathfrak{M}'$.

Because Z is a bisimulation between $\mathfrak{M}, 0$ and $\mathfrak{M}', 0$, it satisfies **back** and **forth**. In order to show that it is an abstract bisimulation between $\mathfrak{A}, 0$ and $\mathfrak{A}', 0$ it therefore suffices to show that Z satisfies **abstract atoms**.

Suppose, therefore, that wZw'. Because Z is a σ-bisimulation, we have $\mathfrak{M}, w \models p$ iff $\mathfrak{M}', w' \models p$ for all $p \in \sigma$. As each $\varphi \in \Phi$ is a propositional formula over σ, it follows that $\mathfrak{M}, w \models \varphi$ iff $\mathfrak{M}', w' \models \varphi$.

We therefore have $f(w) = \varphi \Leftrightarrow \mathfrak{M}, w \models \varphi \Leftrightarrow \mathfrak{M}', w' \models \varphi \Leftrightarrow f'(w') = \varphi$. This holds for all $\varphi \in \Phi$, so **abstract atoms** is satisfied, as was to be shown.

For (2) $\Rightarrow$ (1), suppose that $Z : (\mathfrak{A}, 0) \sim_\Phi (\mathfrak{A}', 0)$. As Z is an abstract bisimulation between the roots of $\mathfrak{A}$ and $\mathfrak{A}'$, it must be a total relation. For any $w' \in W'$, fix a single $w_{w'} \in W$ such that $w_{w'}Zw$, and let $\mathfrak{M}' = (\mathfrak{F}', V')$ be the model such that $V'(w') = V(w_{w'})$. We will show that Z is a total σ-bisimulation between $\mathfrak{M}$ and $\mathfrak{M}'$, that Φ is a σ-cover of $\mathfrak{M}'$ and that $\mathfrak{A}'$ is the Φ-abstraction of $\mathfrak{M}'$.

Because Z is an abstract bisimulation, it satisfies **back** and **forth**. To show that it is a σ-bisimulation, it therefore suffices to show that it satisfies **atoms**. Take any $w \in W, w' \in W'$ such that wZw'.

Then both wZw' and $w_{w'}Zw'$. So $f(w) = f'(w') = f(w_{w'})$, because Z is an abstract bisimulation. Since Φ is a σ-cover of $\mathfrak{M}$ it follows from $f(w) = f(w_{w'})$ that $V_\sigma(w) = V_\sigma(w_{w'})$. Furthermore, by construction, $V(w_{w'}) = V'(w')$. It follows that $V_\sigma(w) = V'_\sigma(w')$, so **atoms** is satisfied, as was to be shown.

Now, to show that Φ is a σ-cover of $\mathfrak{M}'$. Take any w'_1, w'_2 such that $\mathfrak{M}', w'_1 \models \varphi$ and $\mathfrak{M}', w'_2 \models \varphi$, for some $\varphi \in \Phi$. Then $\mathfrak{M}, w_{w'_1} \models \varphi$ and $\mathfrak{M}, w_{w'_2} \models \varphi$, so because Φ is a σ-cover of $\mathfrak{M}$ we have $\mathfrak{M}, w_{w_1'} \equiv_{\mathrm{PL}(\sigma)} \mathfrak{M}, w_{w_2'}$. By construction we also have $\mathfrak{M}, w_{w_1'} \equiv_{\mathrm{PL}(\sigma)} \mathfrak{M}', w_1'$ and $\mathfrak{M}, w_{w_2'} \equiv_{\mathrm{PL}(\sigma)} \mathfrak{M}', w_2'$, so we obtain $\mathfrak{M}', w_1' \equiv_{\mathrm{PL}(\sigma)} \mathfrak{M}', w_2'$. So Φ is a σ-cover of $\mathfrak{M}'$.

Finally, take any $w' \in W'$ and let $\varphi = f'(w')$. Because $w_{w'} Z w'$ and Z is both an abstract bisimulation between $\mathfrak{A}$ and $\mathfrak{A}'$ and a σ-bisimulation between $\mathfrak{M}$ and $\mathfrak{M}'$, we have $f(w_{w'}) = f'(w') = \varphi$, and $\mathfrak{M}, w_{w'} \models \varphi$ iff $\mathfrak{M}', w' \models \varphi$. Also, because $\mathfrak{A}$ is the Φ-abstraction of $\mathfrak{M}$, we have $\mathfrak{M}, w_{w'} \models f(w_{w'})$. Taken together, this shows that $\mathfrak{M}', w' \models f'(w')$, which completes the proof that $\mathfrak{A}'$ is the Φ-abstraction of $\mathfrak{M}'$. ◀

▶ **Proposition 4.5.** *Fix a natural number $N > 0$. Given a σ-EME Φ (with σ a signature) and a rooted abstract Φ-model $(\mathfrak{A}, 0)$ of size at most N, we can compute in polynomial time a formula $\delta_{\mathfrak{A}} \in ML(\Phi)$ such that, for all rooted abstract Φ-models $(\mathfrak{A}', 0)$ of size at most N, $\mathfrak{A}', 0 \models \delta_{\mathfrak{A}}$ iff $\mathfrak{A}', 0 \sim_\Phi \mathfrak{A}, 0$. Moreover, $|\delta_{\mathfrak{A}}| \leq O(2^N \cdot \max_{\varphi\in\Phi} |\varphi|)$.*

Proof. Let $\mathfrak{A} = (\mathfrak{F}, f)$ with $\mathfrak{F} = (W, R)$. We begin by constructing, for every $w \in W$, a characteristic formula $\varphi_w \in \mathrm{ML}(\Phi)$ such that $\mathfrak{A}, w' \models \varphi_w$ if, and only if, $\mathfrak{A}, w$ and $\mathfrak{A}, w'$ are abstractly Φ-bisimilar. We do this through the standard partition refinement procedure. For every $\varphi \in \Phi$, we let $W_\varphi = \{w \in W \mid f(w) = \varphi\}$, and we take $\mathcal{I}_0 = \{(W_\varphi, \varphi) \mid \varphi \in \Phi\}$. Then, from $\mathcal{I}_m$ we obtain $\mathcal{I}_{m+1}$ by taking any two $(J, \psi), (J', \psi') \in \mathcal{I}_m$ and checking whether there are $j_1, j_2 \in J$ such that there exists $j_1' \in J'$ with $(j_1, j_1') \in R$ but $(j_2, j_2') \notin R$ for all $j_2 \in J'$. If so, then in $\mathcal{I}_{m+1}$ we replace the pair (J, ψ) by the two pairs

$$(\{j \in J \mid \exists j' \in J' : (j, j') \in R\}, \psi \wedge \Diamond\psi') \text{ and } ((\{j \in J \mid \forall j' \in J' : (j, j') \notin R\}, \psi \wedge \neg\Diamond\psi'),$$

otherwise we stop and set $\mathcal{I} = \mathcal{I}_m$. In this way, we obtain a set $\mathcal{I} = \{(I_1, \psi_1), \cdots, (I_k, \psi_k)\}$ such that $I_1, \ldots, I_k$ form a partition of W and $\mathfrak{A}, w \models \psi_i$ iff $w \in I_i$, for $1 \leq i \leq k$.

A straightforward induction on n shows that, for each $(J, \psi) \in \mathcal{I}_n$ $(n \geq 0)$, it holds that $|\psi| \leq O(2^n) \cdot \max_{\varphi\in\Phi} |\varphi|$. Furthermore, since in each iteration the partition is refined and each $\mathcal{I}_n$ can contain at most $|\mathfrak{F}| \leq N$ elements, the above process must stabilize after at most $m \leq N$ iterations. We can therefore conclude that each ψ_i, for $i \leq k$, has size at most $O(2^N) \cdot \max_{\varphi\in\Phi} |\varphi|$.

For $w \in W$, we write $\psi_{[w]}$ for the unique formula such that there is a J such that $(J, \psi_{[w]}) \in \mathcal{I}$ and $w \in J$. Now, let

$$\mathit{char}(\mathcal{I}, \mathfrak{F}) = \Big(\Box^{\leq N} \bigvee_{(J,\psi)\in\mathcal{I}} \psi\Big) \wedge \Big(\bigwedge_{(J,\psi)\in\mathcal{I}} \bigwedge_{(J',\psi')\in\mathcal{I}} \Box^{\leq N}(\psi \to \pm\Diamond\psi')\Big),$$

where $\pm\Diamond\psi'$ is $\Diamond\psi'$ if there are $j \in J, j' \in J'$ such that $(j, j') \in R$ and $\neg\Diamond\psi'$ otherwise, and take

$$\delta = \mathit{char}(\mathcal{I}, \mathfrak{F}) \wedge \psi_{[0]}.$$

It is not difficult to see that $|\delta| \leq O(2^N) \cdot \max_{\varphi\in\Phi} |\varphi|$ and δ can be computed in time $O(2^N) \cdot \Sigma_{\varphi\in\Phi}|\varphi|$.

It remains only to show that for every rooted abstract Φ-model $\mathfrak{A}', 0$ of size at most N, we have $\mathfrak{A}', 0 \models \delta$ iff $\mathfrak{A}', 0 \sim_\Phi \mathfrak{A}, 0$.

($\Rightarrow$) Suppose that $\mathfrak{A}', 0 \models \delta$. Let $Z \subseteq W \times W'$ be the relation such that wZw' iff $\mathfrak{A}', w' \models \psi_{[w]}$. We will show that Z is an abstract Φ-bisimulation between $\mathfrak{A}$ and $\mathfrak{A}'$ and that $(\mathfrak{A}, 0)Z(\mathfrak{A}', 0)$.

That $(\mathfrak{A}, 0)Z(\mathfrak{A}', 0)$ is the easy part: one of the conjuncts of δ is $\psi_{[0]}$, so from $\mathfrak{A}', 0 \models \delta$ it follows that, in particular, $\mathfrak{A}', 0 \models \psi_{[0]}$ and therefore $0Z0$. We continue to show that Z is an abstract Φ-bisimulation.

For **abstract atoms**, suppose that wZw'. We have $\mathfrak{A}, w \models \psi_{[w]}$ by the construction of $\psi_{[w]}$. By the construction of Z it follows from wZw' that we also have $\mathfrak{A}', w' \models \psi_{[w]}$. One conjunct of $\psi_{[w]}$ is some $\varphi \in \Phi$, so we have $\mathfrak{A}, w \models \varphi$ and $\mathfrak{A}', w' \models \varphi$, so $f(w) = f'(w') = \varphi$, as required for **abstract atoms** to hold.

For **forth**, suppose that $w_1 Z w_1'$ and $(w_1, w_2) \in R$. Because any world can be reached from 0 in at most N steps, it follows from $\mathfrak{A}', 0 \models \delta$ that $\mathfrak{A}', w_1' \models \psi_{[s_1]} \to \pm\Diamond\psi_{[w_2]}$. Furthermore, because $(w_1, w_2) \in R$ the $\pm$ is positive, so $\mathfrak{A}', w_1' \models \psi_{[w_1]} \to \Diamond\psi_{[w_2]}$.

From $w_1 Z w_1'$ it follows that $\mathfrak{A}', w_1' \models \psi_{[w_1]}$, which implies that $\mathfrak{A}', w_1' \models \Diamond\psi_{[w_2]}$. Hence there is some $w_2' \in W'$ such that $(w_1', w_2') \in R'$ and $\mathfrak{A}', w_2' \models \psi_{[w_2]}$. This implies that $w_2 Z w_2'$. We now have $(w_1', w_2') \in R'$ and $w_2 Z w_2'$, so **forth** is satisfied.

For **back**, suppose that $w_1 Z w_1'$ and $(w_1', w_2') \in R'$. Because $\mathfrak{M}', 0 \models \mathit{char}(\mathcal{I}, \mathfrak{F})$ and w_2' can be reached from 0 in at most N steps, we have $\mathfrak{M}', w_2' \models \psi$ for some $(J, \psi) \in \mathcal{I}$. This implies that $\mathfrak{M}', w_1' \models \Diamond\psi$ and therefore $\mathfrak{M}', w_1' \not\models \psi_{[w_1]} \to \neg\Diamond\psi$. We must therefore have $\pm\Diamond\psi = \Diamond\psi$, so there is some $w_2 \in J$ such that $\psi_{[w_2]} = \psi$ and $(w_1, w_2) \in R$. This implies that $(w_1, w_2) \in R$ and $w_2 Z w_2'$, so **back** is satisfied. This completes the left-to-right direction of the proof that δ is the abstract class identifier for $[\mathfrak{A}]$.

$(\Leftarrow)$ Suppose that $Z : \mathfrak{A}, 0 \sim_\Phi \mathfrak{A}', 0$. To show is that $\mathfrak{A}', 0 \models \delta$.

In order to do so, first note that because Z is an abstract Φ-bisimulation, and therefore satisfies **abstract atoms**, we have for every $(W_\varphi, \varphi) \in \mathcal{I}_0$, every $w \in W_\varphi$ and every wZw' that $\mathfrak{A}', w' \models \varphi$.

Because of the way the $\mathcal{I}_m$ are created through a partition refinement procedure, it follows from **back** and **forth** that for every $\mathcal{I}_m$, every $(J, \psi) \in \mathcal{I}_m$, every $w \in J$ and every wZw' that $\mathfrak{A}', w' \models \psi$.

It immediately follows that $\mathfrak{A}', 0 \models \psi_{[0]}$. Furthermore, because Z is an abstract bisimulation that relates $\mathfrak{A}, 0$ and $\mathfrak{A}', 0$, it follows from the fact that all worlds are reachable from 0 that Z is a total relation. So for every $w' \in W'$ there is a $w \in W$ such that wZw'. There is also some $(J, \psi) \in \mathcal{I}$ such that $w \in J$. It follows that $\mathfrak{A}', w' \models \psi$, and therefore $\mathfrak{A}', 0 \models \bigwedge_{0 \leq l \leq N} \Box^l \bigvee_{(J,\psi) \in \mathcal{I}} \psi$.

Now, take any $(J, \psi), (J', \psi') \in \mathcal{I}$. We need to show that, for any $w_1' \in W'$, we have $\mathfrak{A}', w' \models \psi \to \pm\Diamond\psi'$. So suppose that $\mathfrak{A}', w' \models \psi$. Then wZw' for some $w \in J$. If $\pm\Diamond\psi' = \Diamond\psi'$ then there is some $w_2 \in J'$ such that $(w_1, w_2) \in R$. By **forth** there is then some $w_2' \in W'$ such that $w_2 Z w_2'$ and $(w_1', w_2') \in R'$. Because $w_2 \in J'$ and $w_2 Z w_2'$ we have $\mathfrak{A}', w_2' \models \psi'$, and hence $\mathfrak{A}', w_1' \models \Diamond\psi'$, so $\mathfrak{A}', w_1' \models \psi \to \pm\Diamond\psi'$.

If $\pm\Diamond\psi' = \neg\Diamond\psi'$ then for all $w_2 \in J'$ we have $(w_1, w_2) \notin R$. By **back** it follows that for all $w_2' \in W'$ such that $w_2 Z w_2'$ we have $(w_1', w_2') \notin R'$. This implies that for every w_3' such that $(w_1', w_3') \in R'$ there is some $w_3 \notin J'$ such that $w_3 Z w_3'$. Therefore, $\mathfrak{A}', w_3' \not\models \psi'$. As this holds for all such w_3', we have $\mathfrak{A}', w_1' \models \neg\Diamond\psi'$ and therefore $\mathfrak{A}', w_1' \models \psi \to \pm\Diamond\psi'$.

We have now shown that $\mathfrak{A}', w_1' \models \psi \to \pm\Diamond\psi'$ for all $w_1' \in W'$ and all $(J, \psi), (J', \psi') \in \mathcal{I}$. This implies that

$$\mathfrak{A}', 0 \models \bigwedge_{(J,\psi) \in \mathcal{I}} \bigwedge_{(J',\psi') \in \mathcal{I}} \Box^{\leq N}(\psi \to \pm\Diamond\psi').$$

Together with the previously shown $\mathfrak{A}', 0 \models \Box^{\leq N} \bigvee_{(J,\psi) \in \mathcal{I}} \psi$ and $\mathfrak{A}', 0 \models \psi_{[0]}$ this shows that $\mathfrak{A}', 0 \models \delta$, which completes the right-to-left direction. ◀

▶ **Proposition 4.12.** *Fix a tabular quasi-normal modal logic L. Let $\mathfrak{M}, 0$ be any pointed model based on a pointed frame in $\mathcal{F}_L$. Furthermore, let $\mathfrak{F} \in \mathcal{F}_L$ and let $\xi \in PL(\sigma_{\mathfrak{F}})$. Then the following are equivalent:*

1. $\mathfrak{M}, 0 \models rt_{\mathfrak{F},L}(\xi)$,
2. *there is a model* $\mathfrak{M}'$ *based on the frame* $\mathfrak{F}$ *such that* $\mathfrak{M}, 0 \sim_\sigma \mathfrak{M}', 0$ *and* $v_{\mathfrak{M}'} \models \xi$.

The proof uses the following lemma, which we prove first.

▶ **Lemma B.1.** *Let* Φ *be a* σ*-cover of* $\mathfrak{M}$*,* Z *a* σ*-bisimulation such that* $(\mathfrak{M}, w)Z(\mathfrak{M}', w')$ *and* $\varphi \in \Phi$ *the formula such that* $\mathfrak{M}', w' \models \varphi$*. Then for every* $p \in \sigma$*,* $v_{\mathfrak{M}'} \models p_{w'} \Leftrightarrow \mathfrak{M}, 0 \models \Diamond^{\leq N}(\varphi \wedge p)$*.*

Proof. ($\Rightarrow$) Suppose $v_{\mathfrak{M}'} \models p_{w'}$. Then $\mathfrak{M}', w' \models p$. As Z is a σ-bisimulation, this implies that $\mathfrak{M}, w \models \varphi$. Furthermore, from $\mathfrak{M}', w' \models \varphi$ we also have $\mathfrak{M}, w \models \varphi$. This implies that $\mathfrak{M}, w \models \varphi \wedge p$, and therefore $\mathfrak{M}, 0 \models \Diamond^{\leq N}(\varphi \wedge p)$.

($\Leftarrow$) Suppose $\mathfrak{M}, 0 \models \Diamond^{\leq N}(\varphi \wedge p)$. Because Z is a σ-bisimulation we have $\mathfrak{M}, w \models \varphi$. As Φ is a σ-cover of $\mathfrak{M}$, the value of p is constant across all worlds that satisfy φ. From $\mathfrak{M}, 0 \models \Diamond^{\leq N}(\varphi \wedge p)$ it therefore follows that $\mathfrak{M}, w \models p$, which implies that $\mathfrak{M}', w' \models p$. It follows that $v_{\mathfrak{M}'} \models p_{w'}$. ◀

Proof of Proposition 4.12. ($\Rightarrow$) Suppose $\mathfrak{M}, 0 \models rt_{\mathfrak{F}}(\xi)$. Because $(\Gamma, \{\Delta_\Phi\}_{\Phi \in \Gamma})$ is a σ-encoding, there are $\Phi \in \Gamma, \delta \in \Delta_\Phi$ such that Φ is a σ-cover of $\mathfrak{M}$ and $\mathfrak{M}, 0 \models \delta$. Let $\mathfrak{A}$ be the Φ-abstraction of $\mathfrak{M}$. We then have $\mathfrak{M}, 0 \models cover(\sigma, \Phi) \wedge \delta$, so the antecedent of $rt_{\mathfrak{F}}(\xi)$ is satisfied for these Φ and δ. By assumption, $\mathfrak{M}, 0 \models rt_{\mathfrak{F}}(\xi)$, so the consequent of $rt_{\mathfrak{F}}(\xi)$ must also hold, i.e., $\mathfrak{M}, 0 \models \bigvee_{(\mathfrak{F}, f) \in [\delta]_L} \xi^f$. Let $\mathfrak{A}' = (\mathfrak{F}, f) \in [\delta]_L$ be such that $\mathfrak{M}, 0 \models \xi^f$.

We are given that $\mathfrak{A}' \in [\delta]_L$. Because $\mathfrak{M}, 0 \models cover(N, \sigma, \Phi) \wedge \delta$ we also have $\mathfrak{A} \in [\delta]_L$. So there is an abstract Φ-bisimulation Z between $\mathfrak{A}$ and $\mathfrak{A}'$. It follows that there is some $\mathfrak{M}' = (\mathfrak{F}', V')$ such that $\mathfrak{A}'$ is the Φ-abstraction of $\mathfrak{M}'$ and Z is a σ-bisimulation between $\mathfrak{M}$ and $\mathfrak{M}'$. Because $\mathfrak{A}' = (\mathfrak{F}, f)$ is the Φ-abstraction of $\mathfrak{M}'$, we have $\mathfrak{M}', w' \models f(w')$ for all $w' \in W'$. The conditions of Lemma B.1 are therefore satisfied, which implies that $\mathfrak{M}, 0 \models \Diamond^{\leq N}(f(w') \wedge p) \Leftrightarrow v_{\mathfrak{M}'} \models p_{w'}$. From $\mathfrak{M}, 0 \models \xi^f$ it therefore follows that $v_{\mathfrak{M}'} \models \xi$. We already showed that $\mathfrak{M}, 0$ and $\mathfrak{M}', 0$ are σ-bisimilar, so this complete the left-to-right direction of the proof.

($\Leftarrow$) Suppose that $\mathfrak{M}' = (\mathfrak{F}, V')$ is σ-bisimilar to $\mathfrak{M}$ and that $v_{\mathfrak{M}'} \models \xi$. We must show that $\mathfrak{M}, \models rt_{\mathfrak{F}}(\xi)$. Let $\Phi \in \Gamma$ and $\delta \in \Delta_\Phi$ be such that the antecedent of $rt_{\mathfrak{F}}(\xi)$ is satisfied, i.e., $\mathfrak{M}, 0 \models cover(N, \sigma, \Phi) \wedge \delta$. From $\mathfrak{M}, 0 \models cover(N, \sigma, \Phi)$ it follows that Φ is a σ-cover of $\mathfrak{M}$. Let $\mathfrak{A}$ be the Φ-abstraction of $\mathfrak{M}$. Because $\mathfrak{M}, 0 \models \delta$ we also have $\mathfrak{A}, 0 \models \delta$, so $\mathfrak{A} \in [\delta]_L$.

Since $\mathfrak{M}, 0 \sim_\Phi \mathfrak{M}', 0$, we have that Φ is also a σ-cover of $\mathfrak{M}'$, and that $\mathfrak{M}', 0 \models \delta$. Let $\mathfrak{A}' = (\mathfrak{F}, f)$ be the Φ-abstraction of $\mathfrak{M}'$. Then $\mathfrak{A}', 0 \models \delta$, and therefore $\mathfrak{A}' \in [\delta]_L$. Furthermore, for every $w' \in W'$ we have $\mathfrak{M}', w' \models f(w')$. By Lemma B.1 we now have $v_{\mathfrak{M}'} \models p_{w'}$ if, and only if, $\mathfrak{M}, 0 \models \Diamond^{\leq N}(f(w') \wedge p)$. Since, by assumption, $v_{\mathfrak{M}'} \models \xi$, this implies that $\mathfrak{M}, 0 \models \xi^f$. The consequent of $rt_{\mathfrak{F}}(\xi)$ is therefore also satisfied in $\mathfrak{M}, 0$. As this holds for every Φ and δ for which the antecedent is satisfied, we have $\mathfrak{M}, 0 \models rt_{\mathfrak{F}}(\xi)$, as was to be shown. ◀

▶ **Theorem 1.1.** *Let* L *be a tabular quasi-normal modal logic. We can associate to each signature* σ *of propositional variables a companion signature* $\widehat{\sigma}$ *such that there are poly-time translations*

$$tr_L : ML(\sigma) \to PL(\widehat{\sigma}) \quad \textit{and} \quad rt_L : PL(\widehat{\sigma}) \to ML(\sigma)$$

such that for every modal formula φ *and every uniform* $\widehat{\tau}$*-interpolant* ψ *of* $tr_L(\varphi)$ *in propositional logic,* $rt_L(\psi)$ *is a strongest* $L(\tau)$*-implicate for* φ*.*

Proof. We extend the translation $tr_{\mathfrak{F}}$ to a translation tr_L for any tabular quasi-normal logics L. Fix $L = \mathrm{Log}(\mathcal{F}_L)$ for a finite set $\mathcal{F}_L$ of finite rooted frames. We assume that the worlds of $\mathfrak{F}$ in $(\mathfrak{F}, w) \in \mathcal{F}_L$ are mutually disjoint (and so do not use 0 to denote the roots of frames in $\mathcal{F}_L$). Take atoms $r_{\mathfrak{F},w}$, for $(\mathfrak{F}, w) \in \mathcal{F}_L$. They are used to identify the frame $(\mathfrak{F}, w) \in \mathcal{F}_L$ in which we evaluate a modal formula. For any signature σ, let $\sigma_L = \bigcup_{(\mathfrak{F},w)\in\mathcal{F}_L} \sigma_{\mathfrak{F},w}$. The signatures $\sigma_{\mathfrak{F},w}$ are assumed to be mutually disjoint and also disjoint from $\{r_{\mathfrak{F},w} \mid (\mathfrak{F}, w) \in \mathcal{F}_L\}$. Let $\hat{\sigma} = \sigma_L \cup \{r_{\mathfrak{F},w} \mid (\mathfrak{F}, w) \in \mathcal{F}_L\}$.

Next let for any modal formula φ,

$$tr_L(\varphi) = \mathrm{Unique}(\mathcal{F}_L)) \wedge \bigwedge_{(\mathfrak{F},w)\in\mathcal{F}_L} (r_{\mathfrak{F},w} \to tr_{\mathfrak{F},w}(\varphi)).$$

where

$$\mathrm{Unique}(\mathcal{F}_L) = (\bigvee_{(\mathfrak{F},w)\in\mathcal{F}_L} (r_{\mathfrak{F},w} \wedge \bigwedge_{(\mathfrak{G},v)\in\mathcal{F}_L\setminus\{(\mathfrak{F},w)\}} \neg r_{\mathfrak{G},v}))$$

Note that if $\varphi \in \mathrm{ML}(\sigma)$, then $tr_L(\varphi) \in \mathrm{PL}(\hat{\sigma})$. For any model $\mathfrak{M}$ based on a frame $(\mathfrak{F}, w) \in \mathcal{F}_L$ we obtain a propositional model $v_{\mathfrak{M}}$ defined by setting

- $v_{\mathfrak{M}}(r_{\mathfrak{F},w}) = 1$ and $v_{\mathfrak{M}}(r_{\mathfrak{G},v}) = 0$ for $(\mathfrak{G}, v) \in \mathcal{F}_L \setminus \{(\mathfrak{F}, w)\}$;
- $v_{\mathfrak{M}}(p_{\mathfrak{F},v}) = 1$ iff $\mathfrak{M}, v \models p$, for all v in $\mathfrak{F}$;
- $v_{\mathfrak{M}}(p_{\mathfrak{G},v'}) \in \{0, 1\}$ arbitrary for all worlds v' in any $(\mathfrak{G}, v) \in \mathcal{F}_L \setminus \{(\mathfrak{F}, w)\}$.

▷ Claim B.2. Let $(\mathfrak{F}, w) \in \mathcal{F}_L$ and $\mathfrak{M}$ be a model based on $\mathfrak{F}$. Then

1. for every modal formula φ, we have $\mathfrak{M}, 0 \models \varphi$ iff $v_{\mathfrak{M}} \models tr_L(\varphi)$.
2. For every modal implication $\varphi \to \psi$, we have $\varphi \to \psi \in L$ iff $tr_L(\varphi) \to tr_L(\psi)$ is a tautology.

The proof is straightforward. The following claims imply that we can transform any propositional uniform $\hat{\sigma}$-interpolant of $tr_L(\varphi)$ in polynomial time into an equivalent one of the form $\mathrm{Unique}(\mathcal{F}_L) \wedge \bigwedge_{(\mathfrak{F},w)\in\mathcal{F}_L}(r_{(\mathfrak{F},w)} \to \xi_{\mathfrak{F},w})$ with $\xi_{\mathfrak{F},w}$ a propositional uniform $\sigma_{\mathfrak{F},w}$-interpolant of $tr_{\mathfrak{F},w}(\varphi)$.

▷ Claim B.3. Let φ be a modal formula and $\sigma \subseteq \mathrm{sig}(\varphi)$. Let ξ be any propositional uniform $\hat{\sigma}$-interpolant of $tr_L(\varphi)$ and let $\xi_{\mathfrak{F},w}$ be any propositional uniform $\sigma_{\mathfrak{F},w}$-interpolant of $tr_{\mathfrak{F},w}(\varphi)$, for $(\mathfrak{F}, w) \in \mathcal{F}_L$. Then ξ and $\mathrm{Unique}(\mathcal{F}_L) \wedge \bigwedge_{(\mathfrak{F},w)\in\mathcal{F}_L}(r_{\mathfrak{F},w} \to \xi_{\mathfrak{F},w})$ are logically equivalent.

To show Claim B.3, it suffices to show that $\mathrm{Unique}(\mathcal{F}_L) \wedge \bigwedge_{(\mathfrak{F},w)\in\mathcal{F}_L}(r_{\mathfrak{F},w} \to \xi_{\mathfrak{F},w})$ is a uniform $\hat{\sigma}$-interpolant of $tr_L(\varphi)$ (since uniform interpolants are unique up to logical equivalence). We first show that

$$tr_L(\varphi) \to (\mathrm{Unique}(\mathcal{F}_L) \wedge \bigwedge_{(\mathfrak{F},w)\in\mathcal{F}_L} (r_{\mathfrak{F},w} \to \xi_{\mathfrak{F},w}))$$

is a tautology. Consider a model v satisfying $tr_L(\varphi)$. Then we have a unique $(\mathfrak{F}, w) \in \mathcal{F}_L$ with $v \models r_{\mathfrak{F},w}$. We also have $v \models tr_{\mathfrak{F},w}(\varphi)$. But then also $v \models \mathrm{Unique}(\mathcal{F}_L) \wedge \bigwedge_{(\mathfrak{F},w)\in\mathcal{F}_L}(r_{\mathfrak{F},w} \to \xi_{\mathfrak{F},w})$. Next one can easily show that every model v of $\mathrm{Unique}(\mathcal{F}_L) \wedge \bigwedge_{(\mathfrak{F},w)\in\mathcal{F}_L}(r_{\mathfrak{F},w} \to \xi_{\mathfrak{F},w})$ can be expanded to a model v' of $tr_L(\varphi)$. The claim follows directly.

Let φ be a modal formula and $\sigma \subseteq \mathrm{sig}(\varphi)$. Let ξ be any propositional uniform $\hat{\sigma}$-interpolant of $tr_L(\varphi)$. For $(\mathfrak{F}, w) \in \mathcal{F}$, obtain $\xi^{\uparrow\mathfrak{F},w}$ from ξ by

- replacing $r_{\mathfrak{F},w}$ by $\top$,
- replacing all $r_{\mathfrak{G},v}$ with $(\mathfrak{G}, v) \in \mathcal{F}_L \setminus \{(\mathfrak{F}, w)\}$ by $\bot$, and

- replacing all $p_{\mathfrak{G},v}$ with v not in $(\mathfrak{F}, w)$ by $\bot$.

▷ Claim B.4. $\xi^{\uparrow\mathfrak{F},w}$ is a propositional uniform $\sigma_{\mathfrak{F},w}$-interpolant for $tr_{\mathfrak{F},w}(\varphi)$, for all $(\mathfrak{F}, w) \in \mathcal{F}_L$.

To show Claim B.4, we first show that $tr_{\mathfrak{F},w}(\varphi) \to \xi^{\uparrow\mathfrak{F},w}$ is tautology. Let v be model of $tr_{\mathfrak{F},w}(\varphi)$. We obtain model of $tr_L(\varphi)$ by expanding v to v' with

- $v'(r_{\mathfrak{F},w}) = 1$,
- $v'(r_{\mathfrak{G},v}) = 0$ for $(\mathfrak{G}, v) \in \mathcal{F}_L \setminus \{(\mathfrak{F}, w)\}$, and
- $v'(p_{\mathfrak{G},v}) = 0$ for v not in $(\mathfrak{F}, w)$.

But then v' satisfies ξ (since ξ is a uniform $\hat{\sigma}$-interpolant of $tr_L(\varphi)$). Hence, by definition, v satisfies $\xi^{\uparrow\mathfrak{F},w}$. Next observe that every model of $\xi^{\uparrow\mathfrak{F},w}$ can be expanded to a model of ξ is the obvious way. This model then satisfies $tr_{\mathfrak{F},w}(\varphi)$, by definition of the formulas. The claim follows directly.

Define for any $\xi \in \mathrm{PL}(\hat{\sigma})$,

$$rt_L(\xi) = \bigvee_{(\mathfrak{F},w)\in\mathcal{F}_L} rt_{(\mathfrak{F},w),L}(\xi^{\uparrow\mathfrak{F},w}).$$

It follows from Theorem 4.14 and Claim B.4 that rt_L is as required. ◀

C Illustrating Examples

Consider the following frame $\mathfrak{F}$, with 0 as the root.

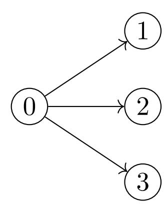

We will consider a few models based on this frame. Let $\mathfrak{M}_1, \cdots, \mathfrak{M}_4$ be as shown below, from left to right.

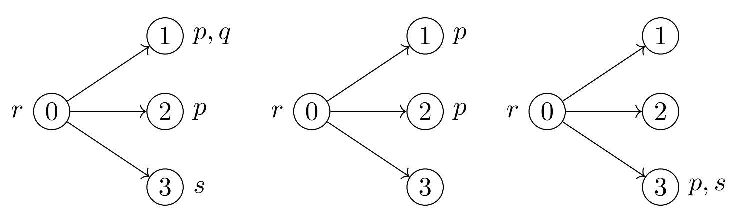

Let $\sigma = \{p, r, s\}$ and consider $\mathfrak{M}_1$. There are several choices for a σ-cover of $\mathfrak{M}_1$. For example, $\Phi = \{p, \neg p \wedge r, \neg p \wedge \neg r\}$ is one, as is $\Phi' = \{\neg r \wedge \neg s, r, \neg r \wedge s\}$. The set Φ is also a σ-cover of $\mathfrak{M}_2$ and $\mathfrak{M}_3$, but Φ' is not.

Neither Φ nor Φ' is a σ-cover of $\mathfrak{M}_4$. In fact, no three-element set is a σ-cover of $\mathfrak{M}_4$, as all four worlds of $\mathfrak{M}_4$ are propositionally distinguishable over σ.

For the remainder of the example, we will consider only the cover Φ, so let $\varphi_1 = p, \varphi_2 = \neg p \wedge r$ and $\varphi_3 = \neg p \wedge \neg r$, so $\Phi = \{\varphi_1, \varphi_2, \varphi_3\}$. Consider the following three abstract Φ-models $\mathfrak{A}_1$, $\mathfrak{A}_2$ and $\mathfrak{A}_3$, respectively.

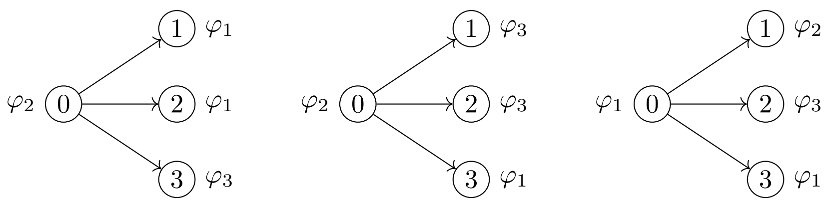

In order for an abstract model $\mathfrak{A}$ to be the Φ-abstraction of a model $\mathfrak{M}$, two things have to be the case. Firstly, Φ must be a σ-cover of $\mathfrak{M}$. Secondly, for each world w we must have $\mathfrak{M}, w \models f(w)$. It is straightforward to verify that $\mathfrak{A}_1$ is the Φ-abstraction of $\mathfrak{M}_1$ and $\mathfrak{M}_2$, and that $\mathfrak{A}_2$ is the Φ-abstraction of $\mathfrak{M}_3$. The model $\mathfrak{M}_4$ doesn't have a Φ-abstraction, since Φ is not a σ-cover of $\mathfrak{M}_4$.

Note that the choice of σ-cover matters. The set Φ' is also a σ-cover of $\mathfrak{M}_1$, but unlike Φ it is not a σ-cover of $\mathfrak{M}_2$ and $\mathfrak{M}_3$. Hence there are no Φ'-abstractions of $\mathfrak{M}_2$ and $\mathfrak{M}_3$.

The models $\mathfrak{M}_1$, $\mathfrak{M}_2$ and $\mathfrak{M}_3$ are not bisimilar to each other. Yet their Φ-abstractions $\mathfrak{A}_1$ and $\mathfrak{A}_2$ are abstractly bisimilar. So abstract bisimilarity is a coarser comparison than bisimilarity. Neither of these abstract models is abstractly bisimilar to $\mathfrak{A}_3$, however.

The proof of Proposition 4.5 shows a systematic way to create an abstract class identifier for any abstract model. Such identifiers are rather unwieldy, however, so here we will consider some hand-designed simpler ones. An abstract class identifier for $[\mathfrak{A}_1] = [\mathfrak{A}_2]$ is $\delta = \varphi_2 \wedge \Box(\varphi_1 \vee \varphi_3) \wedge \Diamond\varphi_1 \wedge \Diamond\varphi_3$. An abstract class identifier for $[\mathfrak{A}_3]$ is $\delta' = \varphi_1 \wedge \Diamond\varphi_1 \wedge \Diamond\varphi_2 \wedge \Diamond\varphi_3$.

Now, let us determine the strongest σ-implicate for a given modal formula. Let $\psi = \Diamond(p \wedge q) \wedge \Diamond(p \wedge \neg q) \wedge \Diamond s$. We are looking for the strongest σ-implicate. Note that $q \notin \sigma$, so we will need to "forget" q.

The first step is to determine the formula $tr_{\mathfrak{F}}(\psi)$. While this formula is of polynomial size with respect to $|\psi|$, it is nonetheless far too long to display here. Fortunately, we can see that ψ expresses that at least one of the worlds $1, 2$ and 3 satisfies $p \wedge q$, at least one satisfies $p \wedge \neg q$ and at least one (not necessarily a different one) satisfies s. The formula $tr_{\mathfrak{F}}(\psi)$ is equivalent to ψ (on the rooted frame $\mathfrak{F}, 0$), so it expresses the same thing.

The uniform σ-interpolant for $tr_{\mathfrak{F}}(\psi)$ therefore expresses merely that at least two out of $0, 1$ and 2 satisfy p and at least one satisfies s, so that interpolant is equivalent to $\xi = ((p_1 \wedge p_2) \vee (p_2 \wedge p_3) \vee (p_3 \wedge p_1)) \wedge (s_1 \vee s_2 \vee s_3)$.

There are many different covers and abstract bisimulation classes, so we cannot feasibly write down all of $rt(\xi)$ here. We can, however, discuss individual conjuncts. For example, consider the conjunct for Φ and δ_1. For any model $\mathfrak{M}$, we have $\mathfrak{M}, 0 \models \mathit{cover}(\sigma, \Phi) \wedge \delta$ if, and only if, Φ is a σ-cover of $\mathfrak{M}$ and the abstraction of $\mathfrak{M}$ is a member of $[\delta]$. So, in particular, this is true for $\mathfrak{M}_1, 0$, $\mathfrak{M}_2, 0$ and $\mathfrak{M}_3, 0$.

In order for $rt(\xi)$ to be true in any of these models, the formula $\bigvee_{(\mathfrak{F}, f') \in C(\delta)} \xi[p_w \mapsto \Diamond^{\leq N}(f'(w) \wedge p)]$ would therefore have to hold there. For $\mathfrak{M}_1, 0$, that is the case. We can simply take $\mathfrak{A}_1 = (\mathfrak{F}, f') \in C(\delta)$, so $f'(0) = \varphi_2 = \neg p \wedge r$, $f'(1) = f'(2) = \varphi_1 = p$ and

$f'(3) = \varphi_3 = \neg p \wedge \neg r$. Then

$$\begin{aligned}
\xi[p_w \mapsto \Diamond^{\le N}(f'(w) \wedge p)] = &((\Diamond^{\le N}(p \wedge f'(1)) \wedge \Diamond^{\le N}(p \wedge f'(2)))\vee \\
&(\Diamond^{\le N}(p \wedge f'(2)) \wedge \Diamond^{\le N}(p \wedge f'(3)))\vee \\
&(\Diamond^{\le N}(p \wedge f'(3)) \wedge \Diamond^{\le N}(p \wedge f'(1))))\wedge \\
&(\Diamond^{\le N}(p \wedge f'(1)) \vee \Diamond^{\le N}(p \wedge f'(2)) \vee \Diamond^{\le N}(p \wedge f'(3))) \\
= &((\Diamond^{\le N}(p \wedge p) \wedge \Diamond^{\le N}(p \wedge p)) \vee (\Diamond^{\le N}(p \wedge p) \wedge \Diamond^{\le N}(p \wedge \neg p \wedge \neg r))\vee \\
&\quad (\Diamond^{\le N}(p \wedge \neg p \wedge r) \wedge \Diamond^{\le N}(p \wedge p)))\wedge \\
&(\Diamond^{\le N}(s \wedge p) \vee \Diamond^{\le N}(s \wedge p) \vee \Diamond^{\le N}(s \wedge \neg p \wedge \neg r))
\end{aligned}$$

which simplifies to $\Diamond^{\le N} p \wedge (\Diamond^{\le N}(s \wedge p) \vee \Diamond^{\le N}(s \wedge \neg p \wedge \neg r))$, which is true in $\mathfrak{M}_1, 0$.

For $\mathfrak{M}_3$, we can also take $\mathfrak{A}_1 = (\mathfrak{F}, f') \in C(\delta)$. We again obtain $\Diamond^{\le N} p \wedge (\Diamond^{\le N}(s \wedge p) \vee \Diamond^{\le N}(s \wedge \neg p \wedge \neg r))$, which holds in $\mathfrak{M}_3, 0$.

There is no $(\mathfrak{F}, f') \in C(\delta)$ such that $\mathfrak{M}_2, 0 \models \xi[p_w \mapsto \Diamond^{\le N}(f'(w) \wedge p)]$, however. This is because $\xi[p_w \mapsto \Diamond^{\le N}(f'(w) \wedge p)]$ has a conjunct of the form $(\Diamond^{\le N}(s \wedge f'(1)) \vee \Diamond^{\le N}(s \wedge f'(2)) \vee \Diamond^{\le N}(s \wedge f'(3)))$, which can't be true in $\mathfrak{M}_2, 0$ since that model contains no worlds satisfying s (and therefore also no worlds satisfying $s \wedge f'(w)$ for any f' and any w).

In fact, if we analyse the above we can see that $rt(\xi)$ is equivalent to $\Diamond p \wedge \Diamond s$. This is unsurprising, this is clearly the strongest σ-implicate of $\Diamond(p \wedge q) \wedge \Diamond(p \wedge \neg q) \wedge \Diamond s$. Note that this strongest implicate is not a uniform interpolant, since there are formulas that are entailed by φ and that do not follow from $\Diamond p \wedge \Diamond s$, e.g., $\Box(\neg p \to t) \vee \Box(\neg p \to \neg t)$.

For contrast, consider also the frame $\mathfrak{F}$ shown below.

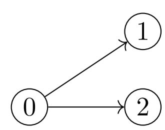

On this frame, ψ implies that one of the worlds 1 and 2 satisfies $p \wedge q$ and the other satisfies $p \wedge \neg q$ (and one of them satisfies s). The uniform σ-interpolant of $tr(\psi)$ is therefore $\xi = p_1 \wedge p_2 \wedge (s_1 \vee s_2)$.

Again we cannot consider each σ-cover Φ and each abstract class identifier δ. Let us look at $\Phi = \{r, \neg r\}$ and $\delta = r \wedge \Diamond r \wedge \Diamond \neg r$. There are exactly two abstract models in $[\delta]$, namely the one given by $f_1(0) = f_1(1) = r$ and $f_1(2) = \neg r$, and the one given by $f_2(0) = f_2(2) = r$ and $f_2(1) = \neg r$.

Let $\mathfrak{M}$ be any model based on this frame such that $\mathfrak{M}, 0 \models \mathit{cover}(\sigma, \Phi) \wedge \delta$. Then in order for $tr(\xi)$ to hold in $\mathfrak{M}, 0$, we must have $\mathfrak{M}, 0 \models \xi[p_w \mapsto \Diamond^{\le N}(f'(w) \wedge p)]$ for some $\mathfrak{A} = (\mathfrak{F}, f') \in C(\delta)$, so for some $f' \in \{f_1, f_2\}$. Taking $f' = f_1$, we obtain $\Diamond^{\le N}(p \wedge r) \wedge \Diamond^{\le N}(p \wedge \neg r) \wedge (\Diamond^{\le N}(s \wedge r) \vee \Diamond^{\le N}(s \wedge \neg r))$. Taking $f' = f_2$ results in the equivalent $\Diamond^{\le N}(p \wedge \neg r) \wedge \Diamond^{\le N}(p \wedge r) \wedge (\Diamond^{\le N}(s \wedge \neg r) \vee \Diamond^{\le N}(s \wedge r))$.

Note that, by assumption, $\mathfrak{M}, 0 \models \delta$, so the root and one of the successors satisfy r while the other successor satisfies $\neg r$. Hence if $\mathfrak{M}, 0 \models \Diamond^{\le N}(p \wedge \neg r)$ then there must be a $p \wedge \neg r$ world, which cannot be the root (since that satisfies r) and therefore has to be one of the successors. If $\mathfrak{M}, 0 \models \Diamond^{\le N}(p \wedge r)$, then either the root satisfies $p \wedge r$ or one of the successors does. Importantly, we also assume that $\mathfrak{M}, 0 \models \mathit{cover}(\sigma, \Phi)$, so any two orlds of $\mathfrak{M}$ that are not distinguished by $\{r, \neg r\}$ must be propositionally indistinguishable over σ. Hence the root and the r-satisfying successor are indistinguishable, which implies that if either of them satisfy $p \wedge r$ then both do.

The formula $\Diamond^{\leq N}(p \wedge r) \wedge \Diamond^{\leq N}(p \wedge \neg r)$ it therefore equivalent, on models that satisfy $cover(\sigma, \Phi) \wedge \delta$, to $\Box p$. Similarly, $\Diamond^{\leq N}(s \wedge r) \vee \Diamond^{\leq N}(s \wedge \neg r)$ is equivalent, on such models, to $\Diamond s$.

If we repeated this procedure on different Φ and δ we would obtain the same results, $\bigvee_{(\mathfrak{F}, f') \in C(\delta)} \xi[p_2 \mapsto \Diamond^{\leq N}(f'(w) \wedge p)]$ is equivalent, on models that satisfy $cover(\sigma, \Phi) \wedge \delta$, to $\Box p \wedge \Diamond q$. It follows that the strongest σ-implicate of $\Diamond(p \wedge q) \wedge \Diamond(p \wedge \neg q) \wedge \Diamond s$ on this frame is $\Box p \wedge \Diamond s$. Unlike the implicate $\Diamond p \wedge \Diamond s$ that we obtained for the model with three successors, this $\Box p \wedge \Diamond s$ is not only the strongest σ-implicate but also the uniform σ-interpolant for $\Diamond(p \wedge q) \wedge \Diamond(p \wedge \neg q) \wedge \Diamond s$.

D Proofs for Section 5

▶ **Theorem 5.1.** *Let $L = Log(\mathcal{F}_L)$ be a tabular quasi-normal logic with $\mathcal{F}_L$ a reduced set of finite rooted frames. Then the following conditions are equivalent:*

1. *L has CIP;*
2. *for all signatures σ: if $\mathfrak{M}, 0 \sim_\sigma \mathfrak{N}, 0$ with $\mathfrak{M}$ a model based on $\mathfrak{F}_i$ and $\mathfrak{N}$ a model based on $\mathfrak{F}_j$, then the σ-reducts of $\mathfrak{M}, 0$ and $\mathfrak{N}, 0$ are isomorphic (in particular, $i = j$).*

Proof. We use the following criterion for CIP for tabular modal logics which follows immediately from known sufficient conditions [25]:

Fact 1. L has CIP iff for all modal formulas φ, ψ and $\sigma = \text{sig}(\varphi) \cap \text{sig}(\psi)$ the following holds: if there are models $\mathfrak{M}, 0 \models \varphi$ and $\mathfrak{N}, 0 \models \psi$ with $\mathfrak{M}, 0 \sim_\sigma \mathfrak{N}, 0$ such that $\mathfrak{M}$ is based on $\mathfrak{F}_i$ and $\mathfrak{N}$ is based on $\mathfrak{F}_j$, then there is a model $\mathfrak{M}', 0$ based on some $\mathfrak{F}_k$ with $\mathfrak{M}', 0 \models \varphi \wedge \psi$.

Fact 1 directly implies (2) $\Rightarrow$ (1). Assume (2) holds. We aim to show CIP. To apply Fact 1, assume φ, ψ are given, $\sigma = \text{sig}(\varphi) \cap \text{sig}(\psi)$ and we have models $\mathfrak{M}, 0 \models \varphi$ and $\mathfrak{N}, 0 \models \psi$ with $\mathfrak{M}, 0 \sim_\sigma \mathfrak{N}, 0$ such that $\mathfrak{M}$ is based on $\mathfrak{F}_i$ and $\mathfrak{N}$ is based on $\mathfrak{F}_j$. By (2), we have $i = j$ and we have an isomorphism $f : \mathfrak{F}_i \to \mathfrak{F}_i$ with $f(0) = 0$ and $\mathfrak{M}, w \models p$ iff $\mathfrak{N}, f(w) \models p$ for all w in $\mathfrak{F}_i$ and $p \in \sigma$. But then we obtain a model $\mathfrak{M}'$ based on $\mathfrak{F}_i$ with $\mathfrak{M}', 0 \models \varphi \wedge \psi$ by taking the $\text{sig}(\varphi)$-reduct of $\mathfrak{M}$ and setting $\mathfrak{M}', w \models p$ iff $\mathfrak{N}, f(w) \models p$ for all $p \in \text{sig}(\psi) \setminus \text{sig}(\varphi)$.

Next we show (1) $\Rightarrow$ (2). Assume $(\mathfrak{F}, 0)$ is a finite rooted frame and $N = \max_{\mathfrak{F} \in \mathcal{F}_L} |\mathfrak{F}|$. Define the *diagram* $\text{diag}_N(\mathfrak{F}, 0)$ of $(\mathfrak{F}, 0)$ (also known as the splitting-, Jankov-, or canonical formula of $(\mathfrak{F}, 0)$ [37]) as the conjunction of the following formulas (with q_w fresh atoms for w in $\mathfrak{F}$):

q_0;
$\Box^{\leq N} \bigvee_{v \in W} q_v$
$\Box^{\leq N}(q_v \to \neg q_{v'})$ if $v \neq v'$
$\Box^{\leq N}(q_v \to \Diamond q_{v'})$ if vRv'
$\Box^{\leq N}(q_v \to \neg \Diamond q_{v'})$ if not vRv'.

The following fact is straightforward:

Fact 2. (1) If there exists a model $\mathfrak{M}$ based on $\mathfrak{F}_j$ with $\mathfrak{M}, 0 \models \text{diag}_N(\mathfrak{F}_i, 0)$, then $\text{Log}(\mathfrak{F}_i, 0) \supseteq \text{Log}(\mathfrak{F}_j, 0)$ and so $i = j$.

(2) If $\mathfrak{M}, 0 \models \text{diag}_N(\mathfrak{F}_i, 0)$ for a model $\mathfrak{M}$ based on $\mathfrak{F}_i$, then for every w in $\mathfrak{F}_i$ there is exactly one $f(w)$ in $\mathfrak{F}_i$ with $\mathfrak{M}, f(w) \models q_w$. f is an automorphism on $(\mathfrak{F}_i, 0)$.

Now assume that L has CIP. To show (2), assume a signature σ and $\mathfrak{M}, 0 \sim_\sigma \mathfrak{N}, 0$ with $\mathfrak{M}$ a model based on $\mathfrak{F}_i$ and $\mathfrak{N}$ a model based on $\mathfrak{F}_j$ are given. We take disjoint sets of atoms $q_{w,1}$, $w \in \mathfrak{F}_i$, and $q_{w,2}$, w in $\mathfrak{F}_j$. Denote by $\text{diag}^1_N(\mathfrak{F}_i, 0)$ the diagram of $(\mathfrak{F}_i, 0)$ using the atoms $q_{w,1}$ and by $\text{diag}^2_N(\mathfrak{F}_j, 0)$ the diagram of $(\mathfrak{F}_j, 0)$ using the atoms $q_{w,2}$. Note that $i = j$

is allowed. We also assume that the atoms of the diagrams are disjoint from σ. We now define descriptions of $\mathfrak{M}, 0$ and $\mathfrak{N}, 0$ using the corresponding diagrams: Let

$$\delta(\mathfrak{M}, 0) = \mathrm{diag}(\mathfrak{F}_i, 0) \wedge \Box^{\leq N}(\bigwedge_{\mathfrak{M},w\models p, p\in\sigma} (q_{w,1} \to p) \wedge \bigwedge_{\mathfrak{M},w\models \neg p, p\in\sigma} (q_{w,1} \to \neg p))$$

and

$$\delta(\mathfrak{N}, 0) = \mathrm{diag}(\mathfrak{F}_i, 0) \wedge \Box^{\leq N}(\bigwedge_{\mathfrak{N},w\models p, p\in\sigma} (q_{w,2} \to p) \wedge \bigwedge_{\mathfrak{N},w\models \neg p, p\in\sigma} (q_{w,2} \to \neg p))$$

It should be clear that $\mathfrak{M}, 0 \models \delta(\mathfrak{M}, 0)$ and $\mathfrak{N}, 0 \models \delta(\mathfrak{N}, 0)$. Since L has CIP we can apply Fact 1 because $\sigma = \mathrm{sig}(\delta(\mathfrak{M}, 0)) \cap \mathrm{sig}(\delta(\mathfrak{N}, 0))$. Hence there is a model $\mathfrak{M}'$ based on a frame $\mathfrak{F}_k$ in $\mathcal{F}_L$ such that $\mathfrak{M}' \models \delta(\mathfrak{M}, 0) \wedge \delta(\mathfrak{N}, 0)$. By Fact 2 (1), $i = j = k$.

By Fact 2 (2), we have automorphisms $f_1 : (\mathfrak{F}_i, 0) \to (\mathfrak{F}_i, 0)$ and $f_2 : (\mathfrak{F}_i, 0) \to (\mathfrak{F}_i, 0)$ with $\mathfrak{M}', f_1(w) \models q_{w,1}$ and $\mathfrak{M}', f_2(w) \models q_{w,2}$ for all w in $\mathfrak{F}_i$. f_1 and f_2 are invariant for σ-atoms: for all w in $\mathfrak{F}_i$ and $p \in \sigma$:

- $\mathfrak{M}, w \models p$ iff $\mathfrak{M}', f_1(w) \models p$;
- $\mathfrak{N}, w \models p$ iff $\mathfrak{M}', f_2(w) \models p$.

To see this for $\mathfrak{M}$, let w in $\mathfrak{F}_i$ and $p \in \sigma$. Observe that $\mathfrak{M}, w \models p$ iff $q_{w,1} \to p$ is a conjunct of $\delta(\mathfrak{M}, 0)$. We also have $\mathfrak{M}', f_1(w) \models q_{w,1}$. Hence, if $\mathfrak{M}, w \models p$, then $\mathfrak{M}', f_1(w) \models p$ since $\mathfrak{M}', 0 \models \delta(\mathfrak{M}, 0)$. The same argument applies to $\neg p$ and so the equivalence for $\mathfrak{M}$ follows. It follows that $f_1 \circ f_2^{-1}$ is an isomorphism between the σ-reducts of $(\mathfrak{M}, 0)$ and $(\mathfrak{N}, 0)$, as required. ◀

▶ **Theorem 1.3.** *Let L be a tabular quasi-normal modal logic that has CIP. Given a valid modal implication $\varphi \to \psi$ in L, and given a propositional Craig interpolant χ for $tr_L(\varphi) \to tr_L(\psi)$, it holds that $rt_L(\chi)$ is a Craig interpolant for $\varphi \to \psi$ with respect L.*

Proof. The proof is similar to the proof of Theorem 1.1 and uses Theorem 5.3. Fix $L = \mathrm{Log}(\mathcal{F}_L)$ with CIP and for a finite set $\mathcal{F}_L$ of finite rooted frames. We make the same assumptions as in the proof of Theorem 1.1. For any signature σ, define $\hat{\sigma}$ as before. Recall

$$tr_L(\varphi) = \mathrm{Unique}(\mathcal{F}_L)) \wedge \bigwedge_{(\mathfrak{F},w)\in\mathcal{F}_L} (r_{\mathfrak{F},w} \to tr_{\mathfrak{F},w}(\varphi)).$$

Now we show that the formulas $\xi^{\uparrow\mathfrak{F},w}$ introduced in the proof of Theorem 1.1 for $\sigma = \mathrm{sig}(\varphi) \cap \mathrm{sig}(\psi)$ are interpolants for $tr_{\mathfrak{F},w}(\varphi) \to tr_{\mathfrak{F},w}(\psi)$ for $(\mathfrak{F}, w) \in \mathcal{F}_L$ if ξ is an interpolant for $tr_L(\varphi) \to tr_L(\psi)$.

▷ Claim D.1. Let $\varphi \to \psi$ be a modal implication and ξ an interpolant for $tr_L(\varphi) \to tr_L(\psi)$ in $\sigma = \mathrm{sig}(\varphi) \cap \mathrm{sig}(\psi)$. Then $\xi^{\uparrow\mathfrak{F},w}$ is an interpolant for $tr_{\mathfrak{F},w}(\varphi) \to tr_{\mathfrak{F},w}(\psi)$, for every $(\mathfrak{F}, w) \in \mathcal{F}_L$.

The proof of Claim D.1 is similar to the proof of Claim B.4. One can first show in the same way that $tr_{\mathfrak{F},w}(\varphi) \to \xi^{\uparrow\mathfrak{F},w}$ is tautology, for all $(\mathfrak{F}, w) \in \mathcal{F}_L$. To show that $\xi^{\uparrow\mathfrak{F},w} \to tr_{\mathfrak{F},w}(\psi)$ is a tautology, recall that every model of $\xi^{\uparrow\mathfrak{F},w}$ can be expanded to a model of ξ in the obvious way. This model then satisfies $tr_L(\psi)$ since ξ is an interpolant. But then it satisfies $tr_{\mathfrak{F},w}(\psi)$, as required.

It follows from Theorem 5.3 and Claim D.1 that rt_L is as required. ◀

E Proofs for Section 6

▶ **Theorem E.1.** *Let L be any normal modal logic that contained in a pre-tabular normal modal logic containing S4 or GL. There are modal formulas $(\varphi_n)_{n=1,2,\dots}$ and $(\psi_n)_{n=1,2,\dots}$ of size polynomial in n such that $\varphi \to \psi \in L$ such that a Craig interpolant for $\varphi \to \psi$ exists, and every Craig interpolant for $\varphi \to \psi$ in L has size at least 2^n.*

Proof. The proof requires slightly different arguments for each pre-tabular normal modal logic. Consider $\sigma_n = \{r\} \cup \{p_0, \dots, p_{n-1}\}$ and let T_n denote the set of types $\bigwedge_{i=0}^{n-1} l_i$ with $l_i \in \{p_i, \neg p_i\}$ for $i = 0, \dots, n-1$. Let $q_0, \dots, q_{n-1}, q'_0, \dots, q'_{n-1}$ be a set of atoms disjoint from σ_n. We aim to enforce the Craig interpolants $I_n = \bigvee_{t \in T_n}(r \wedge t \wedge \Diamond(\neg r \wedge t)$. To this end, define

$$\chi(q,p) = \bigwedge_{0 \le i < n} \big((p_i \leftrightarrow q_i) \wedge (q_i \to \Box q_i) \wedge (\neg q_i \to \Box \neg q_i)\big)$$

and obtain $\chi(q', p)$ from $\chi(q,p)$ by replacing q_i by q'_i, for all $i < n$. Then let

$$\varphi_n = r \wedge \chi(p,q) \wedge \Diamond(\neg r \wedge \bigwedge_{0 \le i < n} (p_i \leftrightarrow q_i))$$

and

$$\psi_n = r \wedge (\chi(q',p) \to \Diamond(\neg r \wedge \bigwedge_{0 \le i < n} (p_i \leftrightarrow q'_i)))$$

Observe that for every pointed frame $\mathfrak{F}, 0$ we have $\mathfrak{F}, 0 \models \varphi_n \to \psi_n$ and I_n is a Craig interpolant for $\varphi_n \to \psi_n$ for every quasi-normal modal logic. The following observation follows directly from the fact that every model of I_n can be expanded to a model of φ_n and every model of $\neg I_n$ can be expanded to a model of $\neg \psi_n$.

Fact 1. If L is a quasi-normal modal logic, then every Craig interpolant for $\varphi_n \to \psi_n$ is equivalent to I_n in L.

We now argue that for every L contained in a normal pre-tabular logic containing S4 or GL, there is no formula equivalent to I_n on frames validating L of size $< 2^{\frac{1}{2}n}$, for sufficiently large n.

Given any transitive model $\mathfrak{M} = (\mathfrak{F}, R), V)$ without infinite strictly ascending R-chains, a world $w \in W$ is called *maximal for a formula χ in $\mathfrak{M}$* if $\mathfrak{M}, w \models \chi$ and for any w' with $\mathfrak{M}, w' \models \chi$ and wRw' it follows that $w' = w$ or $w'Rw$.

We begin with the pre-tabular normal modal logic containing S4.

Case 1. $L \subseteq Log(\mathcal{F}_{\mathrm{Grz.3}})$. Assume for a proof by contradiction that χ_n is equivalent to I_n in L and of size $< 2^{\frac{1}{2}n}$, for sufficiently large n.

We construct a model $\mathfrak{M}_{0,n}$ based on $\mathfrak{F}_{0,n} = (W_{0,n}, R_{0,n})$ as follows: let $W_{0,n} = M_{0,n} \cup N_{0,n}$ with $M_{0,n} = \{a_0, \dots, a_{2^n-1}\}$ and $N_{0,n} = \{b_0, \dots, b_{2^n-1}\}$. Define $R_{0,n}$ as the reflexive and transitive closure of

$$\{(a_i, b_0) \mid i < 2^n\} \cup \{(b_i, b_{i+1})\} \mid i < 2^n - 1\}.$$

Then $\mathfrak{F}_{0,n}$ validates L. Let $\mathfrak{M}_{0,n}$ be defined in such a way that

- every type $t \in T_n$ is satisfied in exactly one a_i, $i < 2^n$;
- r is satisfied in a_i for all $i < 2^n$ and not in any b_i with $i < 2^n$;
- every type $t \in T_n$ is satisfied in exactly one b_i. $i < 2^n$.

We have $\mathfrak{M}_{0,n}, a_i \models I_n$ for all $i < 2^n$. Hence we have $\mathfrak{M}_{0,n}, a_i \models \chi_n$. Pick for every $\xi \in \text{sub}(\chi_n)$ satisfied in some world in $N_{0,n}$ a maximal world b_ξ in $N_{0,n}$ satisfying ξ in $\mathfrak{M}_{0,n}$. Let $\mathfrak{M}_i$ denote the restriction of $\mathfrak{M}_{0,n}$ to a_i and the selected b_ξ. Then each pointed $\mathfrak{M}_i, c$ with c in $\mathfrak{M}_i$ satisfies the same type $t \in T_n$ as $\mathfrak{M}_{0,n}, c$. Moreover, $\mathfrak{M}_i, a_i \models \chi_n$ for all $i < 2^n$ can be shown by induction. Hence $\mathfrak{M}_i, a_i \models I_n$ since the $\mathfrak{M}_i, a_i$ are based on frames validating L. However, note that there is a type in T_n that is not satisfied in any b_ξ since $\mathfrak{M}_i$ is of size $< 2^n$. Hence there is an a_i with $\mathfrak{M}_i, a_i \not\models I_n$ and we have derived a contradiction.

Case 2. $L \subseteq \mathit{Log}(\mathcal{F}_{\mathrm{S5}})$. Assume for a proof by contradiction that χ_n is equivalent to I_n in L and of size $< 2^{\frac{1}{2}n}$, for sufficiently large n.

We construct a model $\mathfrak{M}_{1,n}$ based on $\mathfrak{F}_{1,n} = (W_{1,n}, R_{1,n})$ as follows. Define $W_{1,n} = M_{1,n} \cup N_{1,n}$ with $M_{1,n} = \{a_0, \ldots, a_{2^n-1}\}$ and $N_{1,n} = \{b_0, \ldots, b_{2^n-1}\}$ and let $R_{1,n} = W_{1,n} \times W_{1,n}$. Then $\mathfrak{F}_{1,n}$ validates L. Let the valuation of $\mathfrak{M}_{1,n}$ be defined in the same way as the valuation of $\mathfrak{M}_{0,n}$.

We have $\mathfrak{M}_{1,n}, a_i \models I_n$ for all $i < 2^n$. Hence we have $\mathfrak{M}_{1,n}, a_i \models \chi_n$. Pick for every $\xi \in \text{sub}(\chi_n)$ satisfied in some world in $N_{1,n}$ a world b_ξ in $N_{1,n}$ satisfying ξ in $\mathfrak{M}_{1,n}$. Also pick for every $\xi \in \text{sub}(\chi_n)$ satisfied in some world in $M_{1,n}$ a world a_ξ in $M_{1,n}$ satisfying ξ in $\mathfrak{M}_{1,n}$. Let $\mathfrak{M}_i$ denote the restriction of $\mathfrak{M}_{1,n}$ to a_i and the selected a_ξ and b_ξ. Now the proof continues as in Case 1: Clearly each $\mathfrak{M}_i, c$ with c in $\mathfrak{M}_i$ satisfies the same type $t \in T_n$ as $\mathfrak{M}_{1,n}, c$ and it is easy to see that $\mathfrak{M}_i, a_i \models \chi_n$ for all $i < 2^n$. Hence $M_i, a_i \models I_n$ for $i < 2^n$. However, there is a type in T_n that is not satisfied in any b_ξ since $\mathfrak{M}_i$ is of size $< 2^n$. Hence there is an a_i with $\mathfrak{M}_i, a_i \not\models I_n$ and we have derived a contradiction.

Case 3. $L \subseteq \mathit{Log}(\mathcal{F}_1)$. Assume for a proof by contradiction that χ_n is equivalent to I_n in L and of size $< 2^{\frac{1}{2}n}$, for all sufficiently large n.

We construct a model $\mathfrak{M}_{2,n}$ based on $\mathfrak{F}_{2,n} = (W_{2,n}, R_{2,n})$ by adding to the frame $\mathfrak{F}_{1,n}$ a world d with $a_i R d$ and $b_i R d$ for all a_i, b_i. Also set dRd. Then $\mathfrak{F}_{2,n}$ validates L. Let $\mathfrak{M}_{2,n}$ be defined as in Case 2 and let d satisfy any type in T_n. We have $\mathfrak{M}_{2,n}, a_i \models I_n$ for all $i < 2^n$. Hence we have $\mathfrak{M}_{2,n}, a_i \models \chi_n$. Pick for every $\xi \in \text{sub}(\chi_n)$ satisfied in some world in $N_{1,n}$ a world b_ξ in $N_{1,n}$ satisfying ξ in $\mathfrak{M}_{2,n}$. Also pick for every $\xi \in \text{sub}(\chi_n)$ satisfied in some world in $M_{1,n}$ a world a_ξ in $M_{1,n}$ satisfying ξ in $\mathfrak{M}_{2,n}$. Let $\mathfrak{M}_i$ denote the restriction of $\mathfrak{M}_{2,n}$ to a_i, d, and the selected a_ξ and b_ξ. Now the proof continues similar to Case 1 and 2: clearly each $\mathfrak{M}_i, c$ with c in $\mathfrak{M}_i$ satisfies the same type $t \in T_n$ as $\mathfrak{M}_{1,n}, c$ and it is easy to see that $\mathfrak{M}_i, a_i \models \chi_n$ for all $i < 2^n$. One can proceed in the same way as in the proof of Cases 1 and 2 and derive a contradiction.

Case 4. $L \subseteq \mathit{Log}(\mathcal{F}_2)$. Assume for a proof by contradiction that χ_n is equivalent to I_n in L and of size $< 2^{\frac{1}{2}n}$, for sufficiently large n.

We construct a model $\mathfrak{M}_{3,n}$ based on $\mathfrak{F}_{3,n} = (W_{3,n}, R_{3,n})$ as follows. Define $W_{1,n} = M_{1,n} \cup N_{1,n}$ with $M_{3,n} = \{a_0, \ldots, a_{2^n-1}\}$ and $N_{3,n} = \{b_0, \ldots, b_{2^n-1}\}$ and let $R_{1,n}$ be the reflexive closure of $M_{3,n} \times N_{3,n}$. Then $\mathfrak{F}_{3,n}$ validates L. Let the valuation of $\mathfrak{M}_{3,n}$ be defined in the same way as the valuation of $\mathfrak{M}_{0,n}$. We have $\mathfrak{M}_{3,n}, a_i \models I_n$ for all $i < 2^n$. Hence we have $\mathfrak{M}_{3,n}, a_i \models \chi_n$. Now one can proceed similar to the previous cases and derive a contradiction.

Case 5. $L \subseteq \mathit{Log}(\mathcal{F}_3)$. Assume for a proof by contradiction that χ_n is equivalent to I_n in L and of size $< 2^{\frac{1}{2}n}$, for sufficiently large n.

We construct a model $\mathfrak{M}_{4,n}$ based on $\mathfrak{F}_{4,n}$ by adding to $\mathfrak{M}_{3,n}$ a world d with $a_i R d$ and $b_i R d$ for all a_i, b_i. Also set dRd. Then $\mathfrak{F}_{4,n}$ validates L. Let $\mathfrak{M}_{4,n}$ be defined as in Case 4 and let d satisfy any type in T_n. We have $\mathfrak{M}_{4,n}, a_i \models I_n$ for all $i < 2^n$. Hence we have $\mathfrak{M}_{4,n}, a_i \models \chi_n$. Now one can proceed similar to the previous cases and derive a contradiction.

We continue with the pre-tabular normal modal logics containing GL.

Case 6. $L \subseteq Log(\mathcal{F}_{\mathrm{GL.3}})$. Assume for a proof by contradiction that χ_n is equivalent to I_n in L and of size $< 2^{\frac{1}{2}n}$, for sufficiently large n.

This case is essentially the same as Case 1 except that we have irreflexive instead of reflexive worlds.

Case 7. $L \subseteq Log(\mathcal{F}_{n_1,n_2})$. Assume for a proof by contradiction that χ_n is equivalent to I_n in L and of size $< 2^{\frac{1}{2}n}$, for sufficiently large n. This case is again similar to previous cases.

Fix $n_1 \geq 1$ and $n_2 \geq 0$ and let $n > n_1 + n_2$. We construct a model $\mathfrak{M}_{5,n}$ based on $\mathfrak{F}_{5,n} = (W_{5,n}, R_{5,n})$ as follows: let $W_{5,n} = M_{5,n} \cup N_{5,n} \cup P_{5,n}$ with $M_{5,n} = \{a_0, \ldots, a_{2^n-1}\}$, $N_{5,n} = \{b_0, \ldots, b_{2^n-1}\}$ and $P_{5,n} = \{d_1, \ldots d_{n_1}, d_{n_1+1}, \ldots, d_{n_1+n_2+1}\}$. Define $R_{5,n}$ as the transitive closure of

$$(M_{5,n} \times (N_{5,n} \cup P_{5,n})) \cup \{(d_i, d_{i+1}) \mid 1 \leq i < n_1 + n_2\} \cup (N_{5,n} \times \{d_{n_1+1}\})$$

Then $\mathfrak{F}_{5,n}$ validates L. Let $\mathfrak{M}_{5,n}$ be defined in such a way that

- every type $t \in T_n$ is satisfied in exactly one a_i, $i < 2^n$;
- r is satisfied in a_i for all $i < 2^n$ and not in any other world;
- every type $t \in T_n$ is satisfied in exactly one b_i, $i < 2^n$;
- the worlds in $P_{5,n}$ all satisfy one fixed type t_0.

We have $\mathfrak{M}_{5,n}, a_i \models I_n$ for all $i < 2^n$. Hence we have $\mathfrak{M}_{5,n}, a_i \models \chi_n$. Pick for every $\xi \in \mathrm{sub}(\chi_n)$ satisfied in some world in $N_{0,n}$ a world b_ξ in $N_{0,n}$ satisfying ξ in $\mathfrak{M}_{5,n}$. Let $\mathfrak{M}_i$ denote the restriction of $\mathfrak{M}_{5,n}$ to a_i, $P_{5,n}$, and the selected b_ξ. Then each pointed $\mathfrak{M}_i, c$ with c in $\mathfrak{M}_i$ satisfies the same type $t \in T_n$ as $\mathfrak{M}_{5,n}, c$. Moreover, $\mathfrak{M}_i, a_i \models \chi_n$ for all $i < 2^n$ can be shown by induction. Hence $\mathfrak{M}_i, a_i \models I_n$ since the $\mathfrak{M}_i, a_i$ are based on frames validating L. However, note that there is a type in T_n that is not satisfied in any b_ξ nor in $P_{5,n}$ since $\mathfrak{M}_i$ is of size $< 2^n$. Hence there is an a_i with $\mathfrak{M}_i, a_i \not\models I_n$ and we have derived a contradiction. ◀

▶ **Theorem 6.4.** *Let σ be a signature, and $\widehat{\sigma} = \sigma \times \mathbb{N}$. There are poly-time translations*

$$tr_{alt} : ML(\sigma) \to PL(\widehat{\sigma}) \quad \textit{and} \quad rt_{alt} : PL(\widehat{\sigma}) \to ML(\sigma)$$

such that for every modal formula φ, every $\tau \subseteq \sigma$ and every uniform $\tau \times \mathbb{N}$-interpolant ψ for $tr_{alt}(\varphi)$ in propositional logic, $rt_{alt}(\psi)$ is a uniform $Alt_1(\tau)$-interpolant for φ.

Here, if $\varphi \in ML(\sigma)$ has modal depth n, then $tr_{alt}(\varphi) \in PL(\sigma \times \{0, \ldots, n\})$.

Proof. Let $\varphi \in \mathrm{ML}(\sigma)$ have modal depth n and assume $\tau \subseteq \sigma$. The translation $tr_{\mathrm{alt}}(\varphi)$ will only be concerned with the frames $\mathcal{F}_n = \{\mathfrak{F}_0, \cdots, \mathfrak{F}_n\}$ where $\mathfrak{F}_k = (W_k, R_k)$ is given by $W_k = \{k, \cdots, 0\}$, $R_k = \{(i, i-1) \mid 0 < i \leq k\}$ and the root of $\mathfrak{F}_i$ is the world i (so here we do not follow the convention that the root is labeled 0). In other words, $\mathcal{F}_n$ is the set of linear chains of length up to n, with the worlds counting down from k. We can use the technique developed in Section 4 to find a strongest $\mathrm{Alt}_1(\tau)$-implicate which, because Alt_1 has CIP, will also be a uniform $\mathrm{Alt}_1(\tau)$-interpolant. Naively applying the method from Section 4 would result in an exponential sized interpolant, however, whereas we are looking for poly-time translations to and from propositional logic. For this purpose, we consider simpler translations that take advantage of the specific structure of $\mathcal{F}_n$.

Let $\mathfrak{F}_i \in \mathcal{F}_n$, and let $\mathfrak{M}$ be any model based on $\mathfrak{F}_i$. We can define the associated propositional model $v_{\mathfrak{M}}$ over $\sigma \times \{0, \cdots, i\}$ by taking $v_{\mathfrak{M}}(p_j) = 1$ iff $\mathfrak{M}, j \models p$. Additionally, we recursively define a set of translation $tr_i : \mathrm{ML}(\sigma) \to \mathrm{PL}(\widehat{\sigma})$, for $i \in \mathbb{N}$, by

$$\begin{aligned} tr_i(p) &= p_i, \\ tr(\neg\varphi) &= \neg tr_i(\varphi), \\ tr_i(\varphi_1 \vee \varphi_2) &= tr_i(\varphi_1) \vee tr_i(\varphi_2), \\ tr_i(\Box\varphi) &= \begin{cases} \top & \text{if } i = 0 \\ tr_{i-1}(\varphi) & \text{if } i > 0 \end{cases} \end{aligned}$$

and a reverse translation $rt_i : \mathrm{PL}(\sigma \times \{0, \cdots, i\}) \to \mathrm{ML}(\sigma)$ by

$$rt_i(\psi) = \psi[p_j \mapsto \Diamond^{i-j} p].$$

It is straightforward to verify that for every $\chi \in \mathrm{ML}(\sigma)$, $v_{\mathfrak{M}} \models tr_i(\chi)$ if, and only if, $\mathfrak{M}, i \models \chi$, and that for every propositional formula $\psi \in \mathrm{PL}(\sigma \times \{0, \cdots, i\})$, $v_{\mathfrak{M}} \models \psi$ if, and only if, $\mathfrak{M}, i \models rt_i(\psi)$.

We next define a strongest $\mathrm{Alt}_1(\tau)$-implicate for φ using uniform propositional $\tau \times \{0, \cdots, i\}$-interpolants for $tr_i(\varphi)$, for $0 \leq i \leq n$. Let, for each $0 \leq i \leq n$, ψ_i be a uniform propositional $\tau \times \{0, \cdots, i\}$-interpolant for $tr_i(\varphi)$. Consider the modal formula

$$\psi = \bigwedge_{0 \leq i \leq n-1} (\Diamond^i \Box \bot \to \psi_i[p_j \mapsto \Diamond^{i-j} p]) \wedge (\Diamond^n \top \to \psi_n[p_j \mapsto \Diamond^{n-j} p])$$

which uses propositional variables in τ only. It is straightforward to show that ψ is a strongest $\mathrm{Alt}_1(\tau)$-implicate for φ, as required.

We have used multiple translations tr_i and rt_i and also multiple uniform propositional interpolants, as was previously done for tabular logics determined by sets of finite frames in Section 4.2. A single translation tr_{alt} and single converse translation rt_{alt} are obtained in the same way as in Section 4.3. ◀

F Proofs and Refutations of CIP

We show claims in Figure 1 for which we have not found direct proofs in the literature.

▶ **Theorem F.1.** *Alt_1 and Alt_2 have UIP, whereas Alt_n for $n \geq 3$ does not have CIP.*

Proof. CIP for Alt_1 is shown in [9]. CIP for Alt_2 was proved in [14] (where this logic is referred to as the logic of $\mathbb{T}_2$). CIP for Alt_n, $n \geq 3$ can be refuted using the same argument as in Example 3.2. The results for UIP follow from the fact that a normal modal logic with CIP axiomatised by modal formulas of modal depth at most one also has UIP [32]. ◀

▶ **Theorem F.2.** *EQ_1 and EQ_2 have UIP, whereas EQ_n for $n \geq 3$ does not have CIP. The logics LO_n for $n \geq 1$ have UIP.*

Proof. It is easy to see that, for $n \leq 2$, whenever two pointed EQ_n-models are bisimilar, they are isomorphic. Likewise for LO_n ($n \geq 1$). It follows by Theorem 5.1 that these logics have CIP, and hence UIP (recall that every tabular quasi-normal modal logic with CIP has UIP). Thus, these logics have UIP.

On the other hand, let $n \geq 3$ and consider the EQ_n-models $M = (\{0, 1, 2\}, \{0, 1, 2\}^2, V)$ and $N = (\{0, 1, 2\}, \{0, 1, 2\}^2, V')$ with $V(p) = \{1\}$ and $V'(p) = \{1, 2\}$. The pointed models $M, 0$ and $N, 0$ are bisimilar but not isomorphic. It follows by Theorem 5.1 that EQ_n lacks CIP, and hence, lacks UIP. ◀