Effective Interpolation and Preservation in Guarded Logics *

Michael Benedikt
University of Oxford
michael.benedikt@cs.ox.ac.uk

Balder ten Cate
LogicBlox and UC Santa Cruz
btencate@ucsc.edu

Michael Vanden Boom
University of Oxford
michael.vandenboom@cs.ox.ac.uk

Abstract

Desirable properties of a logic include decidability, and a model theory that inherits properties of first-order logic, such as interpolation and preservation theorems. It is known that the Guarded Fragment (GF) of first-order logic is decidable and satisfies some preservation properties from first-order model theory; however, it fails to have Craig interpolation. The Guarded Negation Fragment (GNF), a recently-defined extension, is known to be decidable and to have Craig interpolation. Here we give the first results on effective interpolation for extensions of GF. We provide an interpolation procedure for GNF whose complexity matches the doubly exponential upper bound for satisfiability of GNF. We show that the same construction gives not only Craig interpolation, but Lyndon interpolation and Relativized interpolation, which can be used to provide effective proofs of some preservation theorems. We provide upper bounds on the size of GNF interpolants for both GNF and GF input, and complement this with matching lower bounds.

Categories and Subject Descriptors Theory of computation [*Logic*]

1. Introduction

Basic modal logic is known to have many attractive properties, both computationally and in its model theory. Given the limited expressiveness of basic modal logic, logicians have searched for more powerful languages that extend it while preserving its positive features.

Andréka et. al. ([2], 4.7) list "meta-properties" of modal logic that one would desire in an extension, including the following.

- The logic should have decidable satisfiability and validity problems.
- The logic should satisfy *interpolation theorems*, like the *Craig Interpolation Property* (CIP): if φ and ψ are in the logic, and φ entails ψ, then there should be a formula θ in the logic such that φ entails θ and θ entails ψ, and θ mentions only relations common to both φ and ψ.
- The logic should satisfy *preservation theorems* such as the Łoś-Tarski theorem: if a sentence in the logic is preserved under superstructures, it should be equivalent to an existential sentence in the logic.
- The logic should have the *finite model property*: a satisfiable sentence should have a finite model.

Interpolation, in particular, is a desirable property since it implies the well-known *Beth Definability Property* (BDP), as well as its extension, the *Projective Beth Definability Property* (PBDP), stating that implicit specifications can be converted to explicit ones. This conversion from implicit to explicit is important in knowledge representation [24] and databases [20]. Interpolation has many applications, both in simplifying definitions and in verification [18].

Andréka et. al. proposed the *Guarded Fragment* (GF) as a candidate for an attractive logic. They showed that it is decidable, that it has the finite model property, and that it satisfies Łoś-Tarski. Unfortunately Hoogland, Marx, and Otto [14] showed that GF does not satisfy CIP. Later work [7] showed that PBDP fails for GF.

The *Guarded Negation Fragment* (GNF) is an extension of GF introduced by Bárány, ten Cate and Segoufin in [5]. It also extends the *Unary Negation Fragment* (UNF) introduced by ten Cate and Segoufin in [23]. Both UNF and GNF were proven decidable and shown to have the finite model property [5, 23]. They were also shown to satisfy many of the properties mentioned earlier: in [23], UNF was shown to have CIP and PBDP; in [7], GNF was shown to have CIP, PBDP, and several preservation properties including the analog of Łoś-Tarski. The combination of interpolation and the finite model property for these logics means that interpolation holds when validity is considered over finite models only. This is attractive for applications of these logics in databases, such as rewriting queries using views [20], since in databases typically only finite models are considered.

In this work we provide further evidence that GNF and UNF are attractive extensions of modal logic. We show that these logics satisfy the Lyndon interpolation property (LIP) and its generalization, the Relativized interpolation property (RIP), and moreover present the first constructive procedures for GNF and UNF interpolation, including CIP, LIP, and RIP. Roughly speaking, the LIP says that if φ entails ψ, we can find θ as in CIP, but with a relation symbol occurring positively in θ iff it occurs positively in both φ and ψ. The RIP says that the quantification pattern in θ must match the quantification pattern in φ and ψ.

We proceed by extending the *mosaic method* that has been applied in providing decidability for algebras of relations [21], modal logics [19], and temporal logics [17]. We present a variant of the method that gives a sound and complete decision procedure for the validity problem for GNF, but also allows interpolants to be read off from a proof of an entailment.

We show how this algorithm gives interpolants that witness the RIP and hence the LIP and CIP for GNF and UNF. Since the methods of [7, 23] were model-theoretic, this gives the first constructive proof of interpolation for GNF and UNF. We deduce

* Benedikt and Vanden Boom are supported by EPSRC grant EP/H017690/1 (The Engineering and Physical Sciences Research Council, UK), and ten Cate is supported by NSF grant IIS-1217869.

bounds on the size of interpolants for GNF, and also on the size of GNF interpolants for GF formulas. We use the RIP to prove that GNF and GF admit effective preservation theorems, which state that formulas satisfying certain semantic properties can be converted to specific forms. To the best of our knowledge, these are the first effective preservation theorems, even for GF.

We complement the upper bounds with lower bounds for the size of GNF interpolants for GF and UNF.

Due to space limitations, most proofs are deferred to the full version of this paper.

2. Preliminaries

Notation and conventions. We use $\boldsymbol{x}, \boldsymbol{y}, \ldots$ to denote vectors of variables. For a formula ϕ, we write $\phi(\boldsymbol{x})$ to indicate that the free variables in ϕ are contained in $\boldsymbol{x}$, and write $\mathrm{free}(\phi)$ for the actual free variables. The formula ϕ may also use constants $\mathrm{con}(\phi)$. We use $\alpha, \beta, \ldots$ to denote atomic formulas, and if we write $\alpha(\boldsymbol{x})$ then we assume that $\mathrm{free}(\alpha) = \boldsymbol{x}$.

Classical query languages. We will make use of some basic notions of database theory – in particular the following "classical" query classes: conjunctive queries (CQs), unions of conjunctive queries (UCQs), first-order logic formulas (FO), existential FO, and positive existential FO. Abiteboul, Hull, and Vianu [1] is a good reference for these languages.

The basics of GNF and GF. The Guarded Negation Fragment (GNF) is built up inductively according to the grammar:

$$\phi ::= R\,\boldsymbol{t} \mid \exists x.\phi \mid \phi \vee \phi \mid \phi \wedge \phi \mid \alpha(\boldsymbol{x}) \wedge \neg\phi(\boldsymbol{x})$$

where R is either a relation symbol or the equality relation $x = y$, α is an atomic relation (including equality), and $\boldsymbol{t}$ is a tuple over variables and constants. Notice that any use of negation must occur conjoined with an atomic relation that contains all the free variables of the negated formula. Such an atomic relation is a *guard* of the formula.

The purpose of allowing equalities as guards is to ensure that every formula with at most one free variable can always be guarded. Thus GNF includes the Unary Negation Fragment (UNF), which is built up as above, but allowing negation only on formulas with at most one free variable.

A GNF formula is in GN-*normal form* if in its syntax tree, no disjunction is directly below an existential quantifier or a conjunction, and no existential quantifier is directly below a conjunction sign. Such GN-normal form formulas ϕ can be generated by the following recursive definition:

$$\begin{aligned} \phi &::= \textstyle\bigvee_i \exists \boldsymbol{x}. \bigwedge_j \psi_{ij} \\ \psi &::= R\,\boldsymbol{t} \mid \alpha(\boldsymbol{x}) \wedge \neg\phi(\boldsymbol{x}) \end{aligned}$$

where R is either a relation symbol or the equality relation, α is an atomic relation (including equality), and $\boldsymbol{t}$ is a tuple over variables and constants. Thus, GN-normal form formulas can be viewed as formulas built up from atomic formulas using UCQs and guarded negation. Every GNF formula can be brought into GN-normal form, at the cost of an exponential increase in length and linear increase in the number of variables.

GNF should be compared to the *Guarded Fragment* (GF), typically defined via the grammar:

$$\phi ::= R\,\boldsymbol{t} \mid \exists \boldsymbol{x}.\alpha(\boldsymbol{xy}) \wedge \phi(\boldsymbol{xy}) \mid \phi \vee \phi \mid \phi \wedge \phi \mid \neg\phi(\boldsymbol{x})$$

where R is either a relation symbol or the equality relation, α is an atomic relation (including equality), and $\boldsymbol{t}$ is a tuple over variables and constants. Here it is the quantification that is guarded, rather than negation. As in GNF, we allow equality guards by default.

It is easy to see that every union of conjunctive queries is expressible in GNF. It is only slightly more difficult to verify that every GF sentence can be expressed in GNF [5].

Turning to fragments of first-order logic that are common in database theory, one can show that common classes of integrity constraints, such as inclusion dependencies (also known as referential constraints) are expressible in GNF, as well as many of the common dependencies used in data integration and exchange. Going further, the relational translations of many of the common description logic languages used in the semantic web (e.g. $\mathcal{ALC}$ and $\mathcal{ALCHIO}$ [4]) are known to be expressible in GF, and hence in GNF.

Not only is GNF an expressive fragment of FO, but it was shown to be decidable and to have the finite model property.

Theorem 1 ([5])**.** *A GNF sentence is satisfiable over all models iff it is satisfiable over finite models. Satisfiability can be tested effectively (and is* 2EXPTIME*-complete).*

Size of formulas. A formula ϕ given as input (to interpolation, preservation, etc.) is assumed to be in the standard tree representation of a formula, and the size of this representation is the number of symbols in ϕ.

For the output, however, we will usually represent the formula and measure its size as a node-labelled DAG (directed acyclic graph). The nodes represent formulas, and the edge relation connects a formula to its subformulas. We will also use DAGs with multiple roots to represent a set of formulas.

We often abuse notation by identifying a formula with a representation; given a representation of a formula ϕ, $|\phi|$ represents the size of the representation: the number of symbols if it is a formula, and the number of edges and labels if the representation is a DAG. Similarly, for a set S of formulas represented as a DAG, we write $|S|$ to denote the total number of nodes and edges in the DAG.

The *width* of ϕ, denoted $\mathrm{width}(\phi)$, is the maximum number of free variables of any subformula of ϕ. For GF sentences, the width is bounded by the maximum arity of any relation.

Interpolation. For two formulas φ_{L} and φ_{R}, we write $\varphi_{\mathrm{L}} \models \varphi_{\mathrm{R}}$ if every model of φ_{L} is a model of φ_{R}, and in this case say that $\varphi_{\mathrm{L}} \models \varphi_{\mathrm{R}}$ is a *validity*. An *interpolant* for a validity is a formula θ such that $\varphi_{\mathrm{L}} \models \theta \models \varphi_{\mathrm{R}}$ and θ mentions only relations present in both φ_{L} and φ_{R}. The Craig interpolation theorem [8] states that if φ_{L} and φ_{R} are in FO, then there is an FO interpolant θ.

Lyndon proved a stronger interpolation property by considering the polarity of relations in the formulas. Given a formula ϕ, we define $\mathrm{occ}(\phi)$ to be the set describing the *polarity of occurrences of atomic relations* (not including the equality relation). A relation occurs positively (respectively, negatively) in a formula if it is within the scope of an even (respectively, odd) number of negations. For every atomic relation S in ϕ (where S is not the equality relation), $(S, +) \in \mathrm{occ}(\phi)$ (respectively, $(S, -) \in \mathrm{occ}(\phi)$) iff S occurs positively (respectively, negatively) in ϕ. Lyndon's interpolation theorem [15] states that the interpolant θ can be taken so that $\mathrm{occ}(\theta) \subseteq \mathrm{occ}(\varphi_{\mathrm{L}}) \cap \mathrm{occ}(\varphi_{\mathrm{R}})$.

A more recent relativized interpolation theorem of Otto [22] states that if φ_{L} and φ_{R} both use only quantification *relativized to a set of unary predicates* $\mathbb{U}$, then the same is true of the interpolant. A relativized quantifier for $\mathbb{U}$ is $\exists \boldsymbol{x}.U_1 x_1 \wedge \cdots \wedge U_k x_k \wedge \psi(\boldsymbol{xy})$ or $\forall \boldsymbol{x}.U_1 x_1 \wedge \cdots \wedge U_k x_k \rightarrow \psi(\boldsymbol{xy})$ where $\boldsymbol{x} = x_1 \ldots x_k$ and $U_i \in \mathbb{U}$ for $i \in \{1, \ldots, k\}$. As shown in [22], the relativized interpolation theorem gives an alternative means of showing preservation theorems. We explain and expand upon this in Section 7.

Although there are proofs of Craig interpolation for FO that construct interpolants efficiently given proofs that witness the validity (e.g. [10]), these do not give bounds on the size of interpolants relative to the size of φ_{L} and φ_{R}. In fact, it has been shown that for

FO validities, there is no computable bound on the sizes of the interpolants relative to the sizes of the original validities [11].

We say that a logic has the Craig Interpolation Property (CIP) if given φ_{L} and φ_{R} in the logic, there is an interpolant that is again in the logic. Similarly we can talk about the Lyndon Interpolation Property (LIP) and Relativized Interpolation Property (RIP). Hoogland et. al [14] showed that the CIP fails for GF. Marx [16] argued that one of the applications of CIP, the Projective Beth Definability Property (PBDP) holds for GF as well as for an extension called the "Packed Fragment". He conjectured that the proof could be made effective using the "mosaic method". The argument in [16] was flawed, and indeed neither the guarded fragment nor the packed fragment has PBDP [7]. Nevertheless, here we use the mosaic technique suggested by Marx to show effective CIP, LIP, RIP, and PBDP for GNF (see Theorem 9).

3. The Mosaic Method

In its simplest form, a mosaic is just a finite set of formulas. When it satisfies certain coherence properties, it can be viewed as a collection of formulas describing part of a model. For instance, in the mosaic system for GF described in Section 4.6 of [2], a mosaic can describe the complete subformula type of some guarded set of elements in a model.

For modal and guarded logics, it turns out that the existence of a (finite) set of mosaics satisfying certain coherence properties is enough to ensure the existence of a full model. This allows mosaics to be used to decide satisfiability.

We adapt this method to check whether $\varphi_{\mathrm{L}} \wedge \neg\varphi_{\mathrm{R}}$ is unsatisfiable, and if so, extract an interpolant for $\varphi_{\mathrm{L}} \models \varphi_{\mathrm{R}}$ from the mosaics. The purpose of this section is to describe the general framework that we are using. We utilize this framework in later sections with different versions of mosaics depending, for instance, on whether the input formulas are in GF versus GNF, or whether we are aiming for Lyndon interpolation versus Relativized interpolation. As a result, some of the terms in this section are only defined informally; the formal definitions appear in later sections since they depend on the particular mosaic variant that we are using.

We start by fixing some sentences φ_{L} and φ_{R} over signatures σ_{L} and σ_{R}, respectively. The first step is to define a finite set Γ of formulas that are relevant to checking satisfiability of $\varphi_{\mathrm{L}} \wedge \neg\varphi_{\mathrm{R}}$. This is usually some extension of the subformula closure of φ_{L} and φ_{R}. To be suitable for computing an interpolant for $\varphi_{\mathrm{L}} \models \varphi_{\mathrm{R}}$, we also annotate these formulas with a *provenance* $X \in \{\mathrm{L}, \mathrm{R}\}$ indicating that the formulas are related to φ_X. This parallels other constructive interpolation methods, e.g. [10]. Given a provenance X, we write $\widetilde{X}$ for the *dual provenance* (i.e. $\widetilde{\mathrm{L}} = \mathrm{R}$ and $\widetilde{\mathrm{R}} = \mathrm{L}$). Throughout the paper, we use X, Y, Z to range over provenances L and R.

Given some finite set of *parameters* $\boldsymbol{c}$, we define $\Gamma(\boldsymbol{c})$ as the set of formulas from Γ with parameters from $\boldsymbol{c}$ substituted for any free variables. Parameters are just a special type of variable. We write $\boldsymbol{a}, \boldsymbol{b}, \ldots$ to denote vectors of parameters, and given a formula ϕ, we write $\phi(\boldsymbol{xb})$ to indicate that the free variables in ϕ are contained in $\boldsymbol{x}$ and the parameters in ϕ are contained in $\boldsymbol{b}$. We write $\mathrm{free}(\phi)$ (respectively, $\mathrm{par}(\phi)$) for the actual free variables (respectively, parameters). For atomic α, if we write $\alpha(\boldsymbol{xb})$ then we assume that $\mathrm{free}(\alpha) = \boldsymbol{x}$ and $\mathrm{par}(\alpha) = \boldsymbol{b}$.

A *mosaic* τ over some finite set of parameters $\boldsymbol{c}$ (of size bounded by the width of the original sentences) is then defined to be a subset of formulas from $\Gamma(\boldsymbol{c})$. Each mosaic is annotated with L or R, and an X-mosaic has additional restrictions on the $\widetilde{X}$-formulas appearing in it. This helps prove properties about the interpolants that we later extract from these mosaics.

Mosaics can be *linked* together. For instance, if a formula $Y : \exists \boldsymbol{x}.\psi(\boldsymbol{xb})$ appears in an X-mosaic τ, then this can be linked to some Y-mosaic τ' using parameters $\boldsymbol{ab}$ and containing the formula $Y : \psi(\boldsymbol{ab})$ (for some new parameters $\boldsymbol{a}$). In addition, the linked mosaic τ' must contain the maximum amount of information from τ about formulas using only shared parameters $\boldsymbol{b}$ (subject to any constraints on formulas in τ') and cannot contain any new information about formulas using only shared parameters. We say an existential requirement $Y : \exists \boldsymbol{x}.\psi(\boldsymbol{xb})$ in τ is *fulfilled* if there is a formula $Y : \psi(\boldsymbol{ab})$ in τ, or there is some τ' such that τ and τ' are linked via $Y : \exists \boldsymbol{x}.\psi(\boldsymbol{xb})$.

In order for a mosaic to correspond to a part of a model it must satisfy certain coherence properties. Within a single mosaic, this means that it must be *internally consistent*. For instance, we may want to require that if it contains a conjunction, then it contains both conjuncts. The exact definition will depend on the particular version of mosaics that we are considering. A set of mosaics is coherent if every mosaic is internally consistent and it is *saturated*: every existential requirement is fulfilled by a mosaic in the set.

We then fix some finite set $\mathcal{M}$ of mosaics that suffices to prove the following key *mosaic theorem*:

> $\varphi_{\mathrm{L}} \wedge \neg\varphi_{\mathrm{R}}$ is satisfiable if and only if there is a saturated set of internally consistent mosaics from $\mathcal{M}$ containing a mosaic with φ_{L} and $\neg\varphi_{\mathrm{R}}$.

In one direction, the idea is to extract a saturated set of internally consistent mosaics from a model. In the other (harder) direction, we must show that a saturated set of internally consistent mosaics from $\mathcal{M}$ can be used to construct a model of $\varphi_{\mathrm{L}} \wedge \neg\varphi_{\mathrm{R}}$.

This gives us a procedure for deciding whether or not $\varphi_{\mathrm{L}} \wedge \neg\varphi_{\mathrm{R}}$ is satisfiable. Consider the following *mosaic elimination procedure* over $\mathcal{M}$.

- At stage $i = 1$, every internally inconsistent mosaic is eliminated.
- At stage $i > 1$, a mosaic τ is eliminated if it has an existential requirement that is not fulfilled in τ, and every linked τ' fulfilling it was already eliminated.

This process continues until a fixpoint is reached. The mosaic theorem can then be used to show that $\varphi_{\mathrm{L}} \wedge \neg\varphi_{\mathrm{R}}$ is satisfiable iff there is a mosaic containing φ_{L} and $\neg\varphi_{\mathrm{R}}$ that is never eliminated.

For interpolation, we are interested in the case when $\varphi_{\mathrm{L}} \wedge \neg\varphi_{\mathrm{R}}$ is unsatisfiable, since this means $\varphi_{\mathrm{L}} \models \varphi_{\mathrm{R}}$.

It turns out that we can construct a *mosaic interpolant* θ^{τ}_X for each X-mosaic τ eliminated using the procedure above. These mosaic interpolants satisfy the property that $\tau_X \models \theta^{\tau}_X$ and $\theta^{\tau}_X \models \neg\tau_{\widetilde{X}}$, where τ_Y denotes the conjunction of all of the Y-formulas from $\Gamma(\boldsymbol{c})$ in τ and the negation of all of the Y-formulas from $\Gamma(\boldsymbol{c})$ that satisfy the constraints on formulas in τ but are not in τ.

We start by constructing these interpolants for mosaics eliminated because of internal inconsistencies. The interpolants for such mosaics are simple: atomic formulas or the negation of atomic formulas. Proceeding by induction on the stage at which the mosaic is eliminated, we then build up more complicated interpolants for the other mosaics. Intuitively, each mosaic interpolant describes the reason the mosaic was eliminated.

Since the sequence of elimination steps for the mosaics provide a proof that $\varphi_{\mathrm{L}} \wedge \neg\varphi_{\mathrm{R}}$ is unsatisfiable, it is not surprising that the desired interpolant for $\varphi_{\mathrm{L}} \models \varphi_{\mathrm{R}}$ can then be constructed from these mosaic interpolants.

It is helpful to compare this to other constructive interpolation approaches. In [10], an interpolant is constructed from a proof tree for $\varphi_{\mathrm{L}} \wedge \neg\varphi_{\mathrm{R}}$ in a tableau proof system for FO. The construction starts at the leaves (where very simple interpolants suffice), and builds up to an interpolant that holds for the original validity at the root of the proof tree. Because of the "one-step" nature of the proof rules, there are straightforward interpolation construction rules at each stage.

Using the mosaic method, the definition of the interpolant at the inductive stages is more complicated. Unlike a traditional proof system, however, the mosaics allow precise control over the parameters appearing at each stage. This is advantageous for constructing an interpolant in GNF, since we need to be able to guard certain parameters in order to remain in GNF. Indeed, this precise control over the parameters was one of the motivations for building GNF interpolants from mosaics rather than from a more traditional proof system.

4. GF Mosaics

We start by describing the mosaic method for validity testing and interpolation construction when the input consists of GF sentences φ_{L} and φ_{R} over signatures σ_{L} and σ_{R}, respectively. The purpose of this section is to give a more gentle introduction to mosaics, without the complications that come from handling GNF, relativized quantifiers, or equality.

We assume the sentences φ_{L} and $\neg\varphi_{\mathrm{R}}$ are in GF and are in negation normal form (negations are pushed to the inside) and do not use equality. These sentences may use constants, and we let $\boldsymbol{e} := \mathrm{con}(\varphi_{\mathrm{L}}) \cup \mathrm{con}(\varphi_{\mathrm{R}})$.

Closure. For GF satisfiability, we restrict formulas in mosaics to the subformula closure of φ_{L} and $\neg\varphi_{\mathrm{R}}$, with a slight extension to handle atoms. The *closure* $\mathrm{cl}(\varphi)$ of a GF formula φ is the smallest set C of formulas containing φ such that:

- if $\psi_1 \wedge \psi_2 \in C$ or $\psi_1 \vee \psi_2 \in C$, then $\psi_1, \psi_2 \in C$;
- if $\exists \boldsymbol{x}.\alpha(\boldsymbol{xy}) \wedge \psi(\boldsymbol{xy}) \in C$, then $\alpha(\boldsymbol{xy}), \psi(\boldsymbol{xy}) \in C$;
- if $\forall \boldsymbol{x}.\alpha(\boldsymbol{xy}) \to \psi(\boldsymbol{xy}) \in C$, then $\neg\alpha(\boldsymbol{xy}), \psi(\boldsymbol{xy}) \in C$;
- if $\alpha(\boldsymbol{xy}) \in C$, then $\exists \boldsymbol{x}.\alpha(\boldsymbol{xy}) \in C$.

Given parameters $\boldsymbol{c}$, $\mathrm{cl}_{\mathrm{L}}(\boldsymbol{c})$ (respectively, $\mathrm{cl}_{\mathrm{R}}(\boldsymbol{c})$) consists of formulas from $\mathrm{cl}(\varphi_{\mathrm{L}})$ (respectively, $\mathrm{cl}(\neg\varphi_{\mathrm{R}})$) with free variables replaced by parameters from $\boldsymbol{c}$ or constants from $\boldsymbol{e}$ and labelled with provenance L (respectively, R).

Observe that the formulas in $\mathrm{cl}_{\mathrm{L}}(\boldsymbol{c}) \cup \mathrm{cl}_{\mathrm{R}}(\boldsymbol{c})$ respect the polarity of relations in the original formulas. This will be important for extracting Lyndon-style interpolants later, where the polarity of relations needs to be respected.

Guards. We say α is an L-*atom* (respectively, R-*atom*) if $(\alpha, +) \in \mathrm{occ}(\varphi_{\mathrm{L}})$ (respectively, $(\alpha, -) \in \mathrm{occ}(\varphi_{\mathrm{R}})$). Usually, any atom is a possible guard in a GF or GNF formula. For the purposes of interpolation, however, we need *atoms in the common signature* (*catoms* for short). We say α is an L-*catom* (respectively, R-*catom*) if $(\alpha, +) \in \mathrm{occ}(\varphi_{\mathrm{L}}) \cap \mathrm{occ}(\varphi_{\mathrm{R}})$ (respectively, $(\alpha, -) \in \mathrm{occ}(\varphi_{\mathrm{L}}) \cap \mathrm{occ}(\varphi_{\mathrm{R}})$).

Given a collection τ of formulas from $\mathrm{cl}_{\mathrm{L}}(\boldsymbol{c}) \cup \mathrm{cl}_{\mathrm{R}}(\boldsymbol{c})$, we say $\boldsymbol{b} \subseteq \boldsymbol{c}$ *is X-guarded in τ* (respectively, *X-cguarded in τ*) if $\boldsymbol{b}$ is empty, or if there is some $X : \exists \boldsymbol{x}.\alpha(\boldsymbol{xb}) \in \tau$ where α is an X-atom (respectively, X-catom). We allow $\boldsymbol{x}$ to be empty, in which case the X-guard or X-cguard for $\boldsymbol{b}$ is actually of the form $X : \alpha(\boldsymbol{b})$.

GF mosaics. A GF X-*mosaic* τ for $\varphi_{\mathrm{L}} \wedge \neg\varphi_{\mathrm{R}}$ using parameters $\boldsymbol{c}$ is a subset of $\Gamma(\boldsymbol{c}) := \mathrm{cl}_{\mathrm{L}}(\boldsymbol{c}) \cup \mathrm{cl}_{\mathrm{R}}(\boldsymbol{c})$ such that $\boldsymbol{c}$ is X-guarded in τ, and if $\widetilde{X} : \psi \in \tau$, then $\mathrm{par}(\psi)$ is X-cguarded in τ. We write $\tau(\boldsymbol{c})$ to emphasize that τ uses parameters $\boldsymbol{c}$.

Each X-mosaic provides information about an X-guarded set of parameters. Unlike the mosaics in [2], we place further restrictions on the formulas in the mosaics; namely, we restrict $\widetilde{X}$-formulas in an X-mosaic. Roughly speaking, the $\widetilde{X}$-formulas are the "dangerous" formulas that could potentially introduce a negation when we try to extract GNF interpolants from these mosaics later, so we require that any such formulas have an X-cguard in the mosaic (a guard that could be used in the interpolant).

Let τ be a GF X-mosaic over parameters $\boldsymbol{c}$. Assume $Y : \eta \in \tau$ is an *existential requirement* i.e. η is of the form $\exists \boldsymbol{x}.\alpha(\boldsymbol{xb}) \wedge \psi(\boldsymbol{xb})$. Then τ is *linked* to τ' via $Y : \eta$ (written $\tau \to_{Y:\eta} \tau'$) if

- τ' is a GF Y-mosaic over parameters $\boldsymbol{ab}$ with $\boldsymbol{a} \cap \boldsymbol{c} = \emptyset$;
- $Y : \alpha(\boldsymbol{xb})[\boldsymbol{a}/\boldsymbol{x}] \in \tau'$ and $Y : \psi(\boldsymbol{xb})[\boldsymbol{a}/\boldsymbol{x}] \in \tau'$;
- for all $Y : \psi'(\boldsymbol{b}) \in \tau$, $Y : \psi'(\boldsymbol{b}) \in \tau'$;
- for all $\widetilde{Y} : \psi'(\boldsymbol{b}) \in \tau$ such that $\mathrm{par}(\psi')$ is Y-cguarded in τ', $\widetilde{Y} : \psi'(\boldsymbol{b}) \in \tau'$;
- for all $Z : \psi'(\boldsymbol{b}) \in \tau'$, $Z : \psi'(\boldsymbol{b}) \in \tau$.

These conditions ensure that the maximum amount of information about the shared parameters $\boldsymbol{b}$ is passed from τ to τ' (subject to the restrictions on $\widetilde{Y}$ formulas in the Y-mosaic τ'), and no new information about $\boldsymbol{b}$ is added in τ'.

Fix some set P of parameters of size $2 \cdot \mathrm{width}(\varphi_{\mathrm{L}} \wedge \neg\varphi_{\mathrm{R}})$, and let $\mathcal{M}$ be the set of GF mosaics over parameters P. The following bound on the number of mosaics is easy to show.

Proposition 2. *There is a polynomial function p (independent of φ_{L} and φ_{R}) such that if $|\varphi_{\mathrm{L}}| + |\varphi_{\mathrm{R}}| = n$, then there are at most $2^{2^{p(n)}}$ mosaics in $\mathcal{M}$. Moreover, for a fixed bound on the arity of relations, there are at most $2^{p(n)}$ mosaics in $\mathcal{M}$.*

Coherence. We say a GF X-mosaic τ is *internally consistent* if there is no atomic $\alpha \in \sigma_{\mathrm{L}} \cup \sigma_{\mathrm{R}}$ such that $Y : \alpha, Z : \neg\alpha \in \tau$, and τ satisfies the following *downward closure* properties:

- if $Y : \alpha(\boldsymbol{ab}) \in \tau$, then $Y : \exists \boldsymbol{x}.\alpha(\boldsymbol{xb}) \in \tau$;
- if $Y : \exists \boldsymbol{x}.\alpha(\boldsymbol{xab}) \in \tau$, then $Y : \exists \boldsymbol{xy}.\alpha(\boldsymbol{xyb}) \in \tau$;
- if $Y : \psi_1 \wedge \psi_2 \in \tau$, then $Y : \psi_1, Y : \psi_2 \in \tau$;
- if $Y : \psi_1 \vee \psi_2 \in \tau$, then $Y : \psi_1 \in \tau$ or $Y : \psi_2 \in \tau$;
- if $Y : \forall \boldsymbol{x}.\alpha(\boldsymbol{xb}) \to \psi(\boldsymbol{xb}) \in \tau$ and there is $\boldsymbol{a} \subseteq \boldsymbol{c} \cup \boldsymbol{e}$ such that $Z : \alpha(\boldsymbol{xb})[\boldsymbol{a}/\boldsymbol{x}] \in \tau$ for some $Z \in \{\mathrm{L}, \mathrm{R}\}$, then $Y : \psi(\boldsymbol{xb})[\boldsymbol{a}/\boldsymbol{x}] \in \tau$.

A set of GF mosaics is said to be *saturated* if for every mosaic $\tau(\boldsymbol{c})$ in the set and for every existential requirement $Y : \eta \in \tau$ for η of the form $\exists \boldsymbol{x}.\alpha(\boldsymbol{xb}) \wedge \psi(\boldsymbol{xb})$, either (i) $Y : \alpha(\boldsymbol{xb})[\boldsymbol{a}/\boldsymbol{x}]$ and $Y : \psi(\boldsymbol{xb})[\boldsymbol{a}/\boldsymbol{x}]$ are in τ for some $\boldsymbol{a} \subseteq \boldsymbol{c} \cup \boldsymbol{e}$ (so η is fulfilled in τ) or (ii) τ is linked to some mosaic τ' in the set via $Y : \eta$ (so η is fulfilled in τ'). Finding a saturated set of internally consistent mosaics containing the original sentences provides a sound and complete method for GF satisfiability testing.

Theorem 3. *$\varphi_{\mathrm{L}} \wedge \neg\varphi_{\mathrm{R}}$ is satisfiable if and only if there is a saturated set of internally consistent mosaics from $\mathcal{M}$ that contains a mosaic τ such that* $\mathrm{L} : \varphi_{\mathrm{L}}, \mathrm{R} : \neg\varphi_{\mathrm{R}} \in \tau$.

We can now use the mosaic elimination procedure described in the previous section, and Theorem 3 ensures that $\varphi_{\mathrm{L}} \wedge \neg\varphi_{\mathrm{R}}$ is satisfiable iff there is some mosaic containing $\mathrm{L} : \varphi_{\mathrm{L}}$ and $\mathrm{R} : \neg\varphi_{\mathrm{R}}$ that is not eliminated using this process.

Constructive interpolation. Now assume $\varphi_{\mathrm{L}} \models \varphi_{\mathrm{R}}$, and consider the mosaic elimination procedure for the GF mosaics $\mathcal{M}$ for $\varphi_{\mathrm{L}} \wedge \neg\varphi_{\mathrm{R}}$. We write $\mathcal{N}'$ for the set of mosaics for $\varphi_{\mathrm{L}} \wedge \neg\varphi_{\mathrm{R}}$ that are eliminated using this procedure, and $\mathcal{N}_i$ for the set of mosaics that have been removed by stage i. By Theorem 3, we know that every mosaic containing $\mathrm{L} : \varphi_{\mathrm{L}}$ and $\mathrm{R} : \neg\varphi_{\mathrm{R}}$ must be in $\mathcal{N}'$.

We first show how to construct a mosaic interpolant θ^{τ}_X for each X-mosaic $\tau(\boldsymbol{c})$ in $\mathcal{N}'$. Roughly speaking, θ^{τ}_X is an interpolant (in the usual sense) for $\tau_X \models \neg\tau_{\widetilde{X}}$, where

$$\tau_Y := \bigwedge_{Y:\psi\in\tau} \psi \quad \wedge \quad \bigwedge_{Y:\psi\in\overline{\tau}} \neg\psi$$

and $\overline{\tau}$ is the set consisting of formulas $Z : \psi \in \mathrm{cl}_Y(\boldsymbol{c}) \setminus \tau$ such that $Z = X$, or $Z = \widetilde{X}$ and $\mathrm{par}(\psi)$ is X-cguarded in τ. In other words, $\overline{\tau}$ contains the formulas from the closure that satisfy the requirements for being in τ (in terms of X-cguardedness), but are not in τ. Intuitively, θ^τ_X describes the reason τ was eliminated, i.e. it expresses why $\tau_X \wedge \tau_{\widetilde{X}}$ is not satisfiable.

We have the following Lyndon-style mosaic interpolant properties.

Lemma 4. *For each* GF *X-mosaic $\tau(\boldsymbol{c}) \in \mathcal{N}'$ (the set of eliminated mosaics), we can construct a DAG representation of a formula θ^τ_X such that*
- $\tau_X \models \theta^\tau_X$ *and* $\theta^\tau_X \models \neg\tau_{\widetilde{X}}$*;*
- $\mathrm{occ}(\theta^\tau_X) \subseteq \mathrm{occ}(\varphi_\mathrm{L}) \cap \mathrm{occ}(\varphi_\mathrm{R})$ *if* $X = \mathrm{L}$, $\mathrm{occ}(\neg\theta^\tau_X) \subseteq \mathrm{occ}(\varphi_\mathrm{L}) \cap \mathrm{occ}(\varphi_\mathrm{R})$ *if* $X = \mathrm{R}$*;*
- $\mathrm{con}(\theta^\tau_X) \subseteq \mathrm{con}(\varphi_\mathrm{L}) \cup \mathrm{con}(\varphi_\mathrm{R})$*;*
- $\mathrm{par}(\theta^\tau_X) \subseteq \boldsymbol{c}$*;*
- θ^τ_X *is in* GNF*, even when parameters are viewed as free variables.*

Moreover, if $n = |\varphi_\mathrm{L}| + |\varphi_\mathrm{R}|$ then there is a DAG representation of $\Theta_i := \{\theta^\tau_Z : \tau \text{ is a } Z\text{-mosaic in } \mathcal{N}_i\}$ such that $|\Theta_i| \leq i \cdot 2^{2^{p'(n)}}$ ($|\Theta_i| \leq i \cdot 2^{p'(n)}$ for the fixed arity case), where p' is some polynomial function independent of φ_L and φ_R.

Proof. We first introduce some additional notation. We use $\top$ (respectively, $\bot$) as an abbreviation for the GNF sentence $\exists x.x = x$ (respectively, $\exists x.(x = x \wedge \neg(x = x))$). During the interpolant construction, we will need to include guards in order to remain in GNF, so for any tuple $\boldsymbol{b}$ of parameters that are X-cguarded in τ, we define a formula $\mathrm{gdd}^\tau_X(\boldsymbol{b})$. If $\boldsymbol{b} = \emptyset$, then set $\mathrm{gdd}^\tau_X(\boldsymbol{b}) := \top$. Otherwise, there is some formula $X : \exists \boldsymbol{x}.\beta(\boldsymbol{xb}) \in \tau$ for some X-catom β, and we set $\mathrm{gdd}^\tau_X(\boldsymbol{b}) := \exists \boldsymbol{x}.\beta(\boldsymbol{xb})$ (note that $\boldsymbol{x}$ can be empty, in which case this is just an atomic formula). We will also write $\mathrm{gdd}^\tau_X(\boldsymbol{b}) \wedge \neg\psi$ to indicate some formula $\exists \boldsymbol{x}.(\beta(\boldsymbol{xb}) \wedge \neg\psi)$.

We proceed by induction on i, the stage at which τ was eliminated. We define the interpolants and then sketch the proof of correctness in the inductive case.

Base case. Consider an X-mosaic $\tau(\boldsymbol{c}) \in \mathcal{N}_1$. Then τ has an internal inconsistency, and we consider the following cases.
- Assume one of these conditions holds for $Y \in \{\mathrm{L}, \mathrm{R}\}$:
  - $Y : \alpha(\boldsymbol{b}), Y : \neg\alpha(\boldsymbol{b}) \in \tau$;
  - $Y : \alpha(\boldsymbol{ab}) \in \tau$, but $Y : \exists \boldsymbol{x}.\alpha(\boldsymbol{xb}) \notin \tau$;
  - $Y : \exists \boldsymbol{x}.\alpha(\boldsymbol{xab}) \in \tau$, but $Y : \exists \boldsymbol{xy}.\alpha(\boldsymbol{xyb}) \notin \tau$;
  - $Y : \psi_1 \wedge \psi_2 \in \tau$, but $Y : \psi_1 \notin \tau$ or $Y : \psi_2 \notin \tau$;
  - $Y : \psi_1 \vee \psi_2 \in \tau$, but $Y : \psi_1 \notin \tau$ and $Y : \psi_2 \notin \tau$;
  - $Y : \alpha(\boldsymbol{xb})[\boldsymbol{a}/\boldsymbol{x}], Y : \forall \boldsymbol{x}.\alpha(\boldsymbol{xb}) \to \psi(\boldsymbol{xb}) \in \tau$, but $Y : \psi(\boldsymbol{xb})[\boldsymbol{a}/\boldsymbol{x}] \notin \tau$.

  If $Y = X$, then $\theta^\tau_X := \bot$. If $Y = \widetilde{X}$, then $\theta^\tau_X := \top$.
- If $X : \alpha(\boldsymbol{b}), \widetilde{X} : \neg\alpha(\boldsymbol{b}) \in \tau$, then $\theta^\tau_X := \alpha(\boldsymbol{b})$.
- If $X : \neg\alpha(\boldsymbol{b}), \widetilde{X} : \alpha(\boldsymbol{b}) \in \tau$, then $\theta^\tau_X := \mathrm{gdd}^\tau_X(\boldsymbol{b}) \wedge \neg\alpha(\boldsymbol{b})$.
- If $X : \alpha(\boldsymbol{xb})[\boldsymbol{a}/\boldsymbol{x}], \widetilde{X} : \forall \boldsymbol{x}.\alpha(\boldsymbol{xb}) \to \psi(\boldsymbol{xb}) \in \tau$, but $\widetilde{X} : \psi(\boldsymbol{xb})[\boldsymbol{a}/\boldsymbol{x}] \notin \tau$, then $\theta^\tau_X := \alpha(\boldsymbol{xb})[\boldsymbol{a}/\boldsymbol{x}]$.
- If $\widetilde{X} : \alpha(\boldsymbol{xb})[\boldsymbol{a}/\boldsymbol{x}], X : \forall \boldsymbol{x}.\alpha(\boldsymbol{xb}) \to \psi(\boldsymbol{xb}) \in \tau$, but $X : \psi(\boldsymbol{xb})[\boldsymbol{a}/\boldsymbol{x}] \notin \tau$, then $\theta^\tau_X := \mathrm{gdd}^\tau_X(\boldsymbol{a}'\boldsymbol{b}) \wedge \neg\alpha(\boldsymbol{xb})[\boldsymbol{a}/\boldsymbol{x}]$, where $\boldsymbol{a}'\boldsymbol{b} = \mathrm{par}(\alpha(\boldsymbol{xb})[\boldsymbol{a}/\boldsymbol{x}])$.

Inductive case. Now consider an X-mosaic $\tau(\boldsymbol{c}) \in \mathcal{N}_i \setminus \mathcal{N}_{i-1}$. This means there is some $Y : \eta = Y : \exists \boldsymbol{x}.\alpha(\boldsymbol{xb}) \wedge \psi(\boldsymbol{xb}) \in \tau$ that is not fulfilled in τ, and if $\tau \to_{Y:\eta} \tau'$ then $\tau' \in \mathcal{N}_{i-1}$.

Assume $\mathrm{par}(\psi) = \boldsymbol{b} \subseteq \boldsymbol{c}$, and fix some $\boldsymbol{a} \in P \setminus \boldsymbol{c}$.

We introduce some additional notation. For $Z \in \{\mathrm{L}, \mathrm{R}\}$, we define $\tau{\restriction}^{\boldsymbol{b}}_Z$ to be the Z-mosaic that results from taking the Z-formulas in τ that only use parameters from $\boldsymbol{b}$, and then unioning with the $\widetilde{Z}$-formulas in τ whose parameters are contained in $\boldsymbol{b}$ and are Z-cguarded in τ. For mosaics υ, υ' and $Z \in \{\mathrm{L}, \mathrm{R}\}$, we write $\upsilon'(\boldsymbol{ab}) \supseteq_Z \upsilon$ if υ' uses only parameters from $\boldsymbol{ab}$ and is obtained from υ be adding only Z-formulas (so the $\widetilde{Z}$-formulas in υ and υ' are identical).

If $Y = X$, then define θ^τ_X to be

$$\bigvee_{\substack{\upsilon'(\boldsymbol{ab}) \supseteq_X \tau{\restriction}^{\boldsymbol{b}}_X \\ \text{s.t. } \tau \to_{X:\eta} \upsilon'}} \exists \boldsymbol{x}.\Bigg(\bigwedge_{\substack{\tau'(\boldsymbol{ab}) \supseteq_{\widetilde{X}} \upsilon' \\ \text{s.t. } \tau \to_{X:\eta} \tau'}} \theta^{\tau'}_X[\boldsymbol{x}/\boldsymbol{a}] \Bigg).$$

Informally, this interpolant expresses that there is a way to fulfil the existential requirement by adding some X-formulas to $\tau{\restriction}^{\boldsymbol{b}}_X$, but no matter what $\widetilde{X}$-formulas we add, the resulting linked mosaic τ' has already been eliminated.

If $Y = \widetilde{X}$, then define θ^τ_X to be

$$\mathrm{gdd}^\tau_X(\boldsymbol{b}) \wedge \neg\Bigg(\bigvee_{\substack{\upsilon'(\boldsymbol{ab}) \supseteq_{\widetilde{X}} \tau{\restriction}^{\boldsymbol{b}}_{\widetilde{X}} \\ \text{s.t. } \tau \to_{\widetilde{X}:\eta} \upsilon'}} \exists \boldsymbol{x}.\Bigg(\bigwedge_{\substack{\tau'(\boldsymbol{ab}) \supseteq_X \upsilon' \\ \text{s.t. } \tau \to_{\widetilde{X}:\eta} \tau'}} \theta^{\tau'}_{\widetilde{X}}[\boldsymbol{x}/\boldsymbol{a}] \Bigg)\Bigg).$$

In both cases, notice that the parameters $\boldsymbol{ab}$ in the inductively defined interpolants $\theta^{\tau'}_Y$ are guarded by the atom $\alpha(\boldsymbol{ab})$ appearing in the existential requirement $Y : \exists \boldsymbol{x}.\alpha(\boldsymbol{xb}) \wedge \psi(\boldsymbol{xb})$. We cannot necessarily guard the quantification in the interpolant using $\alpha(\boldsymbol{xb})$ since α may not be in the common signature. This is why the resulting interpolant might not be in GF, even when the input sentences are in GF. However, in case $Y = \widetilde{X}$, we can always guard $\boldsymbol{b}$ in the common signature using $\mathrm{gdd}^\tau_X(\boldsymbol{b})$. This guard is guaranteed to exist because the existential requirement is a $\widetilde{X}$-formula using parameters $\boldsymbol{b}$ in an X-mosaic. This ensures that the interpolant is in GNF. The occurrence of constants and the polarity of relations follows from the inductive hypothesis and the fact that the polarity of occurrences of relations are preserved when moving to mosaics.

We now sketch part of the proof of correctness, showing that $\tau_X \models \theta^\tau_X$ and $\theta^\tau_X \models \neg\tau_{\widetilde{X}}$ when $Y = X$.

Assume there is a model $\mathfrak{M}$ for τ_X. Since $X : \eta \in \tau$, this means that $\mathfrak{M} \models \exists \boldsymbol{x}.\alpha(\boldsymbol{xb}) \wedge \psi(\boldsymbol{xb})$, so there are elements $\boldsymbol{a}' \in M$ and an expansion $\mathfrak{M}'$ of $\mathfrak{M}$ with the interpretation $\boldsymbol{a}^{\mathfrak{M}'} := \boldsymbol{a}'$ such that $\mathfrak{M}' \models \alpha(\boldsymbol{ab}) \wedge \psi(\boldsymbol{ab})$.

Take υ' to be the union of $\tau{\restriction}^{\boldsymbol{b}}_X$ and the set of formulas $X : \psi'$ in $\mathrm{cl}_X(\boldsymbol{ab})$ such that $\mathfrak{M}' \models \psi'$. It can be checked that $\tau \to_{X:\eta} \upsilon'$. By construction, $\mathfrak{M}' \models \upsilon'_X$. Moreover, for any choice of $\tau'(\boldsymbol{ab}) \supseteq_{\widetilde{X}} \upsilon'$, $\mathfrak{M}' \models \tau'_X$ (since the X-formulas are identical in υ' and τ'). By the inductive hypothesis, this means that $\mathfrak{M}' \models \theta^{\tau'}_X$. Hence, $\mathfrak{M}' \models \bigwedge_{\substack{\tau'(\boldsymbol{ab}) \supseteq_{\widetilde{X}} \upsilon' \\ \text{s.t. } \tau \to_{X:\eta} \tau'}} \theta^{\tau'}_X$ and $\mathfrak{M} \models \theta^\tau_X$.

Now assume that there is a model $\mathfrak{M}$ of θ^τ_X. Then there is some $\upsilon'(\boldsymbol{ab}) \supseteq_X \tau{\restriction}^{\boldsymbol{b}}_X$ with $\tau \to_{X:\eta} \upsilon'$, elements $\boldsymbol{a}' \in M$, and an expansion $\mathfrak{M}'$ of $\mathfrak{M}$ with the interpretation $\boldsymbol{a}^{\mathfrak{M}'} := \boldsymbol{a}'$, such that $\mathfrak{M}' \models \bigwedge_{\substack{\tau'(\boldsymbol{ab}) \supseteq_{\widetilde{X}} \upsilon' \\ \text{s.t. } \tau \to_{X:\eta} \tau'}} \theta^{\tau'}_X$.

Note that for all $\tau'(\boldsymbol{ab}) \supseteq_{\widetilde{X}} \upsilon'$ such that $\tau \to_{X:\eta} \tau'$, $\mathfrak{M}' \models \theta^{\tau'}_X$ and consequently, the inductive hypothesis implies $\mathfrak{M}' \models \neg\tau'_{\widetilde{X}}$. In particular, consider $\tau' := \upsilon' \cup S'$ where S' is the set of formulas $\widetilde{X} : \psi'$ in $\mathrm{cl}_{\widetilde{X}}(\boldsymbol{ab})$ such that $\mathfrak{M}' \models \psi'$, $\mathrm{par}(\psi') \cap \boldsymbol{a} \neq \emptyset$, and $\widetilde{X} : \psi'$ is X-cguarded in υ' as required in an X-mosaic. It is clear that $\tau'(\boldsymbol{ab}) \supseteq_{\widetilde{X}} \upsilon'$ and $\tau \to_{X:\eta} \tau'$, so $\mathfrak{M}' \models \neg\tau'_{\widetilde{X}}$ as described above.

Since $\mathfrak{M}' \models \neg\tau'_{\widetilde{X}}$, there is some conjunct χ in $\tau'_{\widetilde{X}}$ such that $\mathfrak{M}' \models \neg\chi$. Consider some conjunct $\chi(\boldsymbol{ab})$ in $\tau'_{\widetilde{X}}$ that actually uses some parameters from $\boldsymbol{a}$. Then by choice of S', $\mathfrak{M}' \models \chi(\boldsymbol{ab})$, so this formula cannot witness the fact that $\mathfrak{M}' \models \neg\tau'_{\widetilde{X}}$. This means there must be some conjunct $\chi(\boldsymbol{b})$ in $\tau'_{\widetilde{X}}$ that only uses parameters from $\boldsymbol{b}$ such that $\mathfrak{M}' \models \neg\chi(\boldsymbol{b})$. If $\widetilde{X} : \chi(\boldsymbol{b})$ is in τ', then $\widetilde{X} : \chi(\boldsymbol{b}) \in \tau$ by the definition of $\tau \rightarrow_{X:\eta} \tau'$. Likewise, if $\widetilde{X} : \chi(\boldsymbol{b}) \notin \tau'$, then it must be the case that $\mathrm{par}(\chi)$ is X-cguarded in τ', so $\widetilde{X} : \chi(\boldsymbol{b}) \notin \tau$ (otherwise it would contradict the definition of $\tau \rightarrow_{X:\eta} \tau'$). In either case, this means $\chi(\boldsymbol{b})$ must be a conjunct in $\tau_{\widetilde{X}}$. Hence, $\mathfrak{M} \models \neg\tau_{\widetilde{X}}$. □

We can prove the following constructive Lyndon interpolation result using the mosaic interpolants from Lemma 4.

Theorem 5. *Let φ_{L} and φ_{R} be* GF *sentences without equality over signatures σ_{L} and σ_{R}, respectively. If $\varphi_{\mathrm{L}} \models \varphi_{\mathrm{R}}$ and $|\varphi_{\mathrm{L}}| + |\varphi_{\mathrm{R}}| = n$, then we can construct a DAG representation of a* GNF *interpolant θ such that*

- $\varphi_{\mathrm{L}} \models \theta$ *and* $\theta \models \varphi_{\mathrm{R}}$*;*
- $\mathrm{occ}(\theta) \subseteq \mathrm{occ}(\varphi_{\mathrm{L}}) \cap \mathrm{occ}(\varphi_{\mathrm{R}})$*;*
- $\mathrm{con}(\theta) \subseteq \mathrm{con}(\varphi_{\mathrm{L}}) \cup \mathrm{con}(\varphi_{\mathrm{R}})$*;*
- *the DAG representation of θ is of size at most $2^{2^{p(n)}}$ for some polynomial function p independent of φ_{L} and φ_{R} (and this can be improved to a size of at most $2^{p(n)}$ when the bound on the arity of relations is fixed).*

For brevity, *in this theorem and throughout the paper, we give only bounds on the output size, not the running time of the algorithms.* However the proofs will show that the worst-case running time is bounded by a polynomial in the output size. In particular, the running time of the interpolation algorithm above is doubly-exponential in the input (but only singly-exponential for bounded arity GF sentences without equality).

Proof of Theorem 5. Let

$$\theta := \bigvee_{\tau(\emptyset) \supseteq_{\mathrm{L}} \{\mathrm{L}:\varphi_{\mathrm{L}}\}} \ \bigwedge_{\substack{\tau'(\emptyset) \supseteq_{\mathrm{R}} \tau \\ \text{s.t. } \tau' \in \mathcal{N}'}} \theta_{\mathrm{L}}^{\tau'}$$

where $\mathcal{N}'$ is the set of eliminated mosaics and we write $\upsilon'(\emptyset) \supseteq_Z \upsilon$ if υ' is over parameters $\emptyset$ and is obtained from υ be adding only Z-formulas.

We prove that $\varphi_{\mathrm{L}} \models \theta \models \varphi_{\mathrm{R}}$.

Assume there is a structure $\mathfrak{M}$ such that $\mathfrak{M} \models \varphi_{\mathrm{L}}$. Let $\tau(\emptyset) \supseteq_{\mathrm{L}} \{\mathrm{L} : \varphi_{\mathrm{L}}\}$ be the set of sentences from $\mathrm{cl}_{\mathrm{L}}(\emptyset)$ that are true in $\mathfrak{M}$. For any $\tau'(\emptyset) \supseteq_{\mathrm{R}} \tau$ such that $\tau' \in \mathcal{N}'$, $\mathfrak{M} \models \tau'_{\mathrm{L}}$. By Lemma 4, this means that $\mathfrak{M} \models \theta_{\mathrm{L}}^{\tau'}$, so $\mathfrak{M} \models \theta$.

Now assume there is a structure $\mathfrak{M}$ such that $\mathfrak{M} \models \theta$. Then there is some $\tau(\emptyset) \supseteq_{\mathrm{L}} \{\mathrm{L} : \varphi_{\mathrm{L}}\}$ such that $\mathfrak{M} \models \theta_{\mathrm{L}}^{\tau'}$ for all $\tau'(\emptyset) \supseteq_{\mathrm{R}} \tau$ such that $\tau' \in \mathcal{N}'$. Let S' be the set of formulas $\mathrm{R} : \psi$ in $\mathrm{cl}_{\mathrm{R}}(\emptyset)$ such that $\mathfrak{M} \models \psi$. Consider $\tau' := \tau \cup S' \cup \{\mathrm{R} : \neg\varphi_{\mathrm{R}}\} \supseteq_{\mathrm{R}} \tau$. By completeness of the mosaic system (Theorem 3), every mosaic containing $\mathrm{L} : \varphi_{\mathrm{L}}$ and $\mathrm{R} : \neg\varphi_{\mathrm{R}}$ must be in $\mathcal{N}'$. Hence, $\mathfrak{M} \models \theta_{\mathrm{L}}^{\tau'}$. By Lemma 4 this implies that $\mathfrak{M} \models \neg\tau'_{\mathrm{R}}$. But every formula $\mathrm{R} : \psi$ in S' was chosen such that $\mathfrak{M} \models \psi$. The only way that $\mathfrak{M} \models \neg\tau'_{\mathrm{R}}$ is if $\mathfrak{M} \models \varphi_{\mathrm{R}}$.

The other properties in Theorem 5 follow from Lemma 4. □

Although GF does not have CIP, it was shown in [14] that GF does admit the following weaker form of interpolation.

Theorem 6 ([14])**.** *For* GF *sentences $\varphi_{\mathrm{L}} \models \varphi_{\mathrm{R}}$ there is a* GF *sentence θ such that $\varphi_{\mathrm{L}} \models \theta$ and $\theta \models \varphi_{\mathrm{R}}$ and θ only uses relations in the common signature or relations occurring as a guard for some quantification in φ_{L} or φ_{R}.*

We remark that our mosaic interpolation method can be modified to produce GF interpolants like this, with the same DAG-size bounds as in Theorem 5, by modifying the inductive case of Lemma 4 to include the guard that is present already in the existential requirement.

5. GNF Mosaics

We now aim to describe mosaics that can handle φ_{L} and φ_{R} in GNF. For now, we assume φ_{L} and φ_{R} are GNF sentences, over signatures σ_{L} and σ_{R}, that are in GN-normal form, use ordinary quantifiers rather than relativized quantifiers, and do not use equality. Let $\boldsymbol{e} := \mathrm{con}(\varphi_{\mathrm{L}}) \cup \mathrm{con}(\varphi_{\mathrm{R}})$.

Closure. It turns out that just taking the subformula closure is not sufficient to determine satisfiability of $\varphi_{\mathrm{L}} \wedge \neg\varphi_{\mathrm{R}}$ from a set of coherent mosaics. Instead, the mosaics must contain additional information about CQ-shaped subformulas (subformulas of the form $\exists \boldsymbol{x}. \bigwedge_i \psi_i(\boldsymbol{xy})$).

Our solution is to add specializations to the closure. Informally, each specialization describes a way in which the original CQ-shaped formula could be satisfied in some tree-like structure.

Consider a GN-normal form formula $\phi(\boldsymbol{y})$ that is of the form $\exists \boldsymbol{x}. \bigwedge_i \psi_i(\boldsymbol{xy})$. A *specialization* of ϕ is a formula ϕ' obtained from ϕ by the following operations:

- select a subset $\boldsymbol{z}$ of $\boldsymbol{x}$ (call variables from $\boldsymbol{yz}$ the *inside variables* and variables from $\boldsymbol{x} \setminus \boldsymbol{z}$ the *outside variables*);
- select a partition $\boldsymbol{x}_1, \ldots, \boldsymbol{x}_k$ of the outside variables, with the property that for every ψ_j, either ψ_j has no outside variables or all of its outside variables are contained in some partition element $\boldsymbol{x}_i$;
- let χ_0 be the conjunction of the ψ_i using only inside variables, and let χ_j be the conjunction of the ψ_i using outside variables and satisfying $\mathrm{par}(\psi_i) \subseteq \boldsymbol{x}_j\boldsymbol{yz}$;
- set $\phi'(\boldsymbol{yz})$ to be $\chi_0(\boldsymbol{yz}) \wedge \bigwedge_{j \in \{1,\ldots,k\}} \exists \boldsymbol{x}_j.\chi_j(\boldsymbol{x}_j\boldsymbol{yz})$.

We are now ready to define the closure. Each formula is labelled with a provenance $X \in \{\mathrm{L}, \mathrm{R}\}$ and a *polarity* $p \in \{+, -\}$. The *closure* $\mathrm{cl}(X^p : \varphi)$ is the smallest set C of formulas containing $X^p : \varphi$ and such that:

- if $X^+ : \alpha(\boldsymbol{y}) \wedge \neg\psi(\boldsymbol{y}) \in C$, then $X^+ : \alpha(\boldsymbol{y}), X^- : \psi(\boldsymbol{y}) \in C$;
- if $X^- : \alpha(\boldsymbol{y}) \wedge \neg\psi(\boldsymbol{y}) \in C$, then $X^- : \alpha(\boldsymbol{y}), X^+ : \psi(\boldsymbol{y}) \in C$;
- if $X^p : \bigwedge_i \psi_i \in C$ or $X^p : \bigvee_i \psi_i \in C$, then $X^p : \psi_i \in C$ for all i;
- if $X^p : \exists \boldsymbol{x}.\psi(\boldsymbol{xy}) \in C$, then $X^p : \phi' \in C$ for all specializations ϕ' of $\exists \boldsymbol{x}.\psi(\boldsymbol{xy})$;
- if $X^+ : \alpha(\boldsymbol{xy}) \in C$, then $X^+ : \exists \boldsymbol{x}.\alpha(\boldsymbol{xy}) \in C$.

Given parameters $\boldsymbol{c}$, we let $\mathrm{cl}_{\mathrm{L}}(\boldsymbol{c})$ (respectively, $\mathrm{cl}_{\mathrm{R}}(\boldsymbol{c})$) consist of formulas from $\mathrm{cl}(\mathrm{L}^+ : \varphi_{\mathrm{L}})$ (respectively, $\mathrm{cl}(\mathrm{R}^- : \varphi_{\mathrm{R}})$) with free variables replaced by parameters from $\boldsymbol{c}$ or constants from $\boldsymbol{e}$.

Roughly speaking, if $X^p : \psi \in \mathrm{cl}_{\mathrm{L}}(\boldsymbol{c}) \cup \mathrm{cl}_{\mathrm{R}}(\boldsymbol{c})$, then the provenance $X \in \{\mathrm{L}, \mathrm{R}\}$ indicates whether this formula came from a subformula of φ_{L} or $\neg\varphi_{\mathrm{R}}$, and the polarity $p \in \{+, -\}$ indicates whether this subformula occurred positively or negatively (i.e., whether it occurred within the scope of an even or odd number of negations). We write $p\psi$ to mean ψ (respectively, $\neg\psi$) if p is $+$ (respectively, $-$).

As in GF mosaics, we remark that the formulas in $\mathrm{cl}_{\mathrm{L}}(\boldsymbol{c}) \cup \mathrm{cl}_{\mathrm{R}}(\boldsymbol{c})$ respect the polarity of relations in the original formulas.

Guards. Let τ be a collection of formulas from $\mathrm{cl}_{\mathrm{L}}(\boldsymbol{c}) \cup \mathrm{cl}_{\mathrm{R}}(\boldsymbol{c})$. We say $\boldsymbol{b} \subseteq \boldsymbol{c}$ *is X-guarded in τ* (respectively, *X-cguarded in τ*) if $\boldsymbol{b}$ is empty, or if there is some $X^+ : \exists \boldsymbol{x}.\alpha(\boldsymbol{xb}) \in \tau$ where α is an X-atom (respectively, X-catom). (The definition of X-atom and X-catom is the same as in Section 4.)

GNF mosaics. Let $\boldsymbol{c}$ be a tuple of parameters such that $|\boldsymbol{c}| \leq \mathrm{width}(\varphi_{\mathrm{L}} \wedge \neg\varphi_{\mathrm{R}})$. A GNF X*-mosaic* τ for $\varphi_{\mathrm{L}} \wedge \neg\varphi_{\mathrm{R}}$ using parameters $\boldsymbol{c}$ is a subset of $\Gamma(\boldsymbol{c}) := \mathrm{cl}_{\mathrm{L}}(\boldsymbol{c}) \cup \mathrm{cl}_{\mathrm{R}}(\boldsymbol{c})$ such that if $\widetilde{X}^p : \psi \in \tau$, then $\mathrm{par}(\psi)$ is X-cguarded in τ.

Like the GF mosaics, the $\widetilde{X}$-formulas are restricted in X-mosaics. Unlike a GF mosaic, this restriction only applies to positive $\widetilde{X}$-formulas.

Let τ be a GNF X-mosaic over parameters $\boldsymbol{c}$. Assume $Y^+ : \eta \in \tau$ for η an *existential requirement* of the form $\exists \boldsymbol{x}.\psi(\boldsymbol{x}\boldsymbol{b})$. Then τ is *linked* to τ' via $Y^+ : \eta$ (written $\tau \rightarrow_{Y^+:\eta} \tau'$) if

- τ' is a GNF Y-mosaic over $\boldsymbol{a}\boldsymbol{b}$ with $\boldsymbol{a} \cap \boldsymbol{c} = \emptyset$;
- $Y^+ : \psi(\boldsymbol{a}\boldsymbol{b}) \in \tau'$;
- for all $Y^p : \psi'(\boldsymbol{b}) \in \tau$, $Y^p : \psi'(\boldsymbol{b}) \in \tau'$;
- for all $\widetilde{Y}^p : \psi'(\boldsymbol{b}) \in \tau$ such that either $\mathrm{par}(\psi')$ is Y-cguarded in τ' or $p = -$, $\widetilde{Y}^p : \psi'(\boldsymbol{b}) \in \tau'$;
- for all $Z^p : \psi'(\boldsymbol{b}) \in \tau'$, $Z^p : \psi'(\boldsymbol{b}) \in \tau$.

As in the GF mosaics, the maximum amount of information about shared parameters is passed through a link, and no new information about shared parameters is added in τ'.

Fix some set P of parameters of size $2 \cdot \mathrm{width}(\varphi_{\mathrm{L}} \wedge \neg\varphi_{\mathrm{R}})$, and let $\mathcal{M}$ be the set of GNF mosaics over parameters P. The following result is proven by a routine calculation.

Proposition 7. *There is a polynomial function p (independent of φ_{L} and φ_{R}) such that if $|\varphi_{\mathrm{L}}| + |\varphi_{\mathrm{R}}| = n$, then there are at most $2^{2^{p(n)}}$ mosaics in $\mathcal{M}$. This bound holds even if φ_{L} and φ_{R} are* GNF *sentences but are not in* GN*-normal form.*

Coherence. A GNF X-mosaic $\tau(\boldsymbol{c})$ is *downward closed* if the following properties are satisfied:

- if $Y^+ : \alpha(\boldsymbol{b}) \wedge \neg\psi(\boldsymbol{b}) \in \tau$, then $Y^+ : \alpha(\boldsymbol{b}) \in \tau$ and $Y^- : \psi(\boldsymbol{b}) \in \tau$;
- if $Y^- : \alpha(\boldsymbol{b}) \wedge \neg\psi(\boldsymbol{b}) \in \tau$ and $Z^+ : \alpha(\boldsymbol{b}) \in \tau$, then $Y^+ : \psi(\boldsymbol{b}) \in \tau$;
- if $Y^+ : \bigwedge_i \psi_i \in \tau$ (respectively, $Y^- : \bigvee_i \psi_i \in \tau$), then $Y^+ : \psi_i \in \tau$ (respectively, $Y^- : \psi_i \in \tau$) for all i;
- if $Y^+ : \bigvee_i \psi_i \in \tau$ (respectively, $Y^- : \bigwedge_i \psi_i \in \tau$), then $Y^+ : \psi_i \in \tau$ (respectively, $Y^- : \psi_i \in \tau$) for some i;
- if $Y^- : \exists \boldsymbol{x}.\psi(\boldsymbol{x}\boldsymbol{b}) \in \tau$, then for all specializations $\phi'(\boldsymbol{y}\boldsymbol{z})$ of $\exists \boldsymbol{x}.\psi(\boldsymbol{x}\boldsymbol{y})$ and for all $\boldsymbol{c}' \subseteq \boldsymbol{c} \cup \boldsymbol{e}$ with $|\boldsymbol{c}'| = |\boldsymbol{z}|$, $Y^- : \phi'(\boldsymbol{b}\boldsymbol{z})[\boldsymbol{c}'/\boldsymbol{z}] \in \tau$.

The last rule in this definition is a generalization of the rule for universal requirements in the GF mosaics. It helps ensure that when we construct a model from a saturated set of internally consistent mosaics, a CQ-shaped subformula that is asserted to be false in one mosaic, does not become true in the constructed model.

As before, a GNF X-mosaic τ is *internally consistent* if τ is downward closed and there is no atomic formula α and no $Y, Z \in \{\mathrm{L}, \mathrm{R}\}$ such that $Y^+ : \alpha \in \tau$ and $Z^- : \alpha \in \tau$.

A set of GNF mosaics is said to be *saturated* if for every mosaic $\tau(\boldsymbol{c})$ in the set and for every existential requirement $Y^+ : \eta \in \tau$ for η of the form $\exists \boldsymbol{x}.\psi(\boldsymbol{x}\boldsymbol{b})$, either (i) $Y^+ : \psi(\boldsymbol{x}\boldsymbol{b})[\boldsymbol{a}/\boldsymbol{x}]$ is in τ for some $\boldsymbol{a} \subseteq \boldsymbol{c} \cup \boldsymbol{e}$ (so η is fulfilled in τ) or (ii) τ is linked to some mosaic τ' in the set via $Y^+ : \eta$ (so η is fulfilled in τ').

These requirements lead to the following key theorem.

Theorem 8. *$\varphi_{\mathrm{L}} \wedge \neg\varphi_{\mathrm{R}}$ is satisfiable if and only if there is a saturated set of internally consistent* GNF *mosaics from $\mathcal{M}$ that contains a mosaic τ such that $\mathrm{L}^+ : \varphi_{\mathrm{L}}, \mathrm{R}^- : \varphi_{\mathrm{R}} \in \tau$.*

Using the elimination procedure described earlier, this gives an alternative proof of the 2EXPTIME upper bound for satisfiability of GNF (as stated in Theorem 1).

Extensions of GNF mosaics. We can consider various extensions of these mosaics, to handle input with free variables, relativized quantifiers, and equality. We briefly summarize some of these extensions.

Relativization. For input that uses $\mathbb{U}$-relativized quantifiers, a *relativized X-mosaic* $\tau(\boldsymbol{c})$ is a collection of $\mathbb{U}$-relativized GNF sentences with parameters from $\boldsymbol{c}$. Each relativized mosaic also has a set of distinguished parameters $\boldsymbol{d} \subseteq \boldsymbol{c}$ (which we call the relativized parameters in the mosaic). These parameters $\boldsymbol{d}$ in τ will be removed from θ^{τ}_{X} using existential quantification when constructing the mosaic interpolants. Hence, we add further restrictions on the $\widetilde{X}$-formulas that use parameters from $\boldsymbol{d}$ to ensure that we can relativize any quantification over $\boldsymbol{d}$. In particular, if $d \in \boldsymbol{d}$ appears in a $\widetilde{X}$-formula in an X-mosaic τ then there must be some relativizer $X^+ : Ud \in \tau$ such that $U \in \mathbb{U}$ is an X-catom.

Equality. We have ignored the presence of equality in the mosaic construction, but it can be accommodated by normalizing the use of equality in formulas, and then putting additional "equality consistency requirements" on the mosaics.

Constructive interpolation. Using these extensions of the GNF mosaics, we can lift the restrictions placed on φ_{L} and φ_{R} at the beginning of this section, and derive the following constructive interpolation results, subsuming Theorem 5. This proves constructive RIP for GNF, and is a key contribution.

Theorem 9. *Let φ_{L} and φ_{R} be* GNF *(respectively, equality-free* UNF*) formulas over signatures σ_{L} and σ_{R}, respectively. If $\varphi_{\mathrm{L}} \models \varphi_{\mathrm{R}}$ and $|\varphi_{\mathrm{L}}| + |\varphi_{\mathrm{R}}| = n$, then we can construct a DAG representation of a* GNF *(respectively,* UNF*) interpolant θ such that*

- *$\varphi_{\mathrm{L}} \models \theta$ and $\theta \models \varphi_{\mathrm{R}}$;*
- *$\mathrm{occ}(\theta) \subseteq \mathrm{occ}(\varphi_{\mathrm{L}}) \cap \mathrm{occ}(\varphi_{\mathrm{R}})$;*
- *$\mathrm{free}(\theta) \subseteq \mathrm{free}(\varphi_{\mathrm{L}}) \cap \mathrm{free}(\varphi_{\mathrm{R}})$;*
- *$\mathrm{con}(\theta) \subseteq \mathrm{con}(\varphi_{\mathrm{L}}) \cup \mathrm{con}(\varphi_{\mathrm{R}})$;*
- *if φ_{L} and φ_{R} are $\mathbb{U}$-relativized, for $\mathbb{U}$ a distinguished set of unary relations from $\sigma_{\mathrm{L}} \cup \sigma_{\mathrm{R}}$, then θ is $\mathbb{U}$-relativized (when treating $\top$ and $\bot$ as atomic formulas);*
- *the DAG representation of θ is of size at most $2^{2^{p(n)}}$ for some polynomial function p independent of φ_{L} and φ_{R} (and this can be improved to a size of at most $2^{p(n)}$ when φ_{L} and φ_{R} are in* GF *without equality and the bound on the arity of relations is fixed).*

Note that using this interpolation method, we can also construct an *FO* interpolant that can be represented by a doubly-exponential sized formula – e.g. by applying the theorem above, and then converting efficiently from a DAG representation to a formula [3].

Beth definability. As a corollary of Theorem 9, we get doubly exponential bounds on the DAG size of explicit GNF definitions coming from Projective Beth Definability. The argument is standard (see [13]).

6. Lower Bounds

We have matching lower bounds for the interpolation results stated in Theorem 9.

Theorem 10. *There is a polynomial function p and a family of* UNF *sentences $\chi_n \models \chi'_n$ without equality such that $|\chi_n| + |\chi'_n| \leq p(n)$, and there is no* GNF *interpolant θ_n for $\chi_n \models \chi'_n$ with size at most $2^{2^n - 1}$, even when θ_n is represented via a DAG.*

The sentences χ_n and χ'_n describe certain *indexed AND/OR trees*, finite trees where the leaves are labelled with 0 or 1, and other nodes are labelled with AND or OR. Moreover, each node is indexed by a tree that describes the depth of the node (up to depth 2^{2^n}).

Formally, an indexed AND/OR tree is a structure over signature ρ_n consisting of the following predicates: 1. Unary predicates $\mathrm{IndValOne}$, $\mathrm{IndDepth}_0, \ldots, \mathrm{IndDepth}_n$, and a ternary connective $\mathrm{IndChild}$ describe the index of one node. Informally, $\mathrm{IndDepth}_0, \ldots, \mathrm{IndDepth}_n$ and $\mathrm{IndChild}$ describe the structure of a binary tree of depth n (with $\mathrm{IndDepth}_0$ labelling the root), while $\mathrm{IndValOne}$ distinguishes the leaves of this tree with value one. Each such tree can be assigned a value, when viewing the leaves as 2^n bits of a binary number. 2. Unary predicates $\mathrm{InputOne}$ and IsOr along with binary predicate $\mathrm{AOChild}$ describe the AND/OR structure, with IsOr representing an OR node and $\mathrm{InputOne}$ distinguishing the inputs on the leaves that have value one.

The signatures of χ_n and χ'_n extend ρ_n using two copies of a predicate $\mathrm{AOValOne}$ for the intermediate value of the circuit computation at each node. Hence ρ_n will be the common signature for the interpolant.

Roughly speaking, the formula χ_n asserts

> there is an indexed AND/OR tree of depth 2^{2^n}, and the value (using unary relation $\mathrm{AOValOne}$) of any AND/OR tree in the structure is 1.

Likewise, χ'_n asserts

> the value (using unary relation $\mathrm{AOValOne}'$) of any AND/OR tree in the structure is 1.

We have $\chi_n \models \chi'_n$, and there is a polynomial function p such that for all n, $|\chi_n| + |\chi'_n| \leq p(n)$.

GN bisimulation game. In order to prove that there is no "small" GNF interpolant for $\chi_n \models \chi'_n$, we use a variation of GN bisimulations [5].

A position in the k-width version of the GN bisimulation game between $\mathfrak{A}$ and $\mathfrak{B}$ (relative to some signature σ) is a *partial rigid homomorphism* f from $\mathfrak{A}$ to $\mathfrak{B}$ or vice versa with $|\operatorname{dom}(f)| \leq k$. This means that f is a partial homomorphism and for any guarded tuple $\boldsymbol{c} \in \operatorname{dom}(f)$, $f \restriction_{\boldsymbol{c}}$ (the restriction of f to $\boldsymbol{c}$) is a partial isomorphism. We say the *active structure* is the structure containing $\operatorname{dom}(f)$.

Starting in position f, one round of the game consists of the following: (i) Spoiler can restrict to some subset $\boldsymbol{c} \subseteq \operatorname{dom}(f)$, and then the game proceeds from position $f \restriction_{\boldsymbol{c}}$; or (ii) if $\operatorname{dom}(f)$ is a guarded tuple $\boldsymbol{c}$, then Spoiler can choose to switch structures, and the game proceeds from position $(f \restriction_{\boldsymbol{c}})^{-1}$; or (iii) if $|\operatorname{dom}(f)| < k$, Spoiler can select an element c such that c is in the active structure but $c \notin \operatorname{dom}(f)$, Duplicator can choose d in the inactive structure such that $f[c \mapsto d]$ is a partial rigid homomorphism, and then the game proceeds from position $f[c \mapsto d]$ (Duplicator immediately loses if she is unable to choose a d such that $f[c \mapsto d]$ is a partial rigid homomorphism).

We write $\mathfrak{A}, \boldsymbol{a} \rightarrow^{k,m}_{\mathrm{gn}[\sigma]} \mathfrak{B}, \boldsymbol{b}$ if Duplicator has a winning strategy in the k-width m-round guarded negation bisimulation game relative to signature σ starting from a partial rigid homomorphism $g : \boldsymbol{a} \mapsto \boldsymbol{b}$. A winning strategy for Duplicator implies agreement between $\mathfrak{A}$ and $\mathfrak{B}$ on certain GNF formulas.

Proposition 11. *Assume $\mathfrak{A}, \boldsymbol{a} \rightarrow^{k,m}_{\mathrm{gn}[\sigma]} \mathfrak{B}, \boldsymbol{b}$, i.e. Duplicator has a winning strategy in the k-width, m-round GN-bisimulation game between $\mathfrak{A}$ and $\mathfrak{B}$ (relative to signature σ) starting from position $\boldsymbol{a} \mapsto \boldsymbol{b}$. If $\varphi(\boldsymbol{x})$ is a DAG representation of a* GNF *formula over σ such that $|\varphi| \leq m$, $\mathrm{width}(\varphi) \leq k$, and $\mathrm{free}(\varphi) = \boldsymbol{x}$, then $\mathfrak{A} \models \varphi(\boldsymbol{a})$ implies $\mathfrak{B} \models \varphi(\boldsymbol{b})$.*

Counterexamples. The specific family of counterexamples that we use to prove Theorem 10 are an inductively defined family of indexed AND/OR trees $\mathfrak{A}^n_i$ and $\mathfrak{B}^n_i$ of increasing depth. The value of $\mathfrak{A}^n_i$ (respectively, $\mathfrak{B}^n_i$) is 1 (respectively, 0), but the trees have some identical subtrees that make them hard to distinguish with "small" GNF sentences.

Roughly speaking, we show that for $i = 2^{2^n-1}$, Duplicator has a winning strategy in the i-round i-width GN bisimulation game relative to signature ρ_n between the indexed AND/OR trees $\mathfrak{A}^n_i$ and $\mathfrak{B}^n_i$ (starting from the empty homomorphism). By Proposition 11, this is enough to ensure that $\mathfrak{A}^n_i$ and $\mathfrak{B}^n_i$ cannot be distinguished by GNF sentences with a DAG representations of size at most i.

Supposing that there is a "small" (size at most 2^{2^n-1}) DAG representation of an interpolant θ_n for $\chi_n \models \chi'_n$ implies that $\mathfrak{A}^n_i \models \theta_n$ since $\mathfrak{A}^n_i \models \chi_n$ by construction. But $\mathfrak{A}^n_i$ and $\mathfrak{B}^n_i$ cannot be distinguished by such "small" formulas, so this means that $\mathfrak{B}^n_i \models \theta_n$, and hence $\mathfrak{B}^n_i \models \chi'_n$. This is impossible since $\mathfrak{B}^n_i$ is an indexed AND/OR tree with value 0, not value 1 as asserted by χ'_n.

We now describe in more detail this family of counterexamples $\mathfrak{A}^n_i$ and $\mathfrak{B}^n_i$ for all $i \leq 2^{2^n-1}$. For notational simplicity in the description below, we write $\mathfrak{A}_i$ and $\mathfrak{B}_i$ instead of $\mathfrak{A}^n_i$ and $\mathfrak{B}^n_i$.

The basic structure of these AND/OR trees is defined recursively as follows, where we write a to denote a single a-labelled node, and $a(T_1, T_2)$ to denote a tree where the root is labelled a and the children are the roots of trees T_1 and T_2:

- $\mathfrak{A}_0 := 1$ and $\mathfrak{B}_0 := 0$;
- for $i > 0$, $\mathfrak{A}_i := \mathrm{AND}(\mathrm{OR}(\mathfrak{B}_{i-1}, \mathfrak{A}_{i-1}), (\mathrm{OR}(\mathfrak{B}_{i-1}, \mathfrak{A}_{i-1}))$ and $\mathfrak{B}_i := \mathrm{AND}(\mathrm{OR}(\mathfrak{B}_{i-1}, \mathfrak{A}_{i-1}), (\mathrm{OR}(\mathfrak{B}_{i-1}, \mathfrak{B}_{i-1}))$.

Evaluating $\mathfrak{A}_i$ (respectively, $\mathfrak{B}_i$) gives value 1 (respectively, 0).

A separate index tree (binary tree of depth n) is then attached to each node in the trees described above, where the index correctly describes the depth of this node in the AND/OR tree.

We view these indexed AND/OR trees as structures over the signature $\rho_n = \sigma_n \cap \sigma'_n$. In particular, this means that there are no internal computations using $\mathrm{AOValOne}$ or $\mathrm{AOValOne}'$.

Let $h : \boldsymbol{b} \mapsto \boldsymbol{a}$ be a position (partial rigid homomorphism) in the k-width GN bisimulation game between some $\mathfrak{B}_i$ and $\mathfrak{A}_i$. We can assume that h is restricted to elements in the AND/OR tree (rather than the elements in the index trees). This is without loss of generality since h can always be chosen to preserve the depth of positions, and the local structure of the index trees is identical in $\mathfrak{A}_i$ and $\mathfrak{B}_i$.

As mentioned earlier, we want to show that Duplicator has a winning strategy in certain GN bisimulation games between $\mathfrak{B}_i$ and $\mathfrak{A}_i$. In order to prove this inductively, we need to introduce the notion of "safe" positions h, which represent good starting positions for Duplicator.

We first introduce some additional notation. We write $\mathfrak{B}_i(b)$ for the restriction of $\mathfrak{B}_i$ to the subtree rooted at b. For $i > 0$, let C_i (respectively, D_i) denote the set of nodes in $\mathfrak{A}_i$ (respectively, $\mathfrak{B}_i$) at depth 2 (these are the roots of the $\mathfrak{A}_{i-1}$ and $\mathfrak{B}_{i-1}$ subtrees of $\mathfrak{A}_i$ and $\mathfrak{B}_i$). We write $\boldsymbol{b} \restriction_b$ for the restriction of $\boldsymbol{b}$ to those elements appearing in $\mathfrak{B}_i(b)$.

We now define inductively what it means for h to be (i, k, m)-*safe*. We say a strategy in the m-round game is (i, k, m)-safe if it uses only (i, k, m')-safe positions where m' is the number of remaining moves in the game. We say $h : \boldsymbol{b} \mapsto \boldsymbol{a}$ is (i, k, m)-safe if it satisfies the following properties:

(S1) for all $b \in \boldsymbol{b}$, $\mathrm{depth}(b) = \mathrm{depth}(h(b))$;

(S2) if $i > 0$, then there exists a mapping $f : D_i \rightarrow C_i$ such that for all $b, b' \in D_i$,
- if b, b' are siblings, then $f(b), f(b')$ are siblings;
- if $d \in \boldsymbol{b}$ is the parent of b, then $h(d)$ is the parent of $f(b)$;
- if $d \in \boldsymbol{b}$ appears in $\mathfrak{B}_i(b)$, then $h(d)$ appears in $\mathfrak{A}_i(f(b))$;
- $\mathfrak{B}_i(b), \boldsymbol{b} \restriction_b \rightarrow^{k,m}_{\mathrm{gn}[\rho_n]} \mathfrak{A}_i(f(b)), h(\boldsymbol{b} \restriction_b)$ via a strategy for Duplicator that is $(i-1, k, m)$-safe.

Likewise, we say a position $h : \boldsymbol{a} \mapsto \boldsymbol{b}$ is (i,k,m)-safe if the conditions (S1) and (S2) hold with $\mathfrak{B}_i$, $\boldsymbol{b}$, D_i exchanged with $\mathfrak{A}_i$, $\boldsymbol{a}$, C_i, respectively.

Observe that the empty partial rigid homomorphism h from $\mathfrak{B}_i$ to $\mathfrak{A}_i$ is (i,k,m)-safe: (S1) vacuously holds, and (S2) holds since (if $i > 0$) we can choose f such that sibling relationships are preserved, and $\mathfrak{A}_{i-1}$ (respectively, $\mathfrak{B}_{i-1}$) subtrees in $\mathfrak{B}_i$ are mapped to $\mathfrak{A}_{i-1}$ (respectively, $\mathfrak{B}_{i-1}$) subtrees in $\mathfrak{A}_i$. Similarly for the empty partial rigid homomorphism from $\mathfrak{A}_i$ to $\mathfrak{B}_i$.

More importantly, we can show that Duplicator has a winning (i,k,m)-safe strategy when starting from an (i,k,m)-safe position.

Lemma 12. *For all k, $i \leq 2^{2^n-1}$, and $m \leq i$,*
- *if $\boldsymbol{b} \mapsto \boldsymbol{a}$ is (i,k,m)-safe then $\mathfrak{B}_i, \boldsymbol{b} \rightarrow^{k,m}_{\mathrm{gn}[\rho_n]} \mathfrak{A}_i, \boldsymbol{a}$ via a strategy for Duplicator that is (i,k,m)-safe;*
- *if $\boldsymbol{a} \mapsto \boldsymbol{b}$ is (i,k,m)-safe then $\mathfrak{A}_i, \boldsymbol{a} \rightarrow^{k,m}_{\mathrm{gn}[\rho_n]} \mathfrak{B}_i, \boldsymbol{b}$ via a strategy for Duplicator that is (i,k,m)-safe.*

This implies that for $i = 2^{2^n-1}$, Duplicator has a winning strategy in the i-round i-width GN bisimulation game relative to signature ρ_n between the indexed AND/OR trees $\mathfrak{A}^n_i$ and $\mathfrak{B}^n_i$ (starting from the empty homomorphism). As argued above, this can be combined with Proposition 11 to prove the lower bound stated in Theorem 10.

Guarded fragment. A variant of this construction inspired by [12] can be used to get a lower bound for the size of interpolants of GF validities.

Theorem 13. *There is a polynomial function p and a family of* GF *sentences $\chi_n \models \chi'_n$ with relations of unbounded arity (respectively, bounded arity) such that $|\chi_n| + |\chi'_n| \leq p(n)$, and there is no* GNF *interpolant θ_n for $\chi_n \models \chi'_n$ with size at most 2^{2^n} (respectively, size at most 2^n), even when θ_n is represented by a DAG.*

Beth definability. A similar argument can be used to prove doubly-exponential lower bounds on the DAG-size of GNF explicit definitions coming from Beth Definability.

7. Preservation Theorems

In this section, we consider some preservation theorems for GNF and GF that follow from the constructive interpolation results of Theorem 9. We remark that the results in this section hold when preservation and equivalence is considered over finite structures only, as well as over general structures.

We first consider the analog of the Łoś-Tarski theorem. A formula φ is *preserved under extensions* if $\mathfrak{A} \subseteq \mathfrak{B}$ and $\mathfrak{A} \models \varphi$ together imply $\mathfrak{B} \models \varphi$ (see, e.g., [9, 22]). It was shown in [7] that every GNF formula preserved under extensions has an equivalent *existential* GNF *formula*, which is a GNF formula where no quantifier is in the scope of a negation symbol. We can use relativized GNF interpolation to get a bound on the size of this existential formula. This contrasts with the result of Dawar et. al. [9] that for general first-order formulas there is no elementary bound for the Łoś-Tarski theorem.

Corollary 14 (Analog of Łoś-Tarski Theorem)**.** *Let φ be a* GNF *formula over signature σ with $|\varphi| = n$. If φ is preserved under extensions, then we can construct a DAG representation of an equivalent existential* GNF *formula φ' such that $|\varphi'| \leq 2^{2^{p(n)}}$ for some polynomial function p independent of φ.*

Proof. We follow Otto [22]. Let $\mathbb{U} := \{U^1, U^2\}$ for distinct unary predicates U^1, U^2 not occurring in σ.

Since φ is preserved under extensions, we have

$$\neg\exists y.(U^1 y \wedge \neg U^2 y) \wedge \bigwedge_{t \in \boldsymbol{x} \cup \boldsymbol{e}} U^1 t \wedge \varphi^1 \models \varphi^2$$

where $\mathrm{free}(\varphi) = \boldsymbol{x}$, $\mathrm{con}(\varphi) = \boldsymbol{e}$, and φ^i denotes the relativization of φ to U^i. For readability, here and in the similar arguments below, the equality guard for the negation of the first conjunct in the antecedent has been omitted.

Let N be the size of this validity, which is linear in n. By Theorem 9, the previous GNF validity admits a $\mathbb{U}$-relativized GNF interpolant θ over signature $\sigma \cup \mathbb{U}$ with a DAG representation of size at most $2^{2^{p(N)}}$ for some polynomial function p independent of φ.

Since U^1 does not occur on both sides of the validity, U^1 cannot appear at all in θ. Since U^2 does not occur negatively on the left hand side of the original validity, U^2 cannot appear negatively in θ. This means that every existential quantifier in θ must occur positively (since every quantifier must be relativized by U^2, and U^2 can only occur positively).

Replacing every occurrence of some $U^2 z$ in θ by $z = z$ results in a GNF formula φ' over σ that is equivalent to φ, and such that any quantifier occurs within the scope of an even number of negations. By pushing the negations to the atomic level (and potentially duplicating guards to cover any negated atoms), we have a doubly-exponential size DAG representation of an existential GNF formula. □

Similarly, we can use the Lyndon interpolation result to prove a relationship between monotonicity and positivity. We say a formula φ is *monotone* if it is preserved when a structure is modified via adding tuples to the interpretation of some relation. We say a formula φ is *positive* if all relation symbols (except possibly equality) occur positively (within the scope of an even number of negations). A formula is *domain-independent* (see [1]) if it depends only on the interpretations of relation symbols.

Corollary 15 (Monotone=Positive)**.** *Let φ be a* GNF *formula over signature σ with $|\varphi| = n$ such that φ is monotone. Then we can construct a DAG representation of an equivalent positive* GNF *formula φ' such that $|\varphi'| \leq 2^{2^{p(n)}}$ for some polynomial function p independent of φ. Furthermore, if φ is domain-independent, then φ' can be taken to be positive and in existential* GNF*.*

Proof. Since φ is monotone, we have

$$\bigwedge_{R \in \sigma} \neg\exists \boldsymbol{y}.(R^1 \boldsymbol{y} \wedge \neg R^2 \boldsymbol{y}) \wedge \varphi^1 \models \varphi^2 \qquad (1)$$

where φ^i denotes φ with each $R \in \sigma$ replaced with R^i. Applying Theorem 9 yields an interpolant only using relations of the form R^2 for $R \in \sigma$. Moreover, these relations must occur positively. Replacing each R^2 with R, yields φ'.

The above argument is modified in the domain-independent case by using additional predicates for the domains. More precisely, we relativize the antecedent and consequent of the entailment (1) by fresh unary predicates U^1 and U^2, respectively; and we add $\bigwedge_{R \in \sigma; i,j \in \{1,2\}} \neg\exists \boldsymbol{x}.(\bigwedge_{x \in \boldsymbol{x}} U^i x \wedge R^1 \boldsymbol{x} \wedge \neg \bigwedge_{x \in \boldsymbol{x}} U^j x)$ as conjunct to the antecedent. Applying Theorem 9 as before, and replacing each R^2 with R, as well as $U^2 z$ by $z = z$, yields the desired formula φ'. □

Guarded fragment. We now turn to effective preservation for GF. Recall that Andréka et. al. [2] prove a Łoś-Tarski theorem for GF, but with no bound. Here we state a constructive version.

Theorem 16. *Let φ be a* GF *formula over signature σ with $|\varphi| = n$. If φ is preserved under extensions, then we can construct an equivalent existential* GF *formula φ' such that $|\varphi'| \leq 2^{2^{2^{2^{p(n)}}}}$ for a polynomial function p independent of φ.*

Proof. We sketch the argument for the boolean case. The proof proceeds by applying Corollary 14 to obtain a DAG representation of some GNF existential χ' equivalent to φ of doubly-exponential size. We can convert χ' to χ'' that is a union of χ''_i, where each χ''_i is a "GNCQ" – a conjunctive query with atomic negation, where every negated atom is guarded. Further, each GNCQ χ''_i within χ'' is at most exponential in the size of χ', and the overall size of χ'' (as a formula) is at most doubly exponential in the size of χ'. Let A_i be the set of queries in GF of the form $\exists \boldsymbol{y}. \bigwedge_j A_{ij}$, where $\bigwedge A_{ij}$ is obtained by identifying variables in χ''_i and then adding on at most $2^{|\varphi|}$ additional positive atoms. We claim that χ' is equivalent to $\bigvee_i \bigvee_{Q \in A_i} Q$. Since the set of indices i is doubly-exponential in the size of χ' and each A_i has size doubly-exponential in χ', this gives the desired bound. Clearly each query in A_i implies χ''_i and hence implies φ. On the other hand, consider a model M satisfying φ. Given any model M there is another structure M^* agreeing with M on all GF sentences that has a guarded tree decomposition. Indeed, the guarded unravelling of M (see, e.g., [2]) gives such an M^*. M^* must satisfy φ and therefore must satisfy some χ''_i via a homomorphism h. Let C be the image of h and Q_C be the conjunctive query that describes the restriction of M^* to the bags containing C – such a query can be written in GF by constructing it inductively from the leaves of the tree up. The size of Q_C is at most the size of χ''_i times the maximal size of bags – the latter being a factor which is at most exponential in the size of the schema. Note that every element $c \in C$ corresponds to exactly one variable y_c of Q_C. We let Q'_C extend Q_C be adding, for every negated atom in χ''_i, the result of replacing each variable x_i by $y_{h(x_i)}$. Then Q'_C also holds in M^* and $Q'_C \in$ GF. Hence Q'_C holds in M and is in A_i, which completes the argument. □

We get similar bounds for characterizing monotone domain-independent GF formulas.

Theorem 17. *Let φ be a* GF *formula over signature σ with $|\varphi| = n$. If φ is monotone and domain-independent then we can construct an equivalent formula φ' that is positive and is in existential* GF*, and such that $|\varphi'| \leq 2^{2^{2^{2^{p(n)}}}}$ for a polynomial function p independent of φ.*

Proof. Again we focus on the case where φ is boolean.

By Corollary 15, φ is equivalent to a positive and existential GNF formula χ' whose DAG representation has doubly-exponential size. As above, we can convert this to a union $\chi'' = \bigvee_i \chi''_i$, with χ''_i in this case being a conjunctive query, but possibly including guarded inequalities. As before, this results in a doubly-exponential blow up in the overall size, but only a singly-exponential blow up in the size of each CQ.

Again letting A_i be the set of queries in GF of the form $\exists \boldsymbol{y}. \bigwedge_j A_{ij}$, where $\bigwedge A_{ij}$ is obtained by identifying variables in χ''_i and then adding on at most $2^{|\varphi|}$ additional positive atoms, we can show that χ' is equivalent to $\bigvee_i \bigvee_{Q \in A_i} Q$. □

8. Conclusions

We have shown that GNF is a rich logic that admits interpolation and preservation theorems with elementary complexity, in sharp contrast to full first-order logic. In doing this, we developed the mosaic method for use in interpolation. In future work, we will look at revisiting the applications of GNF interpolation to databases (as in [6]) in light of these results.

The bounds for interpolation are tight for GNF and UNF. We leave open the question of tightness of our bounds for preservation, as well as the tightness of interpolation for GF with restricted arity and equality. We also do not know the status of many other preservation theorems for GNF, such as preservation under surjective homomorphism, and the status of monotone=positive for non-domain independent GF formulas.

References

[1] S. Abiteboul, R. Hull, and V. Vianu. *Foundations of Databases*. Addison-Wesley, 1995.

[2] H. Andréka, I. Németi, and J. van Benthem. Modal languages and bounded fragments of predicate logic. *J. Phil. Logic*, 27(3):217–274, 1998.

[3] J. Avigad. Eliminating definitions and skolem functions in first-order logic. *ACM Trans. Comput. Log.*, 4(3), 2003.

[4] F. Baader, D. Calvanese, D. L. McGuinness, D. Nardi, and P. F. Patel-Schneider, editors. *The description logic handbook*. Cambridge, 2003.

[5] V. Bárány, B. ten Cate, and L. Segoufin. Guarded negation. In *ICALP*, 2011.

[6] V. Bárány, M. Benedikt, and P. Bourhis. Access patterns and integrity constraints revisited. In *ICDT*, 2013.

[7] V. Bárány, M. Benedikt, and B. ten Cate. Rewriting guarded negation queries. In *MFCS*, 2013.

[8] W. Craig. Three uses of the Herbrand-Gentzen theorem in relating model theory and proof theory. *JSL*, 22(3):269–285, 1957.

[9] A. Dawar, M. Grohe, S. Kreutzer, and N. Schweikardt. Model theory makes formulas large. In *ICALP*. 2007.

[10] M. Fitting. *First-order logic and automated reasoning*. Springer, 1996.

[11] H. Friedman. The complexity of explicit definitions. *Advances in Math.*, 20(1):18 – 29, 1976.

[12] E. Grädel. On the restraining power of guards. *JSL*, 64(4):1719–1742, 1999.

[13] E. Hoogland. *Definability and interpolation: model-theoretic investigations*. PhD thesis, University of Amsterdam, 2000.

[14] E. Hoogland, M. Marx, and M. Otto. Beth definability for the guarded fragment. In *LPAR*, 1999.

[15] R. C. Lyndon. An interpolation theorem in the predicate calculus. *Pacific J. Math.*, 9:129–142, 1959. ISSN 0030-8730.

[16] M. Marx. Queries determined by views: pack your views. In *PODS*, 2007.

[17] M. Marx, S. Mikulas, and M. Reynolds. The mosaic method for temporal logics. In *TABLEAUX*, 2000.

[18] K. L. McMillan. Applications of Craig Interpolation to model checking. In *CSL*, 2004.

[19] S. Mikulás. Taming first-order logic. *Journal of the IGPL*, 6(2):305–316, 1998.

[20] A. Nash, L. Segoufin, and V. Vianu. Views and queries: Determinacy and rewriting. *ACM Trans. Database Syst.*, 35(3), 2010.

[21] I. Németi. *Free Algebras and Decidability in Algebraic Logic*. PhD thesis, Hungarian Academy of Sciences, Budapest, 1986.

[22] M. Otto. An interpolation theorem. *BSL*, 6(4):447–462, 2000.

[23] B. ten Cate and L. Segoufin. Unary negation. In *STACS*, 2011.

[24] B. ten Cate, E. Franconi, and I. Seylan. Beth definability in expressive description logics. In *IJCAI*, 2011.

A. Mosaic Method and Interpolation using GF Mosaics

In this section, we present details of the GF mosaics defined in Section 4. The purpose of this section is to demonstrate (in the simpler context of GF, rather than GNF) the proof techniques used to prove soundness and completeness of the mosaic method, and how to extract GNF interpolants from it.

A.1 Soundness and completeness of GF mosaic method (Proof of Theorem 3)

We now give detailed proofs of the soundness and completeness of the GF mosaic method.

We must first fix some additional definitions.

We say a formula $\widetilde{X} : \psi$ is X-*safe* in τ if $\mathrm{par}(\psi)$ is X-cguarded in τ.

Fix parameters $\boldsymbol{c}$. Given a structure $\mathfrak{M}$ over $\sigma_{\mathrm{L}} \cup \sigma_{\mathrm{R}} \cup \boldsymbol{c}$ such that $\boldsymbol{c}$ is guarded by some X-atom, let $\upsilon_Y^{\mathfrak{M}}$ be the unique Y-mosaic over $\boldsymbol{c}$ such that $\upsilon_Y^{\mathfrak{M}} := \{Y : \psi \mid Y : \psi \in \mathrm{cl}_Y(\boldsymbol{c}) \text{ and } \mathfrak{M} \models \psi\}$. The X-mosaic $\tau := \upsilon_X^{\mathfrak{M}} \cup \left\{ \widetilde{X} : \psi \mid \widetilde{X} : \psi \in \upsilon_{\widetilde{X}}^{\mathfrak{M}} \text{ and } \widetilde{X} : \psi \text{ is } X\text{-safe in } \upsilon_X^{\mathfrak{M}} \right\}$ is called the X-*mosaic* τ *of* $\boldsymbol{c}$ *in* $\mathfrak{M}$.

We say an X-mosaic τ over $\boldsymbol{c}$ is *realizable* if there is some structure $\mathfrak{M}$ such that τ is the X-mosaic of $\boldsymbol{c}$ in $\mathfrak{M}$.

For convenience in the following proofs, we also restate here the definition of linked GF mosaics (updated to use the notion of Y-safety). Let τ be a GF X-mosaic over parameters $\boldsymbol{c}$ with $Y^+ : \eta \in \tau$ for η of the form $\exists \boldsymbol{x}.\psi(\boldsymbol{xb})$. We say τ' is linked with τ via $Y^+ : \eta$ if

(E1) τ' is a GF Y-mosaic over parameters $\boldsymbol{ab}$ with $\boldsymbol{a} \cap \boldsymbol{c} = \emptyset$;
(E2) $Y : \alpha(\boldsymbol{xb})[\boldsymbol{a}/\boldsymbol{x}] \in \tau'$ and $Y : \psi(\boldsymbol{xb})[\boldsymbol{a}/\boldsymbol{x}] \in \tau'$;
(E3) for all $Y : \psi'(\boldsymbol{b}) \in \tau$, $Y : \psi'(\boldsymbol{b}) \in \tau'$;
(E4) for all $\widetilde{Y} : \psi'(\boldsymbol{b}) \in \tau$ such that $\widetilde{Y} : \psi'(\boldsymbol{b})$ is Y-safe in τ', $\widetilde{Y} : \psi'(\boldsymbol{b}) \in \tau'$;
(E5) for all $Z : \psi'(\boldsymbol{b}) \in \tau'$, $Z : \psi'(\boldsymbol{b}) \in \tau$.

We start with two propositions relating realizability to our coherency requirements for GF mosaics (this will be used for soundness).

Proposition 18. *If τ is realizable, then τ is internally consistent.*

Proof. Let $\mathfrak{M}$ be a structure realizing the X-mosaic $\tau(\boldsymbol{c})$. By definition, $Z : \psi \in \tau$ implies $\mathfrak{M} \models \psi$.

This means it is not possible for there to be some $Y : \alpha$ and $Z : \neg\alpha$ in τ.

The downward closure properties must also be satisfied. For instance, consider the last downward closure property. Assume $Y : \forall \boldsymbol{x}.\alpha(\boldsymbol{xb}) \to \psi(\boldsymbol{xb}) \in \tau$, and there is $\boldsymbol{a} \subseteq \boldsymbol{c} \cup \boldsymbol{e}$ such that $Z : \alpha(\boldsymbol{xb})[\boldsymbol{a}/\boldsymbol{x}] \in \tau$. We know that $\mathfrak{M} \models \forall \boldsymbol{x}.\alpha(\boldsymbol{xb}) \to \psi(\boldsymbol{xb})$ and $\mathfrak{M} \models \alpha(\boldsymbol{xb})[\boldsymbol{a}/\boldsymbol{x}]$, so $\mathfrak{M} \models \psi(\boldsymbol{xb})[\boldsymbol{a}/\boldsymbol{x}]$. If $Y = X$, then this is enough to ensure that $Y : \psi(\boldsymbol{xb})[\boldsymbol{a}/\boldsymbol{x}] \in \tau$. If $Y = \widetilde{X}$, then it remains to show that $Y : \psi(\boldsymbol{xb})[\boldsymbol{a}/\boldsymbol{x}]$ is X-safe in $\upsilon_X^{\mathfrak{M}}$.

- If $Y = Z$, then $Z : \alpha(\boldsymbol{xb})[\boldsymbol{a}/\boldsymbol{x}] \in \tau$ implies that $\mathrm{par}(\alpha(\boldsymbol{xb})[\boldsymbol{a}/\boldsymbol{x}])$ is X-cguarded in $\upsilon_X^{\mathfrak{M}}$, so $Y : \psi(\boldsymbol{xb})[\boldsymbol{a}/\boldsymbol{x}]$ is X-safe in $\upsilon_X^{\mathfrak{M}}$.
- If $Y = \widetilde{Z}$, then α appears positively in an X-formula (since $Z = X$) and negatively in a $\widetilde{X}$-formula (since $Y = \widetilde{X}$). Making use of the fact that occurrences are preserved in moving to mosaics, this means that α is an X-catom, so $\mathrm{par}(\psi(\boldsymbol{xb})[\boldsymbol{a}/\boldsymbol{x}])$ is X-cguarded by $\alpha(\boldsymbol{xb})[\boldsymbol{a}/\boldsymbol{x}] \in \upsilon_X^{\mathfrak{M}}$, and $Y : \psi(\boldsymbol{xb})[\boldsymbol{a}/\boldsymbol{x}]$ is X-safe in $\upsilon_X^{\mathfrak{M}}$.

It is straightforward to check the other downward closure properties. □

Proposition 19. *If υ is realizable, then there is a saturated set $\mathcal{S} \subseteq \mathcal{M}$ of realizable mosaics that includes υ.*

Proof. We inductively construct sets $\mathcal{S}_i$, ensuring at each stage i that $\mathcal{S}_i$ contains υ and other realizable mosaics, and that any existential requirement in $\mathcal{S}_{i-1}$ is fulfilled in $\mathcal{S}_i$. This is enough to ensure that the set $\mathcal{S} := \bigcup \mathcal{S}_i$ is a saturated set of realizable mosaics that includes υ.

At stage $i = 1$, we set $\mathcal{S}_1 := \{\upsilon\}$, which is realizable by assumption.

At stage $i > 1$, if every existential requirement in $\mathcal{S}_{i-1}$ is fulfilled, then set $\mathcal{S}_i := \mathcal{S}_{i-1}$.

Otherwise, consider some X-mosaic $\tau(\boldsymbol{c})$ in $\mathcal{S}_{i-1}$ with $Y : \eta \in \tau$ for η of the form $\exists \boldsymbol{x}.\alpha(\boldsymbol{xb}) \wedge \psi(\boldsymbol{xb})$ for which there is no $\boldsymbol{a} \subseteq \boldsymbol{c} \cup \boldsymbol{e}$ with $Y : (\alpha(\boldsymbol{xb}) \wedge \psi(\boldsymbol{xb}))[\boldsymbol{a}/\boldsymbol{x}] \in \tau$, and for all τ' such that $\tau \to_{Y:\eta} \tau'$, $\tau' \notin \mathcal{S}_{i-1}$.

Let $\mathfrak{M}$ be a structure realizing τ. Since $\mathfrak{M} \models \tau$, there must be some elements $\boldsymbol{a}' \in M$, parameters $\boldsymbol{a} \in P$ such that $\boldsymbol{a} \cap \boldsymbol{c} = \emptyset$, and an expansion $\mathfrak{M}'$ of $\mathfrak{M}$ with the interpretation $\boldsymbol{a}^{\mathfrak{M}'} := \boldsymbol{a}'$ such that $\mathfrak{M}' \models \alpha(\boldsymbol{ab}) \wedge \psi(\boldsymbol{ab})$. This is possible since $|\boldsymbol{a}| \leq \mathrm{width}(\varphi_{\mathrm{L}} \wedge \neg\varphi_{\mathrm{R}})$ and $|P| = 2 \cdot \mathrm{width}(\varphi_{\mathrm{L}} \wedge \neg\varphi_{\mathrm{R}})$.

Let τ' be the Y-mosaic of $\boldsymbol{ab}$ in $\mathfrak{M}'$. This means τ' is realizable.

We also claim that $\tau \to_{Y:\eta} \tau'$. (E1) is satisfied by construction. By the choice of $\boldsymbol{a}^{\mathfrak{M}'}$, we have $\mathfrak{M}' \models \alpha(\boldsymbol{ab}) \wedge \psi(\boldsymbol{ab})$, so $Y : \alpha(\boldsymbol{ab}) \wedge \psi(\boldsymbol{ab}) \in \tau'$ as required by (E2).

We must now check that formulas with shared parameters $\boldsymbol{b}$ are preserved between τ and τ' as described by (E3)–(E5).

(E3) and (E4) follow from the definition of realizable mosaics, and the fact that $\mathfrak{M}$ and $\mathfrak{M}'$ are identical with respect to formulas that use only parameters from $\boldsymbol{b}$.

It remains to check (E5). Assume there is some $Z : \psi'(\boldsymbol{b}) \in \tau'$. We have $\mathfrak{M}' \models \psi'$ since τ' is realized by $\mathfrak{M}'$. Recall that τ is an X-mosaic, τ' is a Y-mosaic, and ψ' is a Z-formula. We proceed by cases depending on the provenances X, Y, and Z. The result is immediate if $Z = X$. Otherwise, if $Z = \widetilde{X}$, it remains to show that $Z : \psi'(\boldsymbol{b})$ is is X-safe in τ.

- Assume $Z = \widetilde{X}$ and $Y = X$. Since $Z = \widetilde{X} = \widetilde{Y}$, $Z : \psi' \in \tau'$ implies $\mathrm{par}(\psi') = \boldsymbol{b}' \subseteq \boldsymbol{b}$ is Y-cguarded in τ' by some $Y : \exists \boldsymbol{y}.\beta(\boldsymbol{yb}') \in \tau'$ where β is a Y-catom. But this means $\mathfrak{M}' \models \exists \boldsymbol{y}.\beta(\boldsymbol{yb}')$ and $\mathfrak{M} \models \exists \boldsymbol{y}.\beta(\boldsymbol{yb}')$. Moreover, $X = Y$, so we can conclude that $Y : \exists \boldsymbol{y}.\beta(\boldsymbol{yb}') \in \tau$ and hence $Z : \psi' \in \tau$ is X-safe in τ.
- Assume $Z = \widetilde{X}$ and $Y = \widetilde{X}$. Since $Y : \eta$ in τ uses $\boldsymbol{b}$ and $Y = \widetilde{X}$, we know that $\boldsymbol{b}$ is X-cguarded in τ. But this means $Z : \psi'(\boldsymbol{b})$ is X-safe in τ.

We have constructed a realizable Y-mosaic τ' such that $\tau \to_{Y:\eta} \tau'$ fulfils the existential requirement $Y : \eta \in \tau$, so we add τ' to $\mathcal{S}_i$. Repeating this procedure for all such τ (and all such $Y : \eta \in \tau$) in $\mathcal{S}_{i-1}$, we will end up with a set $\mathcal{S}_i$ of realizable mosaics such that any existential requirement in $\mathcal{S}_{i-1}$ is fulfilled in $\mathcal{S}_i$. $\square$

We now aim towards proving a proposition useful for completeness (a sort of converse of Proposition 19). In order to do this, we first prove two lemmas about the transfer of formulas between linked mosaics.

Lemma 20 (Backwards transfer property)**.** *Let $\tau_0 \to \cdots \to \tau_k$ be linked mosaics such that every τ_i includes parameters $\boldsymbol{b}$. If $X : \psi(\boldsymbol{b}) \in \tau_k$, then $X : \psi(\boldsymbol{b}) \in \tau_0$.*

Proof. We proceed by induction on k. The base case for $k = 0$ is immediate.

Otherwise, assume $k > 0$. Consider $\tau_{k-1} \to_{Y:\eta} \tau_k$. By definition (E5), $X : \psi(\boldsymbol{b}) \in \tau_k$ implies $X : \psi(\boldsymbol{b}) \in \tau_{k-1}$ (since $\boldsymbol{b}$ is shared by τ_{k-1} and τ_k), so we can apply the inductive hypothesis to $\tau_0 \to \cdots \to \tau_{k-1}$ to get the result. $\square$

Lemma 21 (Forward universal transfer property)**.** *Let $\tau_0 \to \cdots \to \tau_k$ be linked mosaics such that every τ_i includes parameters $\boldsymbol{b}$. If $Z : \forall \boldsymbol{x}.\alpha(\boldsymbol{xb}) \to \psi(\boldsymbol{xb}) \in \tau_0$ and $Z' : \exists \boldsymbol{x}.\alpha(\boldsymbol{xb}) \in \tau_k$, then $Z : \forall \boldsymbol{x}.\alpha(\boldsymbol{xb}) \to \psi(\boldsymbol{xb}) \in \tau_k$.*

Proof. We proceed by induction on k. If $k = 0$, then the result is immediate.

Assume $k > 0$, with τ_0 an X-mosaic and τ_1 a Y-mosaic. In order to apply the inductive hypothesis, we need to show that $Z : \forall \boldsymbol{x}.\alpha(\boldsymbol{xb}) \to \psi(\boldsymbol{xb}) \in \tau_1$.

- If $Z = Y$, then $Z : \forall \boldsymbol{x}.\alpha(\boldsymbol{xb}) \to \psi(\boldsymbol{xb}) \in \tau_1$ by (E3).
- If $Z = \widetilde{Y}$ and $Y = X$, then $\boldsymbol{b}$ is X-cguarded in τ by some $X : \exists \boldsymbol{y}.\beta(\boldsymbol{yb}) \in \tau_0$. By (E3), $X : \exists \boldsymbol{y}.\beta(\boldsymbol{yb}) \in \tau_1$, so $Z : \forall \boldsymbol{x}.\alpha(\boldsymbol{xb}) \to \psi(\boldsymbol{xb})$ is X-safe in τ_1, and hence in τ_1 by (E4).
- If $Z = \widetilde{Y}$ and $Y = \widetilde{X}$, then $Z' : \exists \boldsymbol{x}.\alpha(\boldsymbol{xb}) \in \tau_1$ by the backwards transfer property (Lemma 20). If $Z' = \widetilde{Y}$, then this means $\boldsymbol{b}$ must be Y-cguarded in τ_1. If $Z' = Y$, then $Z' : \exists \boldsymbol{x}.\alpha(\boldsymbol{xb}) \in \tau_1$ witnesses the fact that $\boldsymbol{b}$ is Y-cguarded in τ_1. In either case, we can apply (E4) to get $Z : \forall \boldsymbol{x}.\alpha(\boldsymbol{xb}) \to \psi(\boldsymbol{xb}) \in \tau_1$.

$\square$

Proposition 22. *Let $\mathcal{S} \subseteq \mathcal{M}$ be a saturated set of internally consistent mosaics. Then for all $\tau \in \mathcal{S}$, there is a structure $\mathfrak{M}$ such that $\mathfrak{M} \models \psi$ for all $X : \psi \in \tau$ with $\mathrm{par}(\psi) = \emptyset$.*

Proof. Let $\tau \in \mathcal{S}$ be a mosaic containing $\mathrm{L} : \varphi_{\mathrm{L}}, \mathrm{R} : \neg\varphi_{\mathrm{R}}$. Note that τ is internally consistent by the definition of $\mathcal{S}$. Our goal is to construct a model $\mathfrak{M}$ witnessing the satisfiability of $\varphi_{\mathrm{L}} \wedge \neg\varphi_{\mathrm{R}}$ using the mosaics in $\mathcal{S}$ as building blocks.

We do this by building inductively a tree decomposition T of $\mathfrak{M}$, together with a map f from nodes v in the tree decomposition to a mosaic in $\mathcal{S}$. (We also think of f as a map from the elements in a node v, to the parameters in the corresponding mosaic).

At stage 0, let T_0 consist of a single node v_0 containing a copy of the parameters in τ and constants $\boldsymbol{e}$, and set $f(v_0) = \tau$.

At stage i, consider the set

$$\begin{aligned} S = \{(v, Y : \eta) : \ & v \text{ is a leaf in } T_{i-1} \text{ and } Y : \eta(f(\boldsymbol{b})) \in f(v) \text{ is not fulfilled in } f(v) \\ & \text{for } \eta \text{ of the form } \exists \boldsymbol{x}.\alpha(\boldsymbol{xb}) \wedge \psi(\boldsymbol{x}f(\boldsymbol{b}))\}. \end{aligned}$$

For each $(v, Y : \eta) \in S$, we know that there is some $\tau' \in \mathcal{S}$ such that $f(v) \to_{Y:\eta} \tau'$ (since $\mathcal{S}$ is saturated). We construct T_i by extending T_{i-1} according to the following procedure: for each $(v, Y : \eta) \in S$, we construct a child v' of v with $T(v') = \boldsymbol{a} \cup \boldsymbol{b} \cup \boldsymbol{e}$ for new elements $\boldsymbol{a}$, and define $f(v') = \tau'$ such that $Y : \alpha(f(\boldsymbol{ab})) \wedge \psi(f(\boldsymbol{ab})) \in \tau'$.

Let T be the limit of this process, and let $\mathfrak{M}$ be the structure with elements $M = \bigcup_{v \in T} T(v)$, and such that $\mathfrak{M} \models \alpha(\boldsymbol{a})$ iff there exists v such that $\boldsymbol{a}$ appears in $T(v)$ and $\alpha(f(\boldsymbol{a})) \in f(v)$.

We must now show that for all v in T and for all $\boldsymbol{a}$ in $T(v)$ if $Z : \psi(f(\boldsymbol{a})) \in f(v)$, then $\mathfrak{M} \models \psi(\boldsymbol{a})$. The proof is by induction on the structure of ψ.

- Assume $Z : \psi$ is an atom. Then $\mathfrak{M} \models \psi$ by construction.
- Assume $Z : \psi$ is the negation of an atom $\alpha(\boldsymbol{a})$. Suppose for the sake of contradiction that $\mathfrak{M} \models \alpha(\boldsymbol{a})$. Then there is some w such that $\boldsymbol{a} \subseteq T(w)$, $X : \alpha(f(\boldsymbol{a})) \in f(w)$, and parameters $f(\boldsymbol{a})$ appear in $f(w')$, for all w' on the path between v and w. Using the backwards transfer property (Lemma 20), there is a node containing both $Z : \neg\alpha(f(\boldsymbol{a}))$ and $X : \alpha(f(\boldsymbol{a}))$, which contradicts internal consistency of mosaics in $\mathcal{S}$.
- Assume $Z : \psi$ is of the form $\psi_1 \vee \psi_2$. Then by internal consistency of v, we must have $Z : \psi_1 \in f(v)$ or $Z : \psi_2 \in f(v)$, so we can apply the inductive hypothesis to get the desired result.
- Assume $Z : \psi$ is of the form $\psi_1 \wedge \psi_2$. Then by internal consistency of v, we must have $Z : \psi_1 \in f(v)$ and $Z : \psi_2 \in f(v)$, so we can apply the inductive hypothesis to get the desired result.
- Assume $Z : \psi$ is of the form $\forall \boldsymbol{x}.\alpha(\boldsymbol{x}, f(\boldsymbol{b})) \to \psi'(\boldsymbol{x}, f(\boldsymbol{b}))$, and assume for the sake of contradiction that there is some $\boldsymbol{a} \in M$ such that $\mathfrak{M} \models \alpha(\boldsymbol{ab}) \wedge \neg\psi'(\boldsymbol{ab})$. Then there must be some w such that $\boldsymbol{ab} \subseteq T(w)$, $X : \alpha(f(\boldsymbol{ab})) \in f(w)$, and parameters $f(\boldsymbol{b})$ appear in $f(w')$, for all w' on the path between v and w. By internal consistency, $X : \exists \boldsymbol{x}.\alpha(\boldsymbol{x}, f(\boldsymbol{b})) \in f(w)$.
  Using the backwards transfer property (Lemma 20), there is a node $v' \leq v$ such that v' and w are on the same branch in T and $Z : \forall \boldsymbol{x}.\alpha(\boldsymbol{x}, f(\boldsymbol{b})) \to \psi'(\boldsymbol{x}, f(\boldsymbol{b})) \in f(v')$.
  If $w < v'$ (respectively, $v' < w$), then the backwards transfer property in Lemma 20 (respectively, forwards universal transfer property in Lemma 21) means that $Z : \forall \boldsymbol{x}.\alpha(\boldsymbol{x}, f(\boldsymbol{b})) \to \psi'(\boldsymbol{x}, f(\boldsymbol{b})) \in f(w)$. Within $f(w)$, internal consistency implies that $\psi'(f(\boldsymbol{ab})) \in f(w)$ as well, a contradiction.

- Assume $Z : \psi$ is of the form $\exists \boldsymbol{x}.\alpha(\boldsymbol{xb}) \wedge \psi'(\boldsymbol{xb})$. If $Z : \alpha(f(\boldsymbol{ab})) \wedge \psi'(f(\boldsymbol{ab})) \in f(v)$ for some $\boldsymbol{a}$, then we can apply the inductive hypothesis and we are done. Otherwise, there must be some child w of v such that $f(v) \rightarrow_{Z:\psi} f(w)$, with $\boldsymbol{a} \subseteq T(w)$ such that $Z : \alpha(f(\boldsymbol{ab})) \wedge \psi'(f(\boldsymbol{ab})) \in f(w)$. The result easily follows from the inductive hypothesis applied to $Z : \alpha(f(\boldsymbol{ab})) \wedge \psi'(f(\boldsymbol{ab})) \in f(w)$. □

We are now ready to prove Theorem 3. For soundness we must show:

> If $\varphi_{\mathrm{L}} \wedge \neg\varphi_{\mathrm{R}}$ is satisfiable, then there is a saturated set of internally consistent GNF mosaics from $\mathcal{M}$ that contains some τ such that $\mathrm{L} : \varphi_{\mathrm{L}}, \mathrm{R} : \neg\varphi_{\mathrm{R}} \in \tau$.

Since $\varphi_{\mathrm{L}} \wedge \neg\varphi_{\mathrm{R}}$ is satisfiable, there is a model $\mathfrak{M}$ of $\varphi_{\mathrm{L}} \wedge \neg\varphi_{\mathrm{R}}$. Let τ be the L-mosaic (equivalently, R-mosaic) of $\emptyset$ in $\mathfrak{M}$, i.e. the set of sentences from $\mathrm{cl}_{\mathrm{L}}(\emptyset) \cup \mathrm{cl}_{\mathrm{R}}(\emptyset)$ that are true in $\mathfrak{M}$. Since $\varphi_{\mathrm{L}} \wedge \neg\varphi_{\mathrm{R}}$ holds in $\mathfrak{M}$, τ contains $\mathrm{L} : \varphi_{\mathrm{L}}$ and $\mathrm{R} : \neg\varphi_{\mathrm{R}}$. Since τ is a realizable mosaic, Proposition 19 implies that there is a saturated set $\mathcal{S}$ of realizable mosaics including τ. Moreover, by Proposition 18, $\mathcal{S}$ must contain only internally consistent mosaics, as desired.

For completeness of the GF mosaic method, we must show:

> If there is a saturated set of internally consistent mosaics from $\mathcal{M}$ containing τ such that $\mathrm{L} : \varphi_{\mathrm{L}}, \mathrm{R} : \neg\varphi_{\mathrm{R}} \in \tau$, then $\varphi_{\mathrm{L}} \wedge \neg\varphi_{\mathrm{R}}$ is satisfiable.

This follows immediately from Proposition 22.

A.2 Bound on number of GF mosaics (Proposition 2)

Recall the statement.

> There is a polynomial function p (independent of φ_{L} and φ_{R}) such that if $|\varphi_{\mathrm{L}}| + |\varphi_{\mathrm{R}}| = n$, then there are at most $2^{2^{p(n)}}$ mosaics in $\mathcal{M}$. Moreover, for a fixed bound on the arity of relations, there are at most $2^{p(n)}$ mosaics in $\mathcal{M}$.

Let $l := |\boldsymbol{e}|$ and $m := \mathrm{width}(\varphi_{\mathrm{L}} \wedge \neg\varphi_{\mathrm{R}})$. Note that $l, m \leq n$.

We first calculate the size of $\mathrm{cl}_Y(P)$ (recall that P is a fixed set of parameters of size $2m$). There are at most n subformulas of φ_Y, but $\mathrm{cl}(\varphi_Y)$ also contains new formulas formed by taking some atom and quantifying out some subset of the m variables. This means that the size of $\mathrm{cl}(\varphi_Y)$ is at most $2^m n$. Each formula in $\mathrm{cl}(\varphi_Y)$ has at most m variables that can be mapped to parameters from P or constants from $\boldsymbol{e}$, so the size of $\mathrm{cl}_Y(P)$ is $M := 2^m n(2m + l)^m$.

Every mosaic in $\mathcal{M}$ is labelled L or R and contains a subset of the formulas in $\mathrm{cl}_{\mathrm{L}}(P) \cup \mathrm{cl}_{\mathrm{R}}(P)$. Hence, $|\mathcal{M}| \leq 2 \cdot 2^{2M} \leq 2^{2^{p(n)}}$ for some polynomial function p that is independent of φ_{L} and φ_{R}.

In the special case that the arity m is bounded, then M is polynomial in n, and $|\mathcal{M}| \leq 2^{p(n)}$ for some polynomial function p.

A.3 Construction of GNF interpolants for GF (additional details for Lemma 4)

We recall the statement and definition of the mosaic interpolants given already Section 4, and then sketch the proof of correctness for one of the inductive cases.

> For each GF X-mosaic $\tau(\boldsymbol{c}) \in \mathcal{N}'$ (the set of eliminated mosaics), we can construct a DAG representation of a formula θ^τ_X such that
>
> (Imp) $\tau_X \models \theta^\tau_X$ and $\theta^\tau_X \models \neg\tau_{\widetilde{X}}$;
>
> (Occ) $\mathrm{occ}(\theta^\tau_X) \subseteq \mathrm{occ}(\varphi_{\mathrm{L}}) \cap \mathrm{occ}(\varphi_{\mathrm{R}})$ if $X = \mathrm{L}$,
> $\mathrm{occ}(\neg\theta^\tau_X) \subseteq \mathrm{occ}(\varphi_{\mathrm{L}}) \cap \mathrm{occ}(\varphi_{\mathrm{R}})$ if $X = \mathrm{R}$;
> $\mathrm{con}(\theta^\tau_X) \subseteq \mathrm{con}(\varphi_{\mathrm{L}}) \cup \mathrm{con}(\varphi_{\mathrm{R}})$;
>
> (Par) $\mathrm{par}(\theta^\tau_X) \subseteq \boldsymbol{c}$;
>
> (GN) θ^τ_X is in GNF, even when parameters are viewed as free variables.
>
> Moreover, if $n = |\varphi_{\mathrm{L}}| + |\varphi_{\mathrm{R}}|$ then there is a DAG representation of $\Theta_i := \{\theta^\tau_Z : \tau \text{ is a } Z\text{-mosaic in } \mathcal{N}_i\}$ such that $|\Theta_i| \leq i \cdot 2^{2^{p'(n)}}$ ($|\Theta_i| \leq i \cdot 2^{p'(n)}$ for the fixed arity case), where p' is some polynomial function independent of φ_{L} and φ_{R}.

Proof. We proceed by induction on i, the stage at which τ was eliminated. We use $\top$ (respectively, $\bot$) as an abbreviation for the GNF sentence $\exists x.x = x$ (respectively, $\exists x.(x = x \wedge \neg(x = x))$).

During the interpolant construction, we will need to include guards in order to remain in GNF, so we introduce some notation to help with this. For any tuple $\boldsymbol{b}$ of parameters that are X-cguarded in τ, we define a formula $\mathrm{gdd}^\tau_X(\boldsymbol{b})$. If $\boldsymbol{b} = \emptyset$, then set $\mathrm{gdd}^\tau_X(\boldsymbol{b}) := \top$. Otherwise, there is some formula $X : \exists\boldsymbol{x}.\beta(\boldsymbol{xb}) \in \tau$ for some X-catom β, and we set $\mathrm{gdd}^\tau_X(\boldsymbol{b}) := \exists\boldsymbol{x}.\beta(\boldsymbol{xb})$ (note that $\boldsymbol{x}$ can be empty, in which case this is just an atomic formula). We will also write $\mathrm{gdd}^\tau_X(\boldsymbol{b}) \wedge \neg\psi$ to indicate some formula $\exists\boldsymbol{x}.(\beta(\boldsymbol{xb}) \wedge \neg\psi)$.

Base case. Consider an X-mosaic $\tau(\boldsymbol{c}) \in \mathcal{N}_1$. Then τ has an internal inconsistency, and we consider the following cases.

- Assume one of these conditions holds for $Y \in \{\mathrm{L}, \mathrm{R}\}$:
  - $Y : \alpha(\boldsymbol{b}), Y : \neg\alpha(\boldsymbol{b}) \in \tau$;
  - $Y : \alpha(\boldsymbol{ab}) \in \tau$, but $Y : \exists\boldsymbol{x}.\alpha(\boldsymbol{xb}) \notin \tau$;
  - $Y : \exists\boldsymbol{x}.\alpha(\boldsymbol{xab}) \in \tau$, but $Y : \exists\boldsymbol{xy}.\alpha(\boldsymbol{xyb}) \notin \tau$;
  - $Y : \psi_1 \wedge \psi_2 \in \tau$, but $Y : \psi_1 \notin \tau$ or $Y : \psi_2 \notin \tau$;
  - $Y : \psi_1 \vee \psi_2 \in \tau$, but $Y : \psi_1 \notin \tau$ and $Y : \psi_2 \notin \tau$;

- $Y : \alpha(\boldsymbol{xb})[\boldsymbol{a}/\boldsymbol{x}], Y : \forall \boldsymbol{x}.\alpha(\boldsymbol{xb}) \to \psi(\boldsymbol{xb}) \in \tau$, but $Y : \psi(\boldsymbol{xb})[\boldsymbol{a}/\boldsymbol{x}] \notin \tau$.

If $Y = X$, then $\theta^\tau_X := \bot$. If $Y = \widetilde{X}$, then $\theta^\tau_X := \top$.

- If $X : \alpha(\boldsymbol{b}), \widetilde{X} : \neg\alpha(\boldsymbol{b}) \in \tau$, then $\theta^\tau_X := \alpha(\boldsymbol{b})$.
- If $X : \neg\alpha(\boldsymbol{b}), \widetilde{X} : \alpha(\boldsymbol{b}) \in \tau$, then $\theta^\tau_X := \mathrm{gdd}^\tau_X(\boldsymbol{b}) \wedge \neg\alpha(\boldsymbol{b})$.
- If $X : \alpha(\boldsymbol{xb})[\boldsymbol{a}/\boldsymbol{x}], \widetilde{X} : \forall \boldsymbol{x}.\alpha(\boldsymbol{xb}) \to \psi(\boldsymbol{xb}) \in \tau$, but $\widetilde{X} : \psi(\boldsymbol{xb})[\boldsymbol{a}/\boldsymbol{x}] \notin \tau$, then $\theta^\tau_X := \alpha(\boldsymbol{xb})[\boldsymbol{a}/\boldsymbol{x}]$.
- If $\widetilde{X} : \alpha(\boldsymbol{xb})[\boldsymbol{a}/\boldsymbol{x}], X : \forall \boldsymbol{x}.\alpha(\boldsymbol{xb}) \to \psi(\boldsymbol{xb}) \in \tau$, but $X : \psi(\boldsymbol{xb})[\boldsymbol{a}/\boldsymbol{x}] \notin \tau$, then $\theta^\tau_X := \mathrm{gdd}^\tau_X(\mathrm{par}(\alpha(\boldsymbol{xb})[\boldsymbol{a}/\boldsymbol{x}])) \wedge \neg\alpha(\boldsymbol{xb})[\boldsymbol{a}/\boldsymbol{x}]$.

Inductive case. Now consider an X-mosaic $\tau(\boldsymbol{c}) \in \mathcal{N}_i \setminus \mathcal{N}_{i-1}$. This means there is some $Y : \eta = Y : \exists \boldsymbol{x}.\alpha(\boldsymbol{xb}) \wedge \psi(\boldsymbol{xb}) \in \tau$ that is not fulfilled in τ, and if $\tau \to_{Y:\eta} \tau'$ then $\tau' \in \mathcal{N}_{i-1}$.

Assume $\mathrm{par}(\psi) = \boldsymbol{b} \subseteq \boldsymbol{c}$, and fix some $\boldsymbol{a} \in P \setminus \boldsymbol{c}$.

We introduce some additional notation. For $Z \in \{\mathrm{L}, \mathrm{R}\}$, we define $\tau\restriction^{\boldsymbol{b}}_Z$ to be the Z-mosaic that results from taking the Z-formulas in τ that only use parameters from $\boldsymbol{b}$, and then unioning with the $\widetilde{Z}$-formulas in τ whose parameters are contained in $\boldsymbol{b}$ and are Z-cguarded in τ. For mosaics υ, υ' and $Z \in \{\mathrm{L}, \mathrm{R}\}$, we write $\upsilon'(\boldsymbol{ab}) \supseteq_Z \upsilon$ if υ' uses only parameters from $\boldsymbol{ab}$ and is obtained from υ be adding only Z-formulas (so the $\widetilde{Z}$-formulas in υ and υ' are identical).

If $Y = X$, then define θ^τ_X to be

$$\bigvee_{\substack{\upsilon'(\boldsymbol{ab}) \supseteq_X \tau\restriction^{\boldsymbol{b}}_X \\ \text{s.t. } \tau \to_{X:\eta} \upsilon'}} \exists \boldsymbol{x}. \left(\bigwedge_{\substack{\tau'(\boldsymbol{ab}) \supseteq_{\widetilde{X}} \upsilon' \\ \text{s.t. } \tau \to_{X:\eta} \tau'}} \theta^{\tau'}_X[\boldsymbol{x}/\boldsymbol{a}] \right).$$

We now begin the proof of correctness, starting with (Imp). Assume there is a model $\mathfrak{M}$ for τ_X. Since $X : \eta \in \tau$, this means that $\mathfrak{M} \models \exists \boldsymbol{x}.\alpha(\boldsymbol{xb}) \wedge \psi(\boldsymbol{xb})$, so there are elements $\boldsymbol{a}' \in M$ and an expansion $\mathfrak{M}'$ of $\mathfrak{M}$ with the interpretation $\boldsymbol{a}^{\mathfrak{M}'} := \boldsymbol{a}'$ such that $\mathfrak{M}' \models \alpha(\boldsymbol{ab}) \wedge \psi(\boldsymbol{ab})$.

Take υ' to be the union of $\tau\restriction^{\boldsymbol{b}}_X$ and the set of formulas $X : \psi'$ in $\mathrm{cl}_X(\boldsymbol{ab})$ such that $\mathfrak{M}' \models \psi'$. It can be checked that $\tau \to_{X:\eta} \upsilon'$. By construction, $\mathfrak{M}' \models \upsilon'_X$. Moreover, for any choice of $\tau'(\boldsymbol{ab}) \supseteq_{\widetilde{X}} \upsilon'$, $\mathfrak{M}' \models \tau'_X$ (since the X-formulas are identical in υ' and τ'). By the inductive hypothesis, this means that $\mathfrak{M}' \models \theta^{\tau'}_X$. Hence, $\mathfrak{M}' \models \bigwedge_{\substack{\tau'(\boldsymbol{ab}) \supseteq_{\widetilde{X}} \upsilon' \\ \text{s.t. } \tau \to_{X:\eta} \tau'}} \theta^{\tau'}_X$ and $\mathfrak{M} \models \theta^\tau_X$.

Now assume that there is a model $\mathfrak{M}$ of θ^τ_X. Then there is some $\upsilon'(\boldsymbol{ab}) \supseteq_X \tau\restriction^{\boldsymbol{b}}_X$ with $\tau \to_{X:\eta} \upsilon'$, elements $\boldsymbol{a}' \in M$, and an expansion $\mathfrak{M}'$ of $\mathfrak{M}$ with the interpretation $\boldsymbol{a}^{\mathfrak{M}'} := \boldsymbol{a}'$, such that $\mathfrak{M}' \models \bigwedge_{\substack{\tau'(\boldsymbol{ab}) \supseteq_{\widetilde{X}} \upsilon' \\ \text{s.t. } \tau \to_{X:\eta} \tau'}} \theta^{\tau'}_X$.

Note that for all $\tau'(\boldsymbol{ab}) \supseteq_{\widetilde{X}} \upsilon'$ such that $\tau \to_{X:\eta} \tau'$, $\mathfrak{M}' \models \theta^{\tau'}_X$ and consequently, the inductive hypothesis implies $\mathfrak{M}' \models \neg\tau'_{\widetilde{X}}$. In particular, consider $\tau' := \upsilon' \cup S'$ where S' is the set of formulas $\widetilde{X} : \psi'$ in $\mathrm{cl}_{\widetilde{X}}(\boldsymbol{ab})$ such that $\mathfrak{M}' \models \psi'$, $\mathrm{par}(\psi') \cap \boldsymbol{a} \neq \emptyset$, and $\widetilde{X} : \psi'$ is X-safe in υ' as required in an X-mosaic. It is clear that $\tau'(\boldsymbol{ab}) \supseteq_{\widetilde{X}} \upsilon'$ and $\tau \to_{X:\eta} \tau'$, so $\mathfrak{M}' \models \neg\tau'_{\widetilde{X}}$ as described above.

Since $\mathfrak{M}' \models \neg\tau'_{\widetilde{X}}$, there is some conjunct χ in $\tau'_{\widetilde{X}}$ such that $\mathfrak{M}' \models \neg\chi$. Consider some conjunct $\chi(\boldsymbol{ab})$ in $\tau'_{\widetilde{X}}$ that actually uses some parameters from $\boldsymbol{a}$. Then by choice of S', $\mathfrak{M}' \models \chi(\boldsymbol{ab})$, so this formula cannot witness the fact that $\mathfrak{M}' \models \neg\tau'_{\widetilde{X}}$. This means there must be some conjunct $\chi(\boldsymbol{b})$ in $\tau'_{\widetilde{X}}$ that only uses parameters from $\boldsymbol{b}$ such that $\mathfrak{M}' \models \neg\chi(\boldsymbol{b})$. If $\widetilde{X} : \chi(\boldsymbol{b})$ is in τ', then $\widetilde{X} : \chi(\boldsymbol{b}) \in \tau$ by the definition of $\tau \to_{X:\eta} \tau'$. Likewise, if $\widetilde{X} : \chi(\boldsymbol{b}) \notin \tau'$, then it must be the case that $\mathrm{par}(\chi)$ is X-cguarded in τ', so $\widetilde{X} : \chi(\boldsymbol{b}) \notin \tau$ (otherwise it would contradict the definition of $\tau \to_{X:\eta} \tau'$). In either case, this means $\chi(\boldsymbol{b})$ must be a conjunct in $\tau_{\widetilde{X}}$. Hence, $\mathfrak{M} \models \neg\tau_{\widetilde{X}}$.

For (Occ), recall that for every atomic relation S in ψ (where S is not the equality relation), $(S, +) \in \mathrm{occ}(\phi)$ (respectively, $(S, -) \in \mathrm{occ}(\phi)$) iff S occurs positively (respectively, negatively) in ϕ. The desired occurrence of constants and the polarity of relations follows from the inductive hypothesis and the fact that the polarity of occurrences of relations are preserved when moving to mosaics.

The last thing to prove is (Par) and (GN). The inductively defined interpolants $\theta^{\tau'}_X$ are in GNF and only use parameters in τ', i.e. parameters from $\boldsymbol{ab}$. Any parameters from $\boldsymbol{a}$ in $\theta^{\tau'}_X$ are removed in θ^τ_X using existential quantification. Overall, this means that (Par) and (GN) hold for θ^τ_X.

If $Y = \widetilde{X}$, then define θ^τ_X to be

$$\mathrm{gdd}^\tau_X(\boldsymbol{b}) \wedge \neg \left(\bigvee_{\substack{\upsilon'(\boldsymbol{ab}) \supseteq_{\widetilde{X}} \tau\restriction^{\boldsymbol{b}}_{\widetilde{X}} \\ \text{s.t. } \tau \to_{\widetilde{X}:\eta} \upsilon'}} \exists \boldsymbol{x}. \left(\bigwedge_{\substack{\tau'(\boldsymbol{ab}) \supseteq_X \upsilon' \\ \text{s.t. } \tau \to_{\widetilde{X}:\eta} \tau'}} \theta^{\tau'}_{\widetilde{X}}[\boldsymbol{x}/\boldsymbol{a}] \right) \right).$$

(Imp) is shown by arguing that $\tau_{\widetilde{X}} \models \neg\theta^\tau_X$ and $\neg\theta^\tau_X \models \neg\tau_X$.

The main difference is that we need to ensure the negation is guarded in order to prove (GN). Using the inductive hypothesis, it can be shown that inside the negation, the formula satisfies (GN), and only uses parameters from $\boldsymbol{b}$ (any parameters from $\boldsymbol{a}$ are removed using quantification). Because $\widetilde{X} : \eta$ is a $\widetilde{X}$-formula using $\boldsymbol{b}$ in an X-mosaic, we know that $\mathrm{gdd}^\tau_X(\boldsymbol{b})$ exists. This is enough to conclude that θ^τ_X is in GNF, even when the parameters are viewed as free variables.

Size of interpolants. Finally, we seek to bound the size of the shared DAG representation of all θ^τ_X coming from $\tau \in \mathcal{N}_i$.

A routine calculation shows that the number of nodes and edges added to the shared DAG representation at each stage i is polynomial in the number of mosaics in $\mathcal{N}'$. Since the number of mosaics in $\mathcal{N}'$ is at most doubly exponential in n (respectively, singly exponential in n in the fixed arity case) by Proposition 2, this implies that there is some polynomial function p' independent of φ_{L} and φ_{R} such that $|\Theta_i| \leq i \cdot 2^{2^{p'(n)}}$ (respectively, $|\Theta_i| \leq i \cdot 2^{p'(n)}$). $\square$

Quasi-interpolants. We can derive the quasi-interpolant result (Theorem 6). Recall the statement:

> For GF sentences $\varphi_{\mathrm{L}} \models \varphi_{\mathrm{R}}$ there is a GF sentence θ such that $\varphi_{\mathrm{L}} \models \theta$ and $\theta \models \varphi_{\mathrm{R}}$ and θ only uses relations in the common signature or relations occurring as a guard for some quantification in φ_{L} or φ_{R}.

We modify the inductive case of Lemma 4 as follows:
Assume there is an X-mosaic $\tau(\boldsymbol{c}) \in \mathcal{N}_i \setminus \mathcal{N}_{i-1}$. with $Y : \eta = Y : \exists \boldsymbol{x}.\alpha(\boldsymbol{xb}) \wedge \psi(\boldsymbol{xb}) \in \tau$, and such that $\tau' \in \mathcal{N}_{i-1}$ for all $\tau \rightarrow_{Y:\eta} \tau'$.
If $Y = X$, then define θ^{τ}_X to be

$$\bigvee_{\substack{\upsilon'(\boldsymbol{ab}) \supseteq_X \tau\restriction^{\boldsymbol{b}}_X \\ \text{s.t. } \tau \rightarrow_{X:\eta} \upsilon'}} \exists \boldsymbol{x}. \left(\alpha(\boldsymbol{xb}) \wedge \bigwedge_{\substack{\tau'(\boldsymbol{ab}) \supseteq_{\widetilde{X}} \upsilon' \\ \text{s.t. } \tau \rightarrow_{X:\eta} \tau'}} \theta^{\tau'}_X [\boldsymbol{x}/\boldsymbol{a}] \right).$$

If $Y = \widetilde{X}$, then define θ^{τ}_X to be

$$\mathrm{gdd}^{\tau}_X(\boldsymbol{b}) \wedge \neg \left(\bigvee_{\substack{\upsilon'(\boldsymbol{ab}) \supseteq_{\widetilde{X}} \tau\restriction^{\boldsymbol{b}}_{\widetilde{X}} \\ \text{s.t. } \tau \rightarrow_{\widetilde{X}:\eta} \upsilon'}} \exists \boldsymbol{x}. \left(\alpha(\boldsymbol{xb}) \wedge \bigwedge_{\substack{\tau'(\boldsymbol{ab}) \supseteq_X \upsilon' \\ \text{s.t. } \tau \rightarrow_{\widetilde{X}:\eta} \tau'}} \theta^{\tau'}_{\widetilde{X}} [\boldsymbol{x}/\boldsymbol{a}] \right) \right).$$

In other words, we just use the guard that is in the formula $\exists \boldsymbol{x}.\alpha(\boldsymbol{xb}) \wedge \psi(\boldsymbol{xb})$ that is guiding the interpolant construction at that stage. Because of property (Par), each $\theta^{\tau'}_Y$ can only use parameters from $\boldsymbol{ab}$. After $\boldsymbol{a}$ is removed by quantification, this means that $\alpha(\boldsymbol{xb})$ guards $\theta^{\tau'}_Y[\boldsymbol{x}/\boldsymbol{a}]$ for all $\tau'(\boldsymbol{ab})$ with $\tau \rightarrow_{Y:\eta} \tau'$.

B. Relativized GNF Mosaics

In this section, we describe GNF mosaics in more detail. Recall that in Section 5, we defined GNF mosaics for GNF sentences without equality and with ordinary (rather than relativized) quantifiers. For space considerations, we omit the proofs for those mosaics. Instead, we now define a more general form of GNF mosaics for relativized GNF formulas (still without equality), and give full proofs of soundness and completeness for this more complicated case. Note that we overload all of the terminology and notation that was used for the GF mosaics and and GNF mosaics defined in the body.

B.1 Weak GN-normal form

The mosaics described in this section will use formulas that are in *weak* GN-*normal form*, a slight relaxation of the GN-normal form defined in the body. Formally, weak GN-normal form ϕ are formulas that can be generated using the following recursive definition:

$$\phi ::= \textstyle\bigvee_i \exists \boldsymbol{x}. \bigwedge_j \psi_{ij}$$
$$\psi ::= R\,\boldsymbol{t} \mid \alpha(\boldsymbol{x}) \wedge \phi(\boldsymbol{x}) \mid \alpha(\boldsymbol{x}) \wedge \neg\phi(\boldsymbol{x})$$

where R is either a relation symbol or the equality relation, $\alpha(\boldsymbol{x})$ is a guard with $\mathrm{free}(\alpha) = \boldsymbol{x}$, and $\boldsymbol{t}$ is a vector over variables and constants. The idea is that the grammar generates UCQ-shaped formulas, where each conjunct is either an atom, an answer-guarded UCQ-shaped formula, or a guarded negation of a UCQ-shaped formula. Note that a formula in GN-normal form is in trivially in weak GN-normal form.

Given a formula ϕ in GN-normal form or weak GN-normal form, we define $\mathrm{size}_{\mathrm{CQ}}(\phi)$ (respectively, $\mathrm{rank}_{\mathrm{CQ}}(\phi)$) to be the maximum size (respectively, maximum number of conjuncts ψ_i) of any CQ-shaped subformula $\exists \boldsymbol{x}. \bigwedge_i \psi_i$ of ϕ for non-empty $\boldsymbol{x}$. For the purposes of the CQ-rank, note that formulas $\alpha(\boldsymbol{x}) \wedge \phi(\boldsymbol{x})$ and $\alpha(\boldsymbol{x}) \wedge \neg\phi(\boldsymbol{x})$ are treated as single conjuncts in a CQ-shaped subformula.

We write $\mathrm{guards}(\phi)$ to denote the set of atoms that appear as guards of a negation in ϕ (where ϕ is in GNF, weak GN-normal form, or GN-normal form).

Any GNF sentence can be converted to weak GN-normal form with an exponential blow up in overall size, and preserving the polarity of relations.

Lemma 23. *Let ϕ be a* GNF *formula. We can construct an equivalent ϕ' in* GN-*normal form such that*
- *$|\phi'|$ is exponential in $|\phi|$;*
- $\mathrm{width}(\phi') \leq |\phi|$*;*
- $\mathrm{size}_{\mathrm{CQ}}(\phi') \leq |\phi|$*;*
- $\mathrm{rank}_{\mathrm{CQ}}(\phi') \leq |\phi|$*;*
- $\mathrm{occ}(\phi') = \mathrm{occ}(\phi)$*;*
- $\mathrm{guards}(\phi') = \mathrm{guards}(\phi)$.

Proof. Use the transformations described in [5] to convert to GN-normal form. $\square$

The advantage of weak GN-normal form over GN-normal form is GF sentences can be converted to weak GN-normal form with only a polynomial blow-up.

Lemma 24. *Let ϕ be a* GF *sentence. We can construct an equivalent ϕ' in weak* GN-*normal form such that*

- $|\phi'|$ *is polynomial in* $|\phi|$*;*
- $\mathrm{width}(\phi') = \mathrm{width}(\phi)$*;*
- $\mathrm{rank}_{\mathrm{CQ}}(\phi') = 1$*;*
- $\mathrm{occ}(\phi') = \mathrm{occ}(\phi)$.

Proof. The idea is that, for every subformula of ϕ, we find the innermost guard from a guarded quantification, and copy this guard to that subformula. The resulting sentence is in weak GN-normal form. Note that this result only holds for GF sentences (not GF formulas).

Formally, fix a signature σ, and let w be the maximum arity of any relation in σ. We prove a slightly stronger statement:

> For every GF formula $\phi(\boldsymbol{y}) := \alpha(\boldsymbol{y}) \wedge \psi(\boldsymbol{y})$ over signature σ, we can construct $\phi'(\boldsymbol{y}) := \alpha(\boldsymbol{y}) \wedge \psi'(\boldsymbol{y})$ in weak GN-normal form such that $|\psi'| \leq (w|\psi|)^2$, $\mathrm{width}(\psi') = \mathrm{width}(\psi)$, $\mathrm{rank}_{\mathrm{CQ}}(\psi') = 1$.

We proceed by structural induction on ψ.

- Assume ψ is atomic. Then set $\psi' := \psi$.
- Assume ψ is $\neg\eta$. Then set $\psi' := \psi$.
- Assume ψ is $\psi_1 \wedge \psi_2$. Inductively, construct ϕ'_i for $\phi_i := \alpha \wedge \psi_i$, and set $\psi' := (\alpha \wedge \phi'_1) \wedge (\alpha \wedge \phi'_2)$. Then $|\psi'| \leq (w + (w|\psi_1|)^2) + (w + (w|\psi_1|)^2) + 1 \leq (w(|\psi_1| + |\psi_2| + 1))^2 = (w|\psi|)^2$. Although $(\alpha \wedge \phi'_1) \wedge (\alpha \wedge \phi'_2)$ is a CQ-shaped subformula with two conjuncts, it does not contribute to the CQ-rank, since there is no quantification.
  Similarly for ψ of the form $\psi_1 \vee \psi_2$.
- Assume ψ is $\exists \boldsymbol{x}.\beta(\boldsymbol{xy}) \wedge \eta(\boldsymbol{xy})$. Inductively, construct η' for $\beta(\boldsymbol{xy}) \wedge \eta(\boldsymbol{xy})$ and set $\psi' := \exists \boldsymbol{x}.\beta(\boldsymbol{xy} \wedge \eta'(\boldsymbol{xy})$. Then $|\psi'| \leq (w|\eta|)^2 + 2w + 3 \leq (w(|\eta| + 3))^2 \leq (w|\psi|)^2$.

□

B.2 Relativized GNF mosaics

Fix formulas φ_{L} and φ_{R} over signatures σ_{L} and σ_{R} that are in weak GN-normal form and do not use equality. We assume that $\mathrm{free}(\varphi_{\mathrm{L}}) = \mathrm{free}(\varphi_{\mathrm{R}}) = \boldsymbol{z}$ (which may be the empty set if the input were sentences), and that these variable names are never used in any quantification in the formulas.[1] From now on, we view $\boldsymbol{z}$ as parameters (so $\mathrm{par}(\varphi_{\mathrm{L}}) = \mathrm{par}(\varphi_{\mathrm{R}}) = \boldsymbol{z}$). Let $\boldsymbol{e} := \mathrm{con}(\varphi_{\mathrm{L}}) \cup \mathrm{con}(\varphi_{\mathrm{R}})$.

We also assume φ_{L} and φ_{R} are in $\mathbb{U}$-relativized GNF, for $\mathbb{U}$ a distinguished set of unary relations from $\sigma_{\mathrm{L}} \cup \sigma_{\mathrm{R}}$. As shorthand, we write $\exists \boldsymbol{x}.\boldsymbol{U}(\boldsymbol{x}) \wedge \psi(\boldsymbol{xy})$ for a formula with relativized quantification of the form $\exists \boldsymbol{x}.U_1 x_1 \wedge \cdots \wedge U_k x_k \wedge \psi(\boldsymbol{xy})$ where $\boldsymbol{x} = x_1 \cdots x_k$ and $\boldsymbol{U} = U_1 \cdots U_k$ with $\boldsymbol{U} \subseteq \mathbb{U}$.

Specializations. The mosaics must contain additional information about CQ-shaped subformulas of the form $\exists \boldsymbol{x}.\boldsymbol{U}(\boldsymbol{x}) \wedge \bigwedge_i \psi_i(\boldsymbol{xy})$. Consider a weak GN-normal form formula $\phi(\boldsymbol{y})$ of the form $\exists \boldsymbol{x}.\boldsymbol{U}(\boldsymbol{x}) \wedge \bigwedge_i \psi_i(\boldsymbol{xy})$. A *(relativized) specialization* of ϕ is a formula ϕ' obtained from ϕ by the following operations:

- select a subset $\boldsymbol{z}$ of $\boldsymbol{x}$ (call variables from $\boldsymbol{yz}$ the *inside variables* and variables from $\boldsymbol{x} \setminus \boldsymbol{z}$ the *outside variables*);
- select a partition $\boldsymbol{x}_1, \ldots, \boldsymbol{x}_k$ of the outside variables, with the property that for every ψ_j, either ψ_j has no outside variables or all of its outside variables are contained in some partition element $\boldsymbol{x}_i$;
- let $\boldsymbol{U}_0$ be the conjunction of atoms from $\boldsymbol{U}(\boldsymbol{x})$ using only inside variables, and let $\boldsymbol{U}_j$ be the conjunction of atoms from $\boldsymbol{U}(\boldsymbol{x})$ using outside variables in $\boldsymbol{x}_j$;
- let χ_0 be the conjunction of the ψ_i using only inside variables, and let χ_j be the conjunction of the ψ_i using outside variables and satisfying $\mathrm{par}(\psi_i) \subseteq \boldsymbol{x}_j\boldsymbol{yz}$;
- set $\phi'(\boldsymbol{yz})$ to be

$$\boldsymbol{U}_0(\boldsymbol{z}) \wedge \chi_0(\boldsymbol{yz}) \wedge \bigwedge_{j \in \{1,\ldots,k\}} \exists \boldsymbol{x}_j.\boldsymbol{U}_j(\boldsymbol{x}_j) \wedge \chi_j(\boldsymbol{x}_j\boldsymbol{yz}).$$

Roughly speaking, each specialization describes a way in which the original CQ-shaped formula could be satisfied in some tree decomposition.

We say a specialization $\boldsymbol{U}_0(\boldsymbol{z}) \wedge \chi_0(\boldsymbol{yz}) \wedge \bigwedge_{j \in \{1,\ldots,k\}} \exists \boldsymbol{x}_j.\boldsymbol{U}_j(\boldsymbol{x}_j) \wedge \chi_j(\boldsymbol{x}_j\boldsymbol{yz})$ is *non-trivial* if either there are no outside variables (thus the specialization is only $\boldsymbol{U}_0 \wedge \chi_0$), or χ_0 is non-empty, or the partition of the outside variables is non-trivial.

Lemma 25. *Let $\phi(\boldsymbol{y})$ be a formula of the form $\exists \boldsymbol{x}.\boldsymbol{U}(\boldsymbol{x}) \wedge \bigwedge_i \psi_i(\boldsymbol{xy})$. Given a structure $\mathfrak{M}$ and a tree decomposition T of $\mathfrak{M}$, if there exists a node v with $\boldsymbol{b} \subseteq T(v)$ and $\mathfrak{M}, \boldsymbol{b} \models \phi(\boldsymbol{y})$, then there is a non-trivial specialization $\phi'(\boldsymbol{yz})$ of ϕ and a node w with a tuple $\boldsymbol{c}$ such that $\boldsymbol{bc} \in T(w)$ and $\mathfrak{M}, \boldsymbol{bc} \models \phi'(\boldsymbol{yz})$. We say ϕ' is the* specialization of ϕ defined by $\mathfrak{M}$ and $\boldsymbol{b}$.

Proof. There must be some tuple $\boldsymbol{a} = a_1 \ldots a_m$ of elements (corresponding to $\boldsymbol{x} = x_1 \ldots x_m$) such that $\mathfrak{M}, \boldsymbol{ab} \models \boldsymbol{U}(\boldsymbol{x}) \wedge \bigwedge_i \psi_i(\boldsymbol{xy})$.

For a node w, let w_1 denote the parent of w in T, and $w_2, \ldots, w_k$ denote the children of w in T. We say $\boldsymbol{ab}$ is contained in direction $j = 1$ if $\boldsymbol{ab}$ appears in the tree resulting from removing from T the subtree rooted at w, and is contained in direction $j > 1$ if $\boldsymbol{ab}$ appears in the subtree rooted at w_j.

If there is a node w in the tree decomposition T with $\boldsymbol{ab} \subseteq T(w)$, then it is easy (take the specialization where all variables $\boldsymbol{x}$ are inside variables).

[1] We remark that considering only formulas where $\mathrm{free}(\varphi_{\mathrm{L}}) = \mathrm{free}(\varphi_{\mathrm{R}})$ is not actually a restriction for satisfiability or interpolation. Indeed, given $\mathbb{U}$-relativized weak GN-normal form formulas $\varphi'_{\mathrm{L}}(\boldsymbol{xz})$ and $\varphi'_{\mathrm{R}}(\boldsymbol{yz})$, we can existentially quantify out the free variables of φ_{L} that are not free variables of φ_{R}, and we can universally quantify out the free variables of φ_{R} that are not free variables of φ_{L}, introducing a fresh guard relation and fresh relativizers as necessary. That is,

$$\varphi'_{\mathrm{L}} \models \varphi'_{\mathrm{R}} \quad \text{iff} \quad R\,\boldsymbol{z} \wedge \exists \boldsymbol{x}.\,\big(U_{\mathrm{L}}\,x_1 \wedge \cdots \wedge U_{\mathrm{L}}\,x_j \wedge \varphi'_{\mathrm{L}}(\boldsymbol{xz})\big) \models R\,\boldsymbol{z} \wedge \neg\exists \boldsymbol{y}.\,\big(U_{\mathrm{R}}\,y_1 \wedge \cdots \wedge U_{\mathrm{R}}\,y_k \wedge \varphi'_{\mathrm{R}}(\boldsymbol{yz})\big)$$

where $\boldsymbol{x} = x_1 \ldots x_j$, $\boldsymbol{y} = y_1 \ldots y_k$, R is a fresh predicate of arity $|\boldsymbol{z}|$, and U_{L} and U_{R} are fresh unary predicates.

Otherwise, there is a node w in T with $\boldsymbol{b} \subseteq T(w)$ and such that $\boldsymbol{ab}$ is not contained in any direction from w (if not, then starting at the node containing $\boldsymbol{b}$, we could eventually reach a node w' with $\boldsymbol{ab} \subseteq T(w')$, which we are assuming is not possible).

Let $\boldsymbol{z}$ be the tuple of variables from $\boldsymbol{x}$ corresponding to elements in $\boldsymbol{c} := T(w) \cap \boldsymbol{a}$ (i.e. $x_i \in \boldsymbol{z}$ iff $a_i \in T(w)$). Take $\boldsymbol{yz}$ to be the inside variables, corresponding to elements $\boldsymbol{bc}$ (the elements inside $T(w)$.

Let O be the nonempty set of elements from $\boldsymbol{a}$ that are not in $T(w)$ (i.e. the elements that correspond to outside variables). We partition O into $O_1, \ldots, O_k$ such that O_j is the set of elements from O that are contained in direction j. Because $o \in O$ do not appear in $T(w)$, $O_1, \ldots, O_k$ partition O. This induces a partition of the outside variables.

Taking the resulting specialization ϕ', we have $\mathfrak{M}, \boldsymbol{bc} \models \phi'(\boldsymbol{yz})$. $\square$

In the other direction, every specialization of ϕ logically implies ϕ.

Lemma 26. *Let $\phi(\boldsymbol{y})$ be a formula of the form $\exists \boldsymbol{x}.\boldsymbol{U}(\boldsymbol{x}) \wedge \bigwedge_i \psi_i(\boldsymbol{xy})$. For all structures $\mathfrak{M}$ and for all specializations $\phi'(\boldsymbol{yz})$ of ϕ, if $\mathfrak{M}, \boldsymbol{bc} \models \phi'(\boldsymbol{yz})$, then $\mathfrak{M}, \boldsymbol{b} \models \phi(\boldsymbol{y})$.*

Closure. The *closure* $\mathrm{cl}(X^p : \varphi)$ is the smallest set C of formulas containing $X^p : \varphi$ and such that:

- if $X^+ : \alpha(\boldsymbol{y}) \wedge \neg\psi(\boldsymbol{y}) \in C$, then $X^+ : \alpha(\boldsymbol{y}), X^- : \psi(\boldsymbol{y}) \in C$;
- if $X^- : \alpha(\boldsymbol{y}) \wedge \neg\psi(\boldsymbol{y}) \in C$, then $X^- : \alpha(\boldsymbol{y}), X^+ : \psi(\boldsymbol{y}) \in C$;
- if $X^p : \bigwedge_i \psi_i \in C$ or $X^p : \bigvee_i \psi_i \in C$, then $X^p : \psi_i \in C$ for all i;
- if $X^p : \exists \boldsymbol{x}.\boldsymbol{U}(\boldsymbol{x}) \wedge \psi(\boldsymbol{xy}) \in C$, and $\boldsymbol{U}_0(\boldsymbol{z}) \wedge \chi_0(\boldsymbol{yz}) \wedge \bigwedge_k \exists \boldsymbol{x}_k.\boldsymbol{U}_k(\boldsymbol{x}_k) \wedge \chi_k(\boldsymbol{x}_k\boldsymbol{yz})$ for $\boldsymbol{U}_0(\boldsymbol{z}) = U_1 z_1 \wedge \cdots \wedge U_j z_j$ is a specialization of $\exists \boldsymbol{x}.\boldsymbol{U}(\boldsymbol{x}) \wedge \psi(\boldsymbol{xy})$, then $X^p : U_i z_i \in C$ for all $i \in \{1, \ldots, j\}$ and $X^p : \chi_0(\boldsymbol{yz}) \wedge \bigwedge_j \exists \boldsymbol{x}_j.\boldsymbol{U}_j(\boldsymbol{x}_j) \wedge \chi_j(\boldsymbol{x}_j\boldsymbol{yz}) \in C$;
- if $X^+ : \alpha(\boldsymbol{xy}) \in C$ for $\boldsymbol{x} = \{x_1, \ldots, x_k\}$ and $X^+ : U_1 z_1, \ldots, X^+ : U_k z_k \in C$ for $U_i \in \mathbb{U}$, then $X^+ : \exists \boldsymbol{x}.U_1 x_1 \wedge \cdots \wedge U_k x_k \wedge \alpha(\boldsymbol{xy}) \in C$.

Given parameters $\boldsymbol{c}$, we let $\mathrm{cl}_{\mathrm{L}}(\boldsymbol{c})$ (respectively, $\mathrm{cl}_{\mathrm{R}}(\boldsymbol{c})$) consist of formulas $X^p : \psi$ from $\mathrm{cl}(\mathrm{L}^+ : \varphi_{\mathrm{L}})$ (respectively, $\mathrm{cl}(\mathrm{R}^- : \varphi_{\mathrm{R}})$) with the free variables replaced by parameters from $\boldsymbol{c}$ or constants from $\boldsymbol{e}$, and such that the resulting formulas has parameters contained in $\boldsymbol{c}$. Recall that we are viewing any variables $\boldsymbol{z}$ that were free in φ_X as parameters, so these will never be replaced by other parameters. For instance, consider a formula $\mathrm{L}^+ : \psi(\boldsymbol{yz}') \in \mathrm{cl}(\mathrm{L}^+ : \varphi_{\mathrm{L}})$, where $\boldsymbol{z}' \subseteq \boldsymbol{z}$ are parameters corresponding to original free variables. Then a possible substitution would be $\mathrm{L}^+ : \psi(\boldsymbol{yz}')[\boldsymbol{b}/\boldsymbol{y}]$ where $\boldsymbol{b} \subseteq \boldsymbol{c} \cup \boldsymbol{e}$, and $\mathrm{L}^+ : \psi(\boldsymbol{yz}')[\boldsymbol{b}/\boldsymbol{y}] \in \mathrm{cl}_{\mathrm{L}}(\boldsymbol{c})$ if $\boldsymbol{z}' \subseteq \boldsymbol{c}$.

Each formula in $\mathrm{cl}_{\mathrm{L}}(\boldsymbol{c}) \cup \mathrm{cl}_{\mathrm{R}}(\boldsymbol{c})$ is labelled with a provenance L or R and a polarity $+$ or $-$ to indicate whether the formula is built from a subformula of φ_{L} or $\neg\varphi_{\mathrm{R}}$, and whether this subformula occurred positively or negatively (i.e., whether it occurred within the scope of an even or odd number of negations). Given a formula ψ and a polarity p, we write $p\psi$ to mean ψ (respectively, $\neg\psi$) if p is $+$ (respectively, $-$).

The formulas in $\mathrm{cl}_{\mathrm{L}}(\boldsymbol{c}) \cup \mathrm{cl}_{\mathrm{R}}(\boldsymbol{c})$ respect the polarity of relations, guards, and constants in the original formulas.

Proposition 27. *For all $\tau \subseteq \mathrm{cl}_{\mathrm{L}}(\boldsymbol{c}) \cup \mathrm{cl}_{\mathrm{R}}(\boldsymbol{c})$,*

- *if $\mathrm{L}^p : \psi \in \tau$, then $\mathrm{occ}(p\psi) \subseteq \mathrm{occ}(\varphi_{\mathrm{L}})$, $\mathrm{guards}(\psi) \subseteq \mathrm{guards}(\varphi_{\mathrm{L}})$, and $\mathrm{con}(\psi) \subseteq \mathrm{con}(\varphi_{\mathrm{L}}) \cup \mathrm{con}(\varphi_{\mathrm{R}})$;*
- *if $\mathrm{R}^p : \psi \in \tau$, then $\mathrm{occ}(\neg p\psi) \subseteq \mathrm{occ}(\varphi_{\mathrm{R}})$, $\mathrm{guards}(\psi) \subseteq \mathrm{guards}(\varphi_{\mathrm{R}})$, and $\mathrm{con}(\psi) \subseteq \mathrm{con}(\varphi_{\mathrm{L}}) \cup \mathrm{con}(\varphi_{\mathrm{R}})$.*

Guards and relativizers. We say α is an L-*atom* (respectively, R-*atom*) if $(\alpha, +) \in \mathrm{occ}(\varphi_{\mathrm{L}})$ (respectively, $(\alpha, -) \in \mathrm{occ}(\varphi_{\mathrm{R}})$). Likewise, we say α is an L-*catom* (respectively, R-*catom*) if $(\alpha, +) \in \mathrm{occ}(\varphi_{\mathrm{L}}) \cap \mathrm{occ}(\varphi_{\mathrm{R}})$ (respectively, $(\alpha, -) \in \mathrm{occ}(\varphi_{\mathrm{L}}) \cap \mathrm{occ}(\varphi_{\mathrm{R}})$). We say α is an X-*guard* (respectively, X-*cguard*) if α is an X-atom (respectively, X-catom) and $\alpha \in \mathrm{guards}(\varphi_{\mathrm{L}}) \cup \mathrm{guards}(\varphi_{\mathrm{R}})$. We say U is an X-*relativizer* (respectively, X-*crelativizer*) if U is an X-atom (respectively, X-catom) and $U \in \mathbb{U}$.

Let τ be a collection of formulas from $\mathrm{cl}_{\mathrm{L}}(\boldsymbol{c}) \cup \mathrm{cl}_{\mathrm{R}}(\boldsymbol{c})$ and let $\boldsymbol{d} \subseteq \boldsymbol{c}$ be a distinguished subset of these parameters.

We say $b \in \boldsymbol{c}$ *is X-crelativized in* τ if there is some $X^+ : Ub \in \tau$ where $U \in \mathbb{U}$ and U is an X-catom. We say $\boldsymbol{b}$ is X-crelativized if every $b \in \boldsymbol{b}$ is X-crelativized.

We say $\boldsymbol{b} \subseteq \boldsymbol{c}$ *is X-cguarded in* τ if $|\boldsymbol{b}| \leq 1$, or if there is some $X^+ : \exists \boldsymbol{x}.\boldsymbol{U}(\boldsymbol{x}) \wedge \alpha(\boldsymbol{x}\boldsymbol{b}') \in \tau$ where α and $\boldsymbol{U}$ are X-catoms, $\alpha \in \mathrm{guards}(\varphi_{\mathrm{L}}) \cup \mathrm{guards}(\varphi_{\mathrm{R}})$, $\boldsymbol{U} \subseteq \mathbb{U}$, $\boldsymbol{b}' \supseteq \boldsymbol{b}$, and $\boldsymbol{b}' \cap \boldsymbol{d}$ is X-crelativized. Note that this definition is different than in the GF mosaics because of technical requirements when proving soundness and completeness of this GNF mosaic method. In particular, unary relations are always considered to be X-cguarded, which allows X-relativizers to move freely between linked mosaics (which is important for completeness).

We say a formula $\widetilde{X}^p : \psi$ is X-*safe* in τ if

- $\widetilde{X}^p = \widetilde{X}^+$ and $\mathrm{par}(\psi) \cap \boldsymbol{d}$ is X-crelativized in τ and $\mathrm{par}(\psi)$ is X-cguarded in τ, or
- $\widetilde{X}^p = \widetilde{X}^-$ and $\mathrm{par}(\psi) \cap \boldsymbol{d}$ is X-crelativized in τ.

We emphasize that these definitions (X-cguarded, X-crelativized, and X-safety) depend on the distinguished set of parameters $\boldsymbol{d}$ in τ, so when we talk about, e.g., being X-cguarded in τ, we will assume that τ has some distinguished set of parameters $\boldsymbol{d}$.

Relativized GNF mosaics. Let $\boldsymbol{c}$ be a tuple of parameters such that $|\boldsymbol{c}| \leq \mathrm{width}(\varphi_{\mathrm{L}} \wedge \neg\varphi_{\mathrm{R}})$. A *(relativized)* GNF X-*mosaic* τ for $\varphi_{\mathrm{L}} \wedge \neg\varphi_{\mathrm{R}}$ using parameters $\boldsymbol{c}$ is a subset of $\Gamma(\boldsymbol{c}) := \mathrm{cl}_{\mathrm{L}}(\boldsymbol{c}) \cup \mathrm{cl}_{\mathrm{R}}(\boldsymbol{c})$ with some distinguished parameters $\boldsymbol{d} \subseteq \boldsymbol{c}$ (which we call the *relativized parameters* in τ) and satisfying the requirement that if $\widetilde{X}^p : \psi \in \tau$, then $\widetilde{X}^p : \psi$ is X-safe in τ (with respect to $\boldsymbol{d}$). We write $\tau(\boldsymbol{c})$ to emphasize that τ uses parameters $\boldsymbol{c}$.

Let τ be a GNF X-mosaic over parameters $\boldsymbol{c}$ and relativizing $\boldsymbol{d}$. Assume $Y^+ : \eta \in \tau$ for η an *existential requirement* of the form $\exists \boldsymbol{x}.\psi(\boldsymbol{xb})$. Then τ is linked to τ' via $Y^+ : \eta$ (written $\tau \rightarrow_{Y^+:\eta} \tau'$) if

- τ' is a GNF Y-mosaic over parameters $\boldsymbol{ab}$ relativizing $\boldsymbol{a} \cup (\boldsymbol{b} \cap \boldsymbol{d})$ (respectively, $\boldsymbol{a}$) if $Y = X$ (respectively, $Y = \widetilde{X}$), and with $\boldsymbol{a} \cap \boldsymbol{c} = \emptyset$;
- $Y^+ : \psi(\boldsymbol{xb})[\boldsymbol{a}/\boldsymbol{x}] \in \tau'$;
- for all $Y^p : \psi'(\boldsymbol{b}) \in \tau$, $Y^p : \psi'(\boldsymbol{b}) \in \tau'$;
- for all $\widetilde{Y}^p : \psi'(\boldsymbol{b}) \in \tau$ such that $\widetilde{Y}^p : \psi'(\boldsymbol{b})$ is Y-safe in τ', $\widetilde{Y}^p : \psi'(\boldsymbol{b}) \in \tau'$;
- for all $Z^p : \psi'(\boldsymbol{b}) \in \tau'$, $Z^p : \psi'(\boldsymbol{b}) \in \tau$.

As in the GF mosaics, the maximum amount of information about shared parameters is passed through a link, and no new information about shared parameters can be added.

Fix some set P of parameters of size $2 \cdot \mathrm{width}(\varphi_{\mathrm{L}} \wedge \neg\varphi_{\mathrm{R}})$ that includes the parameters $\mathrm{par}(\varphi_{\mathrm{L}}) = \mathrm{par}(\varphi_{\mathrm{R}})$ (the parameters corresponding to the original free variables in the input formulas). Let $\mathcal{M}$ be the set of GNF mosaics over parameters P.

Coherence. A GNF X-mosaic $\tau(\boldsymbol{c})$ is *downward closed* if the following properties are satisfied:

- if $Y^+ : \alpha(\boldsymbol{b}) \wedge \neg\psi(\boldsymbol{b}) \in \tau$, then $Y^+ : \alpha(\boldsymbol{b}) \in \tau$ and $Y^- : \psi(\boldsymbol{b}) \in \tau$;
- if $Y^- : \alpha(\boldsymbol{b}) \wedge \neg\psi(\boldsymbol{b}) \in \tau$ and $Z^+ : \alpha(\boldsymbol{b}) \in \tau$, then $Y^+ : \psi(\boldsymbol{b}) \in \tau$;
- if $Y^+ : \bigwedge_i \psi_i \in \tau$ (respectively, $Y^- : \bigvee_i \psi_i \in \tau$), then $Y^+ : \psi_i \in \tau$ (respectively, $Y^- : \psi_i \in \tau$) for all i;
- if $Y^+ : \bigvee_i \psi_i \in \tau$ (respectively, $Y^- : \bigwedge_i \psi_i \in \tau$), then $Y^+ : \psi_i \in \tau$ (respectively, $Y^- : \psi_i \in \tau$) for some i;
- if $Y^- : \exists\boldsymbol{x}.\psi(\boldsymbol{x}\boldsymbol{b}) \in \tau$, and there is some specialization $\boldsymbol{U}_0(\boldsymbol{c}') \wedge \chi_0(\boldsymbol{b}\boldsymbol{c}') \wedge \bigwedge_k \exists\boldsymbol{x}_k.\boldsymbol{U}_k(\boldsymbol{x}_k) \wedge \chi_k(\boldsymbol{x}_k\boldsymbol{b}\boldsymbol{c}')$ such that $\boldsymbol{U}_0(\boldsymbol{c}') = U_1c_1 \wedge \cdots \wedge U_jc_j$ and $Z_i^+ : U_ic_i \in \tau$ for all $i \in \{1,\ldots,j\}$, then $Y^- : \chi_0(\boldsymbol{b}\boldsymbol{c}') \wedge \bigwedge_k \exists\boldsymbol{x}_k.\boldsymbol{U}_k(\boldsymbol{x}_k) \wedge \chi_k(\boldsymbol{x}_k\boldsymbol{b}\boldsymbol{c}') \in \tau$.

The last rule in this definition is a generalization of the rule for universal statements in the GF mosaics. It helps ensure that when we construct a model from a collection of linked, internally consistent mosaics, a CQ-shaped subformula that is asserted to be false in one mosaic, does not become true in the constructed model.

As before, a GNF X-mosaic τ is *internally consistent* if τ is downward closed and there is no atomic formula α and no $Y, Z \in \{\mathrm{L}, \mathrm{R}\}$ such that $Y^+ : \alpha \in \tau$ and $Z^- : \alpha \in \tau$.

A set of relativized GNF mosaics is said to be *saturated* if for every mosaic $\tau(\boldsymbol{c})$ in the set and for every existential requirement $Y^+ : \eta \in \tau$ for η of the form $\exists\boldsymbol{x}.U_1x_1 \wedge \cdots \wedge U_jx_j \wedge \psi(\boldsymbol{x}\boldsymbol{b})$, either (i) $Y^+ : U_ix_i[\boldsymbol{a}/\boldsymbol{x}] \in \tau$ for all $i \in \{1,\ldots,j\}$ and $Y^+ : \psi(\boldsymbol{x}\boldsymbol{b})[\boldsymbol{a}/\boldsymbol{x}] \in \tau$ for some $\boldsymbol{a} \subseteq \boldsymbol{c} \cup \boldsymbol{e}$ (so η is fulfilled in τ) or (ii) τ is linked to some mosaic τ' in the set via $Y^+ : \eta$ (so η is fulfilled in τ').

These coherence requirements lead to the following key theorem.

Theorem 28. *$\varphi_{\mathrm{L}} \wedge \neg\varphi_{\mathrm{R}}$ is satisfiable if and only if there is a saturated set of internally consistent* GNF *mosaics from $\mathcal{M}$ that contains a mosaic τ such that $\mathrm{L}^+ : \varphi_{\mathrm{L}}, \mathrm{R}^- : \varphi_{\mathrm{R}} \in \tau$.*

We give the proof in the next section.

We first give some intuition behind what the distinguished, relativized parameters represent in a saturated set of internally consistent mosaics. Roughly speaking, the relativized parameters $\boldsymbol{d}$ in an X-mosaic represent parameters that were introduced due to some existential requirement in an X-formula (and hence were originally X-relativized). This helps explain why for linked mosaics τ and τ' over $Y^+ : \exists\boldsymbol{x}.\psi(\boldsymbol{x}\boldsymbol{b})$, the "new" parameters $\boldsymbol{a}$ used in τ' always become part of the relativized parameters in the Y-mosaic τ'.

If we have a series of linked mosaics all due to X existential requirements, we want to remember all of these X-relativized parameters. As soon as we link to a $\widetilde{X}$-mosaic, however, then we can forget about these X-relativized parameters, and concentrate on the new $\widetilde{X}$-relativized parameters. Thus, the set of distinguished relativized parameters $\boldsymbol{b} \cap \boldsymbol{d}$ is carried through a link between X-mosaics, but not carried through a link between an X-mosaic and a $\widetilde{X}$-mosaic. This forgetfulness is acceptable, because in the interpolation construction for the Y-mosaic τ', we will only quantify out (and hence need relativizers for) the most recently introduced Y-parameters.[2]

B.3 Soundness and completeness of GNF mosaic method (Proof of Theorem 28)

We now give the detailed proofs of the soundness and completeness of the GNF mosaic method, in terms of the definitions given in the previous subsection.

We start with some additional definitions.

Fix parameters $\boldsymbol{c}$ and some $\boldsymbol{d} \subseteq \boldsymbol{c}$. Given a structure $\mathfrak{M}$ over $\sigma_{\mathrm{L}} \cup \sigma_{\mathrm{R}} \cup \boldsymbol{c}$ such that $|\boldsymbol{c}| \leq \mathrm{width}(\varphi_{\mathrm{L}} \wedge \neg\varphi_{\mathrm{R}})$, let $\upsilon_Y^{\mathfrak{M}}$ be the unique Y-mosaic over $\boldsymbol{c}$ relativizing $\boldsymbol{d}$ such that $\upsilon_Y^{\mathfrak{M}} := \{Y^p : \psi \mid Y^p : \psi \in \mathrm{cl}_Y(\boldsymbol{c}) \text{ and } \mathfrak{M} \models p\psi\}$, where $\mathfrak{M} \models p\psi$ means $\mathfrak{M} \models \psi$ (respectively, $\mathfrak{M} \models \neg\psi$) if $p = +$ (respectively, $p = -$). The X-mosaic $\tau := \upsilon_X^{\mathfrak{M}} \cup \left\{\widetilde{X}^p : \psi \mid \widetilde{X}^p : \psi \in \upsilon_{\widetilde{X}}^{\mathfrak{M}} \text{ and } \widetilde{X}^p : \psi \text{ is } X\text{-safe in } \upsilon_X^{\mathfrak{M}}\right\}$ is called the *X-mosaic τ of $\boldsymbol{c}$ relativizing $\boldsymbol{d}$ in $\mathfrak{M}$.*

We say an X-mosaic τ over $\boldsymbol{c}$ relativizing $\boldsymbol{d}$ is *realizable* if there is some structure $\mathfrak{M}$ such that τ is X-mosaic of $\boldsymbol{c}$ relativizing $\boldsymbol{d}$ in $\mathfrak{M}$.

For convenience in the following proofs, we also restate the definition of linked GNF mosaics. Let τ be a GNF X-mosaic over parameters $\boldsymbol{c}$ and relativizing $\boldsymbol{d}$ with $Y^+ : \eta \in \tau$ for η of the form $\exists\boldsymbol{x}.\psi(\boldsymbol{x}\boldsymbol{b})$. We say τ' is linked with τ via $Y^+ : \eta$ if

(E1) τ' is a GNF Y-mosaic using parameters $\boldsymbol{a}\boldsymbol{b}$, with $\boldsymbol{a} \cap \boldsymbol{c} = \emptyset$, and relativizing $\boldsymbol{a} \cup (\boldsymbol{b} \cap \boldsymbol{d})$ (respectively, $\boldsymbol{a}$) if $Y = X$ (respectively, $Y = \widetilde{X}$);

(E2) $Y^+ : \psi(\boldsymbol{x}\boldsymbol{b})[\boldsymbol{a}/\boldsymbol{x}] \in \tau'$;

(E3) for all $Y^p : \psi'(\boldsymbol{b}) \in \tau$, $Y^p : \psi'(\boldsymbol{b}) \in \tau'$;

(E4) for all $\widetilde{Y}^p : \psi'(\boldsymbol{b}) \in \tau$ such that $\widetilde{Y}^p : \psi'(\boldsymbol{b})$ is Y-safe in τ', $\widetilde{Y}^p : \psi'(\boldsymbol{b}) \in \tau'$;

(E5) for all $Z^p : \psi'(\boldsymbol{b}) \in \tau'$, $Z^p : \psi'(\boldsymbol{b}) \in \tau$.

We start with two propositions relating realizability to our coherency requirements for GNF mosaics (this will be used for soundness).

Proposition 29. *If τ is realizable, then τ is internally consistent.*

Proof. Let $\mathfrak{M}$ be a structure realizing the X-mosaic $\tau(\boldsymbol{c})$ relativizing $\boldsymbol{d}$. By definition, $Z^p : \psi \in \tau$ implies $\mathfrak{M} \models p\psi$.

This means it is not possible for there to be some $Y^+ : \alpha$ and $Z^- : \alpha$ in τ.

We must check that the downward closure properties are satisfied in τ. We restate them here for ease of reference.

[2] An alternative approach would be to annotate every parameter with a provenance and carry information about all of the parameters through a link. We could then place restrictions on L-formulas that use R-parameters (and vice versa).

(i) if $Y^+ : \alpha(\boldsymbol{b}) \wedge \neg\psi(\boldsymbol{b}) \in \tau$, then $Y^+ : \alpha(\boldsymbol{b}) \in \tau$ and $Y^- : \psi(\boldsymbol{b}) \in \tau$;
(ii) if $Y^- : \alpha(\boldsymbol{b}) \wedge \neg\psi(\boldsymbol{b}) \in \tau$ and $Z^+ : \alpha(\boldsymbol{b}) \in \tau$, then $Y^+ : \psi(\boldsymbol{b}) \in \tau$;
(iii) if $Y^+ : \bigwedge_i \psi_i \in \tau$ (respectively, $Y^- : \bigvee_i \psi_i \in \tau$), then $Y^+ : \psi_i \in \tau$ (respectively, $Y^- : \psi_i \in \tau$) for all i;
(iv) if $Y^+ : \bigvee_i \psi_i \in \tau$ (respectively, $Y^- : \bigwedge_i \psi_i \in \tau$), then $Y^+ : \psi_i \in \tau$ (respectively, $Y^- : \psi_i \in \tau$) for some i;
(v) if $Y^- : \exists \boldsymbol{x}.\boldsymbol{U}(\boldsymbol{x}) \wedge \psi(\boldsymbol{xb}) \in \tau$, and there is some specialization $\boldsymbol{U}_0(\boldsymbol{c}') \wedge \chi_0(\boldsymbol{bc}') \wedge \bigwedge_k \exists \boldsymbol{x}_k.\boldsymbol{U}_k(\boldsymbol{x}_k) \wedge \chi_k(\boldsymbol{x}_k\boldsymbol{bc}')$ such that $\boldsymbol{U}_0(\boldsymbol{c}') = U_1 c_1 \wedge \cdots \wedge U_j c_j$ and $Z_i^+ : U_i c_i \in \tau$ for all $i \in \{1, \ldots, j\}$, then $Y^- : \chi_0(\boldsymbol{bc}') \wedge \bigwedge_k \exists \boldsymbol{x}_k.\boldsymbol{U}_k(\boldsymbol{x}_k) \wedge \chi_k(\boldsymbol{x}_k\boldsymbol{bc}') \in \tau$.

If $Y = X$, then it is straightforward to check that these downward closure properties hold.

For instance, consider (i). If $Y^+ : \alpha(\boldsymbol{b}) \wedge \neg\psi(\boldsymbol{b}) \in \tau$, then $\mathfrak{M} \models \alpha(\boldsymbol{b}) \wedge \neg\psi(\boldsymbol{b})$. Since $\mathfrak{M} \models \alpha(\boldsymbol{b})$ and $\mathfrak{M} \models \neg\psi(\boldsymbol{b})$ and $Y^+ : \alpha(\boldsymbol{b}) \in \mathrm{cl}_Y(\boldsymbol{c})$ and $Y^- : \psi(\boldsymbol{b}) \in \mathrm{cl}_Y(\boldsymbol{c})$, this is enough to ensure that $Y^+ : \alpha(\boldsymbol{b}), Y^- : \psi(\boldsymbol{b}) \in \tau$.

Note that for (v), we make use of Lemma 26.

If $Y = \widetilde{X}$, then we use a similar argument, but in order to conclude that the required formulas appear in τ, we must also show that these formulas are X-safe in $\upsilon_X^{\mathfrak{M}}$.

For instance, consider (ii). Assume $Y^- : \alpha(\boldsymbol{b}) \wedge \neg\psi(\boldsymbol{b}) \in \tau$ and $Z^+ : \alpha(\boldsymbol{b}) \in \tau$. Then we know that $\mathfrak{M} \models \alpha(\boldsymbol{b})$ and $\mathfrak{M} \models \alpha(\boldsymbol{b}) \rightarrow \psi(\boldsymbol{b})$, hence $\mathfrak{M} \models \psi(\boldsymbol{b})$. This means $Y^+ : \psi(\boldsymbol{b}) \in \upsilon_Y^{\mathfrak{M}}$. It remains to show that this formula is X-safe in $\upsilon_X^{\mathfrak{M}}$.

If $Y = Z$, then $Z^+ : \alpha(\boldsymbol{b}) \in \tau$ implies that $Z^+ : \alpha(\boldsymbol{b})$ is X-safe in $\upsilon_X^{\mathfrak{M}}$, so $Y^+ : \psi(\boldsymbol{b})$ is also X-safe in $\upsilon_X^{\mathfrak{M}}$.

If $Y = \widetilde{Z}$, then α appears positively in an X-formula (since $Z = X$), and negatively in a $\widetilde{X}$-formula (since $Y = \widetilde{X}$). Making use of the fact that polarities and occurrences are preserved in moving to mosaics, this means that α is an X-cguard. Moreover, $Y^- : \alpha(\boldsymbol{b}) \wedge \neg\psi(\boldsymbol{b}) \in \tau$ implies that $\boldsymbol{b} \cap \boldsymbol{d}$ is X-crelativized in τ. This means $\boldsymbol{b}$ is X-cguarded by $X^+ : \alpha(\boldsymbol{b})$ in τ, so $Y^+ : \psi(\boldsymbol{b})$ is X-safe in $\upsilon_X^{\mathfrak{M}}$.

Now consider (v), still for $Y = \widetilde{X}$. Assume $Y^- : \exists \boldsymbol{x}.\boldsymbol{U}(\boldsymbol{x}) \wedge \psi(\boldsymbol{xb}) \in \tau$ and there is some specialization $Y^- : \boldsymbol{U}_0(\boldsymbol{c}') \wedge \chi_0(\boldsymbol{bc}') \wedge \bigwedge_k \exists \boldsymbol{x}_k.\boldsymbol{U}_k(\boldsymbol{x}_k) \wedge \chi_k(\boldsymbol{x}_k\boldsymbol{bc}')$ such that $\boldsymbol{U}_0(\boldsymbol{c}') = U_1 c_1 \wedge \cdots \wedge U_j c_j$ and $Z_i^+ : U_i c_i \in \tau$ for all $i \in \{1, \ldots, j\}$.

We first claim that $Y^- : \chi_0(\boldsymbol{bc}') \wedge \bigwedge_{k \in \{1,\ldots,l\}} \exists \boldsymbol{x}_k.\boldsymbol{U}_j(\boldsymbol{x}_k) \wedge \chi_k(\boldsymbol{x}_k\boldsymbol{bc}')$ is X-safe in $\upsilon_X^{\mathfrak{M}}$. Since the formula is negative, we only need to ensure that all parameters in $\boldsymbol{bc}' \cap \boldsymbol{d}$ are X-crelativized in τ. We already know that the parameters in $\boldsymbol{b} \cap \boldsymbol{d}$ are X-crelativized since $Y^- : \exists \boldsymbol{x}.\psi(\boldsymbol{xb}) \in \tau$. Now consider parameters $c_i \in \boldsymbol{c}'$. If $Z_i = \widetilde{X}$ and $c_i \in \boldsymbol{d}$, then $Z_i^+ : U_i c_i \in \tau$ implies that c_i is X-crelativized in τ. If $Z_i = X$ and $c_i \in \boldsymbol{d}$, then $Z_i^+ : U_i c_i \in \tau$ is the X-crelativizer for c_i (U_i is an X-catom, since occurrences are preserved in moving to mosaics).

Assume for the sake of contradiction that $Y^- : \chi_0(\boldsymbol{bc}') \wedge \bigwedge_{k \in \{1,\ldots,l\}} \exists \boldsymbol{x}_k.\boldsymbol{U}_j(\boldsymbol{x}_k) \wedge \chi_k(\boldsymbol{x}_k\boldsymbol{bc}') \notin \tau$. Since this formula is in $\mathrm{cl}_Y(\boldsymbol{c})$ and is X-safe in $\upsilon_X^{\mathfrak{M}}$, the only way it is not in τ is if $\mathfrak{M} \models \chi_0(\boldsymbol{bc}') \wedge \bigwedge_{k \in \{1,\ldots,l\}} \exists \boldsymbol{x}_k.\boldsymbol{U}_j(\boldsymbol{x}_k) \wedge \chi_k(\boldsymbol{x}_k\boldsymbol{bc}')$. But this would mean $\mathfrak{M} \models \boldsymbol{U}_0(\boldsymbol{c}') \wedge \chi_0(\boldsymbol{bc}') \wedge \bigwedge_{k \in \{1,\ldots,l\}} \exists \boldsymbol{x}_k.\boldsymbol{U}_j(\boldsymbol{x}_k) \wedge \chi_k(\boldsymbol{x}_k\boldsymbol{bc}')$ (since $Z_i^+ : U_i c_i \in \tau$). By Lemma 26, every specialization of a formula implies the original formula, contradicting the fact that $\mathfrak{M} \models \neg\exists \boldsymbol{x}.\boldsymbol{U}(\boldsymbol{x}) \wedge \psi(\boldsymbol{xb})$. □

Proposition 30. *If υ is realizable, then there is a saturated set $\mathcal{S} \subseteq \mathcal{M}$ of realizable mosaics that includes υ.*

Proof. We inductively construct sets $\mathcal{S}_i$, ensuring at each stage i that $\mathcal{S}_i$ contains υ and other realizable mosaics, and that any existential requirement in $\mathcal{S}_{i-1}$ is fulfilled in $\mathcal{S}_i$. This is enough to ensure that the set $\mathcal{S} := \bigcup \mathcal{S}_i$ is a saturated set of realizable mosaics that includes υ.

At stage $i = 1$, we set $\mathcal{S}_1 := \{\upsilon\}$, which is realizable by assumption.

At stage $i > 1$, if every existential requirement in $\mathcal{S}_{i-1}$ is fulfilled, then set $\mathcal{S}_i := \mathcal{S}_{i-1}$.

Otherwise, consider some X-mosaic $\tau(\boldsymbol{c})$ relativizing $\boldsymbol{d}$ in $\mathcal{S}_{i-1}$ with $Y^+ : \eta \in \tau$ for η of the form $\exists \boldsymbol{x}.\psi(\boldsymbol{xb})$ and $\mathrm{par}(\eta) = \boldsymbol{b}$ for which there is no $\boldsymbol{a} \subseteq \boldsymbol{c} \cup \boldsymbol{e}$ with $Y^+ : \psi(\boldsymbol{xb})[\boldsymbol{a}/\boldsymbol{x}] \in \tau$, and for all τ' such that $\tau \rightarrow_{Y^+:\eta} \tau'$, we have $\tau' \notin \mathcal{S}_{i-1}$.

Let $\mathfrak{M}$ be a structure realizing τ. Take parameters $\boldsymbol{a} \subseteq P$ corresponding to $\boldsymbol{x}$ such that $\boldsymbol{a} \cap \boldsymbol{c} = \emptyset$. This is possible since $|\boldsymbol{a}| \leq \mathrm{width}(\varphi_{\mathrm{L}} \wedge \neg\varphi_{\mathrm{R}})$ and $|P| = 2 \cdot \mathrm{width}(\varphi_{\mathrm{L}} \wedge \neg\varphi_{\mathrm{R}})$. Since $\mathfrak{M} \models \tau$, there are elements $\boldsymbol{a}' \subseteq M$, and an expansion $\mathfrak{M}'$ of $\mathfrak{M}$ with the interpretation $\boldsymbol{a}^{\mathfrak{M}'} := \boldsymbol{a}'$ such that $\mathfrak{M}' \models \psi(\boldsymbol{ab})$.

Let τ' be the Y-mosaic of $\boldsymbol{ab}$ relativizing $\boldsymbol{a} \cup (\boldsymbol{b} \cap \boldsymbol{d})$ (respectively, $\boldsymbol{a}$) in $\mathfrak{M}'$ if $Y = X$ (respectively, $Y = \widetilde{X}$). This means τ' is realizable.

We also claim that $\tau \rightarrow_{Y^+:\eta} \tau'$. (E1) is satisfied by construction. By the choice of $\boldsymbol{a}^{\mathfrak{M}'}$, we have $\mathfrak{M}' \models \psi(\boldsymbol{ab})$, so $Y^+ : \psi(\boldsymbol{ab}) \in \tau'$ as required by (E2).

We must now check that formulas with shared parameters $\boldsymbol{b}$ are preserved between τ and τ' as described by (E3)–(E5).

(E3) and (E4) follow by the definition of realizable mosaics, and the fact that $\mathfrak{M}$ and $\mathfrak{M}'$ are identical with respect to formulas that use only parameters from $\boldsymbol{b}$.

It remains to check (E5). Assume there is some $Z^p : \psi'(\boldsymbol{b}) \in \tau'$. We have $\mathfrak{M}' \models p\psi'$ since τ' is realized by $\mathfrak{M}'$. Recall that τ is an X-mosaic, τ' is a Y-mosaic, and ψ' is a Z^p-formula. The result is immediate if $Z^p = X^p$. Otherwise, if $Z^p = \widetilde{X}^p$ it suffices to show that $Z^p : \psi'(\boldsymbol{b})$ is X-safe in τ.

- Assume $Z^p = \widetilde{X}^p$ and $X = Y$. Then $\mathrm{par}(\psi') \cap \boldsymbol{d}$ must be X-crelativized in τ', and hence X-crelativized in τ. Moreover, if $p = +$, then since $\widetilde{X}^+ : \psi' \in \tau'$ and τ' is an X-mosaic, $\mathrm{par}(\psi')$ must be X-cguarded in τ' by some $X^+ : \exists \boldsymbol{w}.\boldsymbol{U}(\boldsymbol{w}) \wedge \beta(\boldsymbol{wa}'\boldsymbol{b}')$ where $\boldsymbol{a}' \subseteq \boldsymbol{a}$ and $\mathrm{par}(\psi') \subseteq \boldsymbol{b}' \subseteq \boldsymbol{b}$ and such that parameters in $\boldsymbol{a}'\boldsymbol{b}' \cap (\boldsymbol{a} \cup (\boldsymbol{b} \cap \boldsymbol{d}))$ are X-crelativized in τ'. But this means that $\mathfrak{M} \models \exists \boldsymbol{wx}.\boldsymbol{U}(\boldsymbol{w}) \wedge \boldsymbol{U}'(\boldsymbol{x}) \wedge \beta(\boldsymbol{wxb}')$ for some $\boldsymbol{U}'(\boldsymbol{x})$ conjunction of X-crelativizers. Hence, $X^+ : \exists \boldsymbol{wx}.\boldsymbol{U}(\boldsymbol{w}) \wedge \boldsymbol{U}'(\boldsymbol{x}) \wedge \beta(\boldsymbol{wxb}') \in \tau$. Since the parameters in $\boldsymbol{b}' \cap \boldsymbol{d}$ are X-crelativized in τ, $\mathrm{par}(\psi')$ is X-cguarded in τ. Overall, this means that $\widetilde{X}^p : \psi'$ is X-safe in τ.
- Assume $Z^p = \widetilde{X}^p$ and $X \neq Y$. Then $\mathrm{par}(\psi')$ must be X-cguarded in τ and $\mathrm{par}(\psi') \cap \boldsymbol{d}$ must be X-crelativized in τ, since $Y^+ : \eta \in \tau$ with $\mathrm{par}(\eta) = \boldsymbol{b}$ and $X \neq Y$. This is enough to ensure that $\widetilde{X}^p : \psi'$ is X-safe in τ.

We have constructed a realizable Y-mosaic τ' such that $\tau \to_{Y^+:\eta} \tau'$, so we add τ' to $\mathcal{S}_i$. Repeating this procedure for all such τ (and all such $Y^+ : \eta \in \tau$) in $\mathcal{S}_{i-1}$, we will end up with a set $\mathcal{S}_i$ of realizable mosaics such that any existential requirement in $\mathcal{S}_{i-1}$ is fulfilled in $\mathcal{S}_i$. $\square$

We now aim towards proving a proposition useful for completeness (a sort of converse of Proposition 30). In order to do this, we first prove lemmas about the transfer of formulas between linked mosaics.

Lemma 31 (Backwards transfer property). *Let $\tau_0 \to \cdots \to \tau_k$ be linked mosaics such that every τ_i includes parameters $\boldsymbol{b}$. If $Z^p : \psi(\boldsymbol{b}) \in \tau_k$, then $X^p : \psi(\boldsymbol{b}) \in \tau_0$.*

Proof. We proceed by induction on k. The base case for $k = 0$ is immediate.

Otherwise, assume $k > 0$. Consider $\tau_{k-1} \to \tau_k$. By definition (E5), $Z^p : \psi(\boldsymbol{b}) \in \tau_k$ implies $Z^p : \psi(\boldsymbol{b}) \in \tau_{k-1}$ (since $\boldsymbol{b}$ is shared by τ_{k-1} and τ_k), so we can apply the inductive hypothesis to $\tau_0 \to \cdots \to \tau_{k-1}$ to get the result. $\square$

Lemma 32 (Forward relativizer transfer property). *Let $\tau_0 \to \cdots \to \tau_k$ be linked mosaics such that every τ_i includes parameters $\boldsymbol{b}$. If $Z^+ : Ub \in \tau_0$ for $b \in \boldsymbol{b}$, then $Z^+ : Ub \in \tau_k$.*

Proof. We proceed by induction on k. The base case for $k = 0$ is immediate.

Otherwise, assume $k > 0$. Consider $\tau_0 \to \tau_1$ where τ_0 is an X-mosaic relativizing $\boldsymbol{d}$ and τ_1 is a Y-mosaic.

- If $Z = Y$, then (E3) implies that $Z^+ : Ub \in \tau_1$.
- If $Z = \widetilde{Y}$ and $Y = X$, then $\widetilde{Y}^+ : Ub \in \tau_0$ implies that $b \cap \boldsymbol{d}$ is X-crelativized in τ_0.
  If $b \notin \boldsymbol{d}$, then $\widetilde{Y}^+ : Ub$ trivially satisfies the requirements for Y-safety in τ_1.
  Otherwise, if $b \in \boldsymbol{d}$, then b is X-crelativized by some $X^+ : U'b \in \tau_0$. Since $X^+ : U'b$ is trivially Y-safe in τ_1, (E3) ensures that $X^+ : U'b \in \tau_1$. Hence, b is Y-crelativized in τ_1. Moreover, b is trivially Y-cguarded (since there is only one parameter). This means that $\widetilde{Y}^+ : Ub$ is Y-safe in τ_1.
  Hence, in either case, we can apply (E4) to ensure that $Z^+ : Ub \in \tau_1$.
- If $Z = \widetilde{Y}$ and $Y = \widetilde{X}$, then it is trivially Y-safe: it is Y-cguarded because there is only one parameter, and it does not need to be Y-crelativized (since no shared parameters are relativized when moving from an X to $\widetilde{X}$ mosaic). Hence, we can apply (E4) to get $Z^+ : Ub \in \tau_1$.

Since we have guaranteed that $Z^+ : Ub \in \tau_1$, we can apply the inductive hypothesis to $\tau_1 \to \cdots \to \tau_k$ to get the result. $\square$

Lemma 33 (Forward negative transfer property). *Let $\tau_0 \to \cdots \to \tau_k$ be linked mosaics such that every τ_i includes parameters $\boldsymbol{b}$. If $Z^- : \psi(\boldsymbol{b}) \in \tau_0$, then $Z^- : \psi(\boldsymbol{b}) \in \tau_k$.*

Proof. We proceed by induction on k. If $k = 0$, then the result is immediate.

Otherwise, assume $k > 0$. Consider $\tau_0 \to \tau_1$ where τ_0 is an X-mosaic relativizing $\boldsymbol{d}$ and τ_1 is a Y-mosaic. In order to apply the inductive hypothesis, we must show that $Z^- : \psi(\boldsymbol{b}) \in \tau_1$.

- If $Z = Y$, then $Z^- : \psi(\boldsymbol{b}) \in \tau_1$ by (E3).
- If $Z = \widetilde{Y}$ and $Y = X$, then $\widetilde{Y}^- : \psi(\boldsymbol{b}) \in \tau_0$ implies that $\boldsymbol{b} \cap \boldsymbol{d}$ is X-crelativized in τ_0.
  By the previous lemma, we know that $\boldsymbol{b} \cap \boldsymbol{d}$ is also X-crelativized in τ_1 (i.e. we know that all of the X-crelativizers in τ_0 are in τ_1). This is enough to ensure that $\widetilde{Y}^- : \psi(\boldsymbol{b})$ is Y-safe in τ_1, so (E4) guarantees that $Z^- : \psi(\boldsymbol{b}) \in \tau_1$.
- If $Z = \widetilde{Y}$ and $Y = \widetilde{X}$, then it is trivially Y-safe: no parameters need to be Y-crelativized (since no shared parameters are crelativized when moving from an X to $\widetilde{X}$ mosaic). Hence, we can apply (E4) to get $Z^- : \psi(\boldsymbol{b}) \in \tau_1$

Since $Z^+ : Ub \in \tau_1$, we can apply the inductive hypothesis to $\tau_1 \to \cdots \to \tau_k$ to get the result. $\square$

We are now ready to show that whenever there is a set of saturated mosaics, we can get a satisfying model (completeness).

Proposition 34. *Let $\mathcal{S} \subseteq \mathcal{M}$ be a saturated set of internally consistent mosaics. Then for all $\tau(\boldsymbol{c}) \in \mathcal{S}$, there is a model $\mathfrak{M}$ over $\sigma_{\mathrm{L}} \cup \sigma_{\mathrm{R}} \cup \boldsymbol{c}$ such that for all $X^p : \psi \in \tau$, $\mathfrak{M} \models p\psi$.*

Proof. Fix $\tau \in \mathcal{S}$. Note that τ is internally consistent by assumption. Our goal is to construct a model $\mathfrak{M}$ witnessing the satisfiability of the formulas in τ. We will use the mosaics in $\mathcal{S}$ as building blocks.

We build inductively a tree decomposition T of $\mathfrak{M}$, together with a map f from nodes v in the tree decomposition to a mosaic in $\mathcal{S}$. We also think of f as a map from the elements in a node v, to the parameters in the corresponding mosaic.

At stage 0, let T_0 consist of a single node v_0 containing a copy of the parameters in τ and constants $\boldsymbol{e}$, and set $f(v_0) = \tau$.

At stage i, consider the set

$$S = \{(v, Y^+ : \eta) : v \text{ is a leaf in } T_{i-1} \text{ and } Y^+ : \eta(f(\boldsymbol{b})) \in f(v) \text{ is not fulfilled in } f(v) \\ \text{for } \eta \text{ of the form } \exists \boldsymbol{x}.\psi'(\boldsymbol{x} f(\boldsymbol{b}))\}.$$

For each $(v, Y^+ : \eta) \in S$, we know that there is some $\tau' \in \mathcal{S}$ such that $f(v) \to_{Y^+:\eta} \tau'$ (since $\mathcal{S}$ is saturated). We construct T_i by extending T_{i-1} according to the following procedure: for each $(v, Y^+ : \eta) \in S$, we construct a child v' of v with $T(v') = \boldsymbol{a} \cup \boldsymbol{b} \cup \boldsymbol{e}$ for new elements $\boldsymbol{a}$, and define $f(v') = \tau'$ such that $Y^+ : \psi'(f(\boldsymbol{ab})) \in \tau'$.

Let T be the limit of this process, and let $\mathfrak{M}$ be the structure with elements $M = \bigcup_{v \in T} T(v)$, and such that $\mathfrak{M} \models \alpha(\boldsymbol{a})$ iff there exists v such that $\boldsymbol{a}$ appears in $T(v)$ and $X^+ : \alpha(f(\boldsymbol{a})) \in f(v)$.

We must now show that for all v in T and for all $\boldsymbol{a}$ in $T(v)$, if $Z^p : \psi(f(\boldsymbol{a})) \in f(v)$, then $\mathfrak{M} \models p\psi(\boldsymbol{a})$.

The proof is by induction on the structure of ψ.

- Assume $Z^+ : \psi$ is an atom. Then $\mathfrak{M} \models \psi$ by construction.
- Assume $Z^- : \psi$ is an atom $\alpha(f(\boldsymbol{a}))$, and suppose for the sake of contradiction that $\mathfrak{M} \models \alpha(\boldsymbol{a})$. Then there is some w such that $X^+ : \alpha(\boldsymbol{a}) \in f(w)$ and $f(\boldsymbol{a})$ appears in $f(w')$ for all w' on the path between v and w. Using the backwards transfer property (Lemma 31), there must be some node containing both $Z^- : \alpha(f(\boldsymbol{a}))$ and $X^+ : \alpha(f(\boldsymbol{a}))$, which contradicts internal consistency of mosaics in $\mathcal{S}$.
- Assume $Z^p : \psi$ is of the form $\bigwedge_i \psi_i$ or $\bigvee_i \psi_i$. The result follows using internal consistency and the inductive hypothesis.
- Assume $Z^+ : \psi$ is $\alpha(f(\boldsymbol{b})) \wedge \neg\psi'(f(\boldsymbol{b}))$. By internal consistency, $Z^+ : \alpha(f(\boldsymbol{b})) \in f(v)$ and $Z^- : \psi'(f(\boldsymbol{b})) \in f(v)$. The result follows using the inductive hypothesis.
- Assume $Z^- : \psi$ is $\alpha(f(\boldsymbol{b})) \wedge \neg\psi'(f(\boldsymbol{b}))$, and suppose for the sake of contradiction that $\mathfrak{M} \models \alpha(\boldsymbol{b}) \wedge \neg\psi'(\boldsymbol{b})$. Then there is some w such that $X^+ : \alpha(f(\boldsymbol{b})) \in f(w)$ and $f(\boldsymbol{b})$ appears in $f(w')$ for all w' on the path between v and w. Using the backwards transfer property (Lemma 31), there must be some node containing both $Z^- : \alpha(f(\boldsymbol{b})) \wedge \neg\psi'(f(\boldsymbol{b}))$ and $X^+ : \alpha(f(\boldsymbol{b}))$. By internal consistency, this implies $Z^+ : \psi'(f(\boldsymbol{b}))$ is also in this node, so $\mathfrak{M} \models \psi'(\boldsymbol{b})$ by the inductive hypothesis, which contradicts $\mathfrak{M} \models \alpha(\boldsymbol{b}) \wedge \neg\psi'(\boldsymbol{b})$.
- Assume $Z^+ : \psi$ is of the form $Z^+ : \exists \boldsymbol{x}.\psi'(\boldsymbol{x} f(\boldsymbol{b}))$. Then either $Z^+ : \psi'(f(\boldsymbol{ab})) \in f(v)$ for some $\boldsymbol{a} \subseteq T(v)$, or there is some child w of v with $f(v) \to_{Z^+:\psi} f(w)$ such that $Z^+ : \psi'(f(\boldsymbol{ab})) \in f(w)$ and $\boldsymbol{a} \subseteq T(w)$. In either case, the inductive hypothesis implies $\mathfrak{M} \models \psi'(\boldsymbol{ab})$, so $\mathfrak{M} \models \psi$.
- Assume $Z^- : \psi$ is of the form $Z^- : \exists \boldsymbol{x}.\psi'(\boldsymbol{x} f(\boldsymbol{b}))$, and assume for the sake of contradiction that $\mathfrak{M} \models \exists \boldsymbol{x}.\psi'(\boldsymbol{xb})$.
  By Lemma 25, there is a node w with $\boldsymbol{bc} \subseteq T(w)$ and a non-trivial specialization $\phi'(\boldsymbol{yz})$ of ψ such that

  $$\mathfrak{M}, \boldsymbol{bc} \models (\boldsymbol{U}_0(\boldsymbol{z}) \wedge \chi_0(\boldsymbol{yz})) \wedge \bigwedge_{j \in \{1,\dots,k\}} \exists \boldsymbol{x}_j.\boldsymbol{U}_j(\boldsymbol{x}_j) \wedge \chi_j(\boldsymbol{x}_j \boldsymbol{yz})$$

  By the backward/forward transfer properties (Lemmas 31 and 33), $Z^- : \psi \in f(w)$.
  Let $\boldsymbol{U}_0(\boldsymbol{z}) = U_1 z_1 \wedge \dots \wedge U_l z_l$. Since $\mathfrak{M} \models \boldsymbol{U}_0(\boldsymbol{c})$, for each $i \in \{1, \dots, l\}$, there is a node w_i such that $Z_i^+ : U_i\, f(c_i) \in f(w_i)$, and the parameter $f(c_i)$ is in every mosaic on the path between w_i and w. By the backward transfer property and forward relativizer transfer properties (Lemmas 31 and 32), $Z_i^+ : U_i\, f(c_i) \in f(w)$.
  By internal consistency, $Z^- : \phi'[\boldsymbol{b}/\boldsymbol{y}, \boldsymbol{c}/\boldsymbol{z}] \in f(w)$.
  Likewise, by internal consistency, we have that $Z^- : \boldsymbol{U}_0(f(\boldsymbol{c})) \wedge \chi_0(f(\boldsymbol{bc})) \in f(w)$ or $Z^- : \exists \boldsymbol{x}_j.\boldsymbol{U}_j(\boldsymbol{x}_j) \wedge \chi_j(\boldsymbol{x}_j f(\boldsymbol{bc}))) \in f(w)$ for some j. Since it is a non-trivial specialization, we can apply the inductive hypothesis to this formula in $f(w)$, and get a contradiction. We remark that this is the only case where we need both forward and backward transfer properties, since we need to be able to move $Z^- : \psi$ (and the relativizing formulas) to the node w related to the the non-trivial specialization.

$\square$

We are now ready to prove Theorem 28. For soundness we must show:

> If $\varphi_{\mathrm{L}} \wedge \neg\varphi_{\mathrm{R}}$ is satisfiable, then there is a saturated set of internally consistent GNF mosaics from $\mathcal{M}$ that contains some τ such that $\mathrm{L}^+ : \varphi_{\mathrm{L}}, \mathrm{R}^- : \varphi_{\mathrm{R}} \in \tau$.

Let $\mathrm{par}(\varphi_{\mathrm{L}}) = \mathrm{par}(\varphi_{\mathrm{R}}) = \boldsymbol{z}$. Since $\varphi_{\mathrm{L}} \wedge \neg\varphi_{\mathrm{R}}$ is satisfiable, there is a model $\mathfrak{M}$ over signature $\sigma_{\mathrm{L}} \cup \sigma_{\mathrm{R}} \cup \boldsymbol{z}$ such that $\mathfrak{M} \models \varphi_{\mathrm{L}} \wedge \neg\varphi_{\mathrm{R}}$. Let τ be the L-mosaic of $\boldsymbol{z}$ relativizing $\emptyset$ in $\mathfrak{M}$, i.e. the set of formulas from $\mathrm{cl}_{\mathrm{L}}(\boldsymbol{z})$ that are true in $\mathfrak{M}$, and the set of formulas from $\mathrm{cl}_{\mathrm{R}}(\boldsymbol{z})$ that are true in $\mathfrak{M}$ and satisfy the L-safety conditions on R-formulas in an L-mosaic. Since $\varphi_{\mathrm{L}} \wedge \neg\varphi_{\mathrm{R}}$ holds in $\mathfrak{M}$, τ contains $\mathrm{L}^+ : \varphi_{\mathrm{L}}$. Since $\mathrm{R}^- : \varphi_{\mathrm{R}}$ is trivially L-safe (since it is a negative formula, and there are no relativized parameters), τ also contains $\mathrm{R}^- : \varphi_{\mathrm{R}}$. Since τ is a realizable mosaic, Proposition 30 implies that there is a saturated set $\mathcal{S}$ of realizable mosaics including τ. Moreover, by Proposition 29, $\mathcal{S}$ must contain only internally consistent mosaics, as desired.

For completeness of the mosaic method, we must show:

> If there is a saturated set of internally consistent mosaics from $\mathcal{M}$ containing τ such that $\mathrm{L}^+ : \varphi_{\mathrm{L}}, \mathrm{R}^- : \varphi_{\mathrm{R}} \in \tau$, then $\varphi_{\mathrm{L}} \wedge \neg\varphi_{\mathrm{R}}$ is satisfiable.

This follows immediately from Proposition 34.

B.4 Bound on number of relativized GNF mosaics

Given $\mathbb{U}$-relativized GNF formulas φ_{L} and φ_{R} in weak GN-normal form, we can prove a doubly exponential upper bound for the number of relativized GNF mosaics for $\varphi_{\mathrm{L}} \wedge \neg\varphi_{\mathrm{R}}$ in $\mathcal{M}$. The following proposition states this bound more precisely.

Proposition 35. *There is a polynomial function p such that if $\varphi := \varphi_{\mathrm{L}} \wedge \neg\varphi_{\mathrm{R}}$ for φ_{L} and φ_{R} in weak* GN*-normal form without equality, and $n = |\varphi_{\mathrm{L}}| + |\varphi_{\mathrm{R}}|$, $w = \mathrm{width}(\varphi)$, and $h = \mathrm{rank}_{\mathrm{CQ}}(\varphi)$, then the number of relativized* GNF *mosaics in $\mathcal{M}$ for φ is at most $2^{p(n) \cdot 2^{p(hw)}}$.*

Proof. Fix $\varphi := \varphi_{\mathrm{L}} \wedge \neg\varphi_{\mathrm{R}}$ with $n = |\varphi_{\mathrm{L}}| + |\varphi_{\mathrm{R}}|$, $w = \mathrm{width}(\varphi)$, and $h = \mathrm{rank}_{\mathrm{CQ}}(\varphi)$. Let $l := |\mathrm{con}(\varphi)|$. Note that $w, h, l \leq n$.

We first calculate the size of $\mathrm{cl}(Y^p : \varphi_Y)$ (where $p = +$ if $Y = \mathrm{L}$ and $p = -$ if $Y = \mathrm{R}$).

There are at most $2n$ polarity-labelled subformulas of φ_Y. However, $\mathrm{cl}(Y^p : \varphi_Y)$ contains additional formulas according to the definition of GNF closure.

For each positive atomic formula α in $Y^p : \varphi_Y$, there are at most w free variables, and hence at most 2^w ways to choose variables for quantification. There can be at most n predicates in $\mathbb{U}$, and at most n^w ways to relativize the quantified variables using these unary predicates. Hence, for each positive atomic formula α in $Y^p : \varphi_Y$, there are at most $2^w n^w$ additional formulas that may be added.

For each negative CQ-shaped formula ϕ in $Y^p : \varphi_Y$, there are at most 2^h CQ-shaped subformulas ϕ' obtained by choosing some subset of the conjuncts in ϕ. In each ϕ', there are at most 2^w choices of the inside variables, and w^w ways to partition the outside variables resulting in specializations ϕ'' of ϕ'. Notice that these specializations ϕ'' only have CQ-shaped subformulas resulting from taking some subset of the conjuncts in the original CQ-shaped formula ϕ, so the specializations of these subformulas have already been accounted for. Hence, for each negative CQ-shaped subformula in $Y^p : \varphi_Y$, there are most $2^h 2^w w^w$ specializations that may be added.

This means that the size of $\mathrm{cl}(Y^p : \varphi_Y)$ is at most $2n(2^w n^w + 2^h 2^w w^w)$.

Now recall that P is a fixed set of parameters of size at most $2w$. Each formula in $\mathrm{cl}(Y^p : \varphi_Y)$ has at most w variables that can be mapped to parameters from P or constants from $\boldsymbol{e}$, so the size of $\mathrm{cl}_Y(P)$ is $M := 2n(2^w n^w + 2^h 2^w w^w)(2w + l)^w$.

Every mosaic in $\mathcal{M}$ contains a subset of the formulas in $\mathrm{cl}_{\mathrm{L}}(P) \cup \mathrm{cl}_{\mathrm{R}}(P)$, is labelled L or R, and has some subset of its at most w parameters distinguished as the relativized parameters. Hence, $|\mathcal{M}| \leq 2 \cdot 2^w \cdot 2^{2M} \leq 2^{p(n)\cdot 2^{p(hw)}}$ for some polynomial function p that is independent of φ_{L} and φ_{R}. ☐

C. Interpolation using Relativized GNF Mosaics

In this section, we present details for the interpolation results stated in Theorem 9.

We first prove a mosaic interpolation lemma (like Lemma 4) for the relativized GNF mosaics defined in Appendix B. We then discuss how to obtain the interpolation results stated in Theorem 9 from these mosaic interpolants.

Throughout this section, we continue to restrict our attention to input formulas φ_{L} and φ_{R} that do not use equality. Please see Appendix D for remarks on how to extend this approach to deal with equality. We remark that even with equality-free input, the interpolants may use equality (for guards of negated formulas with at most one free variable).

C.1 Mosaic interpolants for relativized GNF mosaics

Fix $\mathbb{U}$-relativized GNF formulas φ_{L} and φ_{R} without equality as in the previous section. That is, φ_{L} and φ_{R} are $\mathbb{U}$-relativized formulas in weak GN-normal form and without equality over signatures σ_{L} and σ_{R}. Moreover, $\mathrm{par}(\varphi_{\mathrm{L}}) = \mathrm{par}(\varphi_{\mathrm{R}})$.

Let $\mathcal{M}$ be the set of relativized GNF mosaics for $\varphi_{\mathrm{L}} \wedge \neg\varphi_{\mathrm{R}}$ as defined in Appendix B.

Consider the mosaic elimination procedure over $\mathcal{M}$. We write $\mathcal{N}'$ for the set of mosaics for $\varphi_{\mathrm{L}} \wedge \neg\varphi_{\mathrm{R}}$ that are eliminated using this procedure, and $\mathcal{N}_i$ for the set of mosaics that have been removed by stage i. By Theorem 28, we know that every mosaic containing $\mathrm{L}^+ : \varphi_{\mathrm{L}}$ and $\mathrm{R}^- : \varphi_{\mathrm{R}}$ must be in $\mathcal{N}'$.

For an X-mosaic $\tau(\boldsymbol{c})$, let

$$\tau_X := \bigwedge_{X^p : \psi \in \tau} p\psi \quad \wedge \bigwedge_{X^p : \psi \in \mathrm{cl}_X(\boldsymbol{c}) \setminus \tau} \neg p\psi \qquad \text{and} \qquad \tau_{\widetilde{X}} := \bigwedge_{\widetilde{X}^p : \psi \in \tau} p\psi \quad \wedge \bigwedge_{\substack{\widetilde{X}^p : \psi \in \mathrm{cl}_{\widetilde{X}}(\boldsymbol{c}) \setminus \tau \\ \text{s.t. } \widetilde{X}^p \,:\, \psi \text{ is } X\text{-safe in } \tau}} \neg p\psi \,.$$

That is, τ_Y asserts the truth of the Y-formulas in τ, and the negation of the Y-formulas that are not in τ (but satisfy the conditions for being in τ in terms of X-safety if $Y = \widetilde{X}$).

We now define mosaic interpolants for every X-mosaic $\tau(\boldsymbol{c}) \in \mathcal{N}'$.

Lemma 36. *For each* GNF *X-mosaic $\tau(\boldsymbol{c}) \in \mathcal{N}_i$ relativizing $\boldsymbol{d}$, we can construct a DAG representation of a formula θ_X^τ such that*

(Imp) $\tau_X \models \theta_X^\tau$ *and* $\theta_X^\tau \models \neg\tau_{\widetilde{X}}$*;*

(Occ) $\mathrm{occ}(\theta_X^\tau) \subseteq \mathrm{occ}(\varphi_{\mathrm{L}}) \cap \mathrm{occ}(\varphi_{\mathrm{R}})$ *if* $X = \mathrm{L}$*,*
$\mathrm{occ}(\neg\theta_X^\tau) \subseteq \mathrm{occ}(\varphi_{\mathrm{L}}) \cap \mathrm{occ}(\varphi_{\mathrm{R}})$ *if* $X = \mathrm{R}$*;*
$\mathrm{con}(\theta_X^\tau) \subseteq \mathrm{con}(\varphi_{\mathrm{L}}) \cup \mathrm{con}(\varphi_{\mathrm{R}})$*;*

(Par) $\mathrm{par}(\theta_X^\tau) \subseteq \boldsymbol{c}$ *and* $\mathrm{par}(\theta_X^\tau) \cap \boldsymbol{d}$ *is X-crelativized in τ;*

(Rel) θ_X^τ *is in $\mathbb{U}$-relativized* GNF *(when treating $\top$ and $\bot$ as atomic formulas), even when parameters are viewed as free variables.*

Moreover, if $n = |\varphi_{\mathrm{L}}| + |\varphi_{\mathrm{R}}|$, $w = \mathrm{width}(\varphi_{\mathrm{L}} \wedge \neg\varphi_{\mathrm{R}})$, and $h = \mathrm{rank}_{\mathrm{CQ}}(\varphi_{\mathrm{L}} \wedge \neg\varphi_{\mathrm{R}})$ then there is a DAG representation of $\Theta_i := \{\theta_Z^\tau : \tau \text{ is a } Z\text{-mosaic in } \mathcal{N}_i\}$ with $|\Theta_i| \leq i \cdot 2^{q(n)\cdot 2^{q(hw)}}$ where q is a polynomial functions independent of φ_{L} and φ_{R}.

Proof. We fix some notation.

For any tuple $\boldsymbol{b}$ of parameters that is X-cguarded in τ, we define a formula $\mathrm{gdd}_X^\tau(\boldsymbol{b})$. If $\boldsymbol{b} = \emptyset$, then we set $\mathrm{gdd}_X^\tau(\boldsymbol{b}) := \top$. If $\boldsymbol{b} = b$ (a single parameter), then we set $\mathrm{gdd}_X^\tau(\boldsymbol{b}) := (b = b)$ (an equality guard). Otherwise, there is some $X^+ : \exists\boldsymbol{x}.(\boldsymbol{U}(\boldsymbol{x}) \wedge \beta(\boldsymbol{x}\boldsymbol{b}')) \in \tau$ where β and $\boldsymbol{U}$ are X-catoms, $\boldsymbol{U} \subseteq \mathbb{U}$, $\boldsymbol{b}' \supseteq \boldsymbol{b}$, and $\boldsymbol{b}' \cap \boldsymbol{d}$ is X-crelativized, and we set $\mathrm{gdd}_X^\tau(\boldsymbol{b}) := \exists\boldsymbol{x}.(\boldsymbol{U}(\boldsymbol{x}) \wedge \beta(\boldsymbol{x}\boldsymbol{b}'))$. In this last case, when we write $\mathrm{gdd}_X^\tau(\boldsymbol{b}) \wedge \neg\psi$ we mean the formula $\exists\boldsymbol{x}.(\boldsymbol{U}(\boldsymbol{x}) \wedge \beta(\boldsymbol{x}\boldsymbol{b}') \wedge \neg\psi)$.

The proof proceeds by induction on i, the stage at which τ was eliminated. We intersperse the definitions of the interpolants and proofs of correctness.

Base case. Assume there is an X-mosaic $\tau(\boldsymbol{c}) \in \mathcal{N}_1$. Then τ has an internal inconsistency.

Assume one of the following conditions holds for $Y \in \{\mathrm{L}, \mathrm{R}\}$:

- $Y^+ : \alpha(\boldsymbol{b}), Y^- : \alpha(\boldsymbol{b}) \in \tau$;
- $Y^+ : \bigwedge_i \psi_i \in \tau$ (respectively, $Y^- : \bigvee_i \psi_i \in \tau$), but $Y^+ : \psi_i \notin \tau$ (respectively, $Y^- : \psi_i \notin \tau$) for some i;
- $Y^+ : \bigvee_i \psi_i \in \tau$ (respectively, $Y^- : \bigwedge_i \psi_i \in \tau$), but $Y^+ : \psi_i \notin \tau$ (respectively, $Y^- : \psi_i \notin \tau$) for all i;
- $Y^+ : \alpha(\boldsymbol{b}) \wedge \neg\psi(\boldsymbol{b}) \in \tau$, but $Y^+ : \psi(\boldsymbol{b}) \notin \tau$ or $Y^- : \psi(\boldsymbol{b}) \notin \tau$;
- $Y^+ : \alpha(\boldsymbol{b}) \in \tau$ and $Y^- : \alpha(\boldsymbol{b}) \wedge \neg\psi(\boldsymbol{b}) \in \tau$, but $Y^+ : \psi(\boldsymbol{b}) \notin \tau$.

These internal inconsistencies produce very simple interpolants. If $X = Y$, then $\theta_X^\tau := \bot$. If $X \neq Y$, then $\theta_X^\tau := \top$.

The more interesting cases are when the inconsistency comes from interaction between L and R formulas.

Assume $X^+ : \alpha(\boldsymbol{b}), \widetilde{X}^- : \alpha(\boldsymbol{b}) \in \tau$. Then we set $\theta^\tau_X := \alpha(\boldsymbol{b})$. This clearly satisfies (Imp), (Par), and (Rel).

Assume $X^- : \alpha(\boldsymbol{b}), \widetilde{X}^+ : \alpha(\boldsymbol{b}) \in \tau$ then $\theta^\tau_X := \mathrm{gdd}^\tau_X(\boldsymbol{b}) \wedge \neg\alpha(\boldsymbol{b})$. Although $\neg\alpha(\boldsymbol{b})$ would suffice for (Imp), we need to guard this formula with $\mathrm{gdd}^\tau_X(\boldsymbol{b})$ in order to ensure that we remain in GNF. Because $\boldsymbol{b}$ occurs in a $\widetilde{X}^+$-formula in the X-mosaic τ, $\mathrm{gdd}^\tau_X(\boldsymbol{b})$ exists with $X^+ : \mathrm{gdd}^\tau_X(\boldsymbol{b}) \in \tau$. Hence, (Imp) holds, and θ^τ_X is in GNF as required for (Rel). Property (Par) is immediate.

Assume $X^+ : \alpha(\boldsymbol{b})$ and $\widetilde{X}^- : \alpha(\boldsymbol{b}) \wedge \neg\psi(\boldsymbol{b}) \in \tau$, but $\widetilde{X}^+ : \psi(\boldsymbol{b}) \notin \tau$. Then we set $\theta^\tau_X := \alpha(\boldsymbol{b})$. (Par) and (Rel) clearly hold. For (Imp), it is clear that $\tau_X \models \theta^\tau_X$. To prove the other part of (Imp), we argue that $\tau_{\widetilde{X}} \models \neg\theta^\tau_X$. Suppose for the sake of contradiction that there is some model $\mathfrak{M}$ of $\tau_{\widetilde{X}} \wedge \alpha(\boldsymbol{b})$. Since $\widetilde{X}^- : \alpha(\boldsymbol{b}) \wedge \neg\psi(\boldsymbol{b}) \in \tau$, this means that $\mathfrak{M} \models \neg(\alpha(\boldsymbol{b}) \wedge \neg\psi(\boldsymbol{b})) \wedge \alpha(\boldsymbol{b})$, so $\mathfrak{M} \models \psi(\boldsymbol{b})$. By assumption, $\widetilde{X}^+ : \psi(\boldsymbol{b}) \notin \tau$. Notice $X^+ : \alpha(\boldsymbol{b}) \in \tau$ for α appearing positively in an X-formula and negatively in a $\widetilde{X}$-formula. Since occurrences are preserved in moving to mosaics (Proposition 27), this means that α is an X-cguard. Likewise, $\boldsymbol{b} \cap \boldsymbol{d}$ is X-crelativized in τ, since $\widetilde{X}^- : \alpha(\boldsymbol{b}) \wedge \neg\psi(\boldsymbol{b}) \in \tau$. Hence, $\boldsymbol{b}$ is X-cguarded in τ by $X^+ : \alpha(\boldsymbol{b})$. We can conclude that $\neg\psi(\boldsymbol{b})$ is a conjunct in $\tau_{\widetilde{X}}$, so $\mathfrak{M} \models \neg\psi(\boldsymbol{b})$, a contradiction.

Assuming $\widetilde{X}^+ : \alpha(\boldsymbol{b})$ and $X^- : \alpha(\boldsymbol{b}) \wedge \neg\psi(\boldsymbol{b}) \in \tau$, but $X^+ : \psi(\boldsymbol{b}) \notin \tau$, a similar argument shows that if we set $\theta^\tau_X := \mathrm{gdd}^\tau_X(\boldsymbol{b}) \wedge \neg\alpha(\boldsymbol{b})$ the desired properties are satisfied.

Assume $\widetilde{X}^- : \phi \in \tau$ for ϕ of the form $\exists \boldsymbol{x}.\psi(\boldsymbol{x}\boldsymbol{b}) \in \tau$. Assume there is some specialization $\boldsymbol{U}_0(\boldsymbol{c}') \wedge \chi_0(\boldsymbol{b}\boldsymbol{c}') \wedge \bigwedge_k \exists \boldsymbol{x}_k.\boldsymbol{U}_k(\boldsymbol{x}_k) \wedge \chi_k(\boldsymbol{x}_k\boldsymbol{b}\boldsymbol{c}')$ of ϕ such that $\boldsymbol{U}_0(\boldsymbol{c}') = U_1 c_1 \wedge \cdots \wedge U_j c_j$ and $Z^+_i : U_i c_i \in \tau$ for all $i \in \{1, \ldots, j\}$, but $\widetilde{X}^- : \chi_0(\boldsymbol{b}\boldsymbol{c}') \wedge \bigwedge_k \exists \boldsymbol{x}_k.\boldsymbol{U}_k(\boldsymbol{x}_k) \wedge \chi_k(\boldsymbol{x}_k\boldsymbol{b}\boldsymbol{c}') \notin \tau$. Set θ^τ_X to be the conjunction of $U_i c_i$ for all i such that such that $Z_i = X$ (the empty conjunction is $\top$). Note that (Par) is satisfied since for all i such that $Z_i = X$, $c_i \in \boldsymbol{c}$ and $X^+ : U_i c_i \in \tau$ witnesses the fact that c_i is X-crelativized in τ (this follows from Proposition 27 since U_i appears positively in an X-formula, and negatively in a $\widetilde{X}$-formula). (Rel) is trivially satisfied. For (Imp), it is clear that $\tau_X \models \theta^\tau_X$. Now suppose for the sake of contradiction that there is some model $\mathfrak{M}$ such that $\mathfrak{M} \models \theta^\tau_X \wedge \tau_{\widetilde{X}}$. Then $\mathfrak{M} \models U_i c_i$ for all i and $\mathfrak{M} \models \chi_0(\boldsymbol{b}\boldsymbol{c}') \wedge \bigwedge_k \exists \boldsymbol{x}_k.\boldsymbol{U}_k(\boldsymbol{x}_k) \wedge \chi_k(\boldsymbol{x}_k\boldsymbol{b}\boldsymbol{c}')$, so the specialization of ϕ holds in $\mathfrak{M}$. But it is straightforward to show that every specialization of ϕ implies ϕ (Lemma 26), a contradiction of $\mathfrak{M} \models \tau_{\widetilde{X}}$.

Assume $X^- : \phi \in \tau$ for ϕ of the form $\exists \boldsymbol{x}.\psi(\boldsymbol{x}\boldsymbol{b}) \in \tau$. Assume there is some specialization $X^- : \boldsymbol{U}_0(\boldsymbol{c}') \wedge \chi_0(\boldsymbol{b}\boldsymbol{c}') \wedge \bigwedge_k \exists \boldsymbol{x}_k.\boldsymbol{U}_k(\boldsymbol{x}_k) \wedge \chi_k(\boldsymbol{x}_k\boldsymbol{b}\boldsymbol{c}')$ of ϕ such that $\boldsymbol{U}_0(\boldsymbol{c}') = U_1 c_1 \wedge \cdots \wedge U_j c_j$ and $Z^+_i : U_i c_i \in \tau$ for all $i \in \{1, \ldots, j\}$, but $X^- : \chi_0(\boldsymbol{b}\boldsymbol{c}') \wedge \bigwedge_k \exists \boldsymbol{x}_k.\boldsymbol{U}_k(\boldsymbol{x}_k) \wedge \chi_k(\boldsymbol{x}_k\boldsymbol{b}\boldsymbol{c}') \notin \tau$. Set θ^τ_X to be the disjunction of $\neg U_i c_i$ for all i such that such that $Z_i = \widetilde{X}$ ($\perp$ for the empty disjunction). Note that (Par) is satisfied since for all i such that $Z_i = \widetilde{X}$, $c_i \in \boldsymbol{c}$ and $\widetilde{X}^+ : U_i c_i \in \tau$ in an X-mosaic implies that c_i is X-crelativized in τ if $c_i \in \boldsymbol{d}$. (Rel) is trivially satisfied. For (Imp), it is clear that $\tau_{\widetilde{X}} \models \neg\theta^\tau_X$. Now suppose for the sake of contradiction that there is some model $\mathfrak{M}$ such that $\mathfrak{M} \models \neg\theta^\tau_X \wedge \tau_X$. Then $\mathfrak{M} \models U_i c_i$ for all i and $\mathfrak{M} \models \chi_0(\boldsymbol{b}\boldsymbol{c}') \wedge \bigwedge_k \exists \boldsymbol{x}_k.\boldsymbol{U}_k(\boldsymbol{x}_k) \wedge \chi_k(\boldsymbol{x}_k\boldsymbol{b}\boldsymbol{c}')$, so the specialization of ϕ holds in $\mathfrak{M}$. But it is straightforward to show that every specialization of ϕ implies ϕ (Lemma 26), a contradiction of $\mathfrak{M} \models \tau_X$.

In all of these cases, (Occ) follows from Proposition 27.

Inductive case. Now assume there is an X-mosaic $\tau(\boldsymbol{c}) \in \mathcal{N}_i \setminus \mathcal{N}_{i-1}$ relativizing $\boldsymbol{d} \subseteq \boldsymbol{c}$. Then there is some $Y^+ : \eta = Y^+ : \exists \boldsymbol{x}.\psi(\boldsymbol{x}\boldsymbol{b}) \in \tau$, and if $\tau \to_{Y^+:\eta} \tau'$ then $\tau' \in \mathcal{N}_{i-1}$.

Let $\boldsymbol{r}$ be the parameters that are relativized in any such τ' (that is, $\boldsymbol{r}$ is $\boldsymbol{a} \cup (\boldsymbol{b} \cap \boldsymbol{d})$ if $Y = X$, and $\boldsymbol{r}$ is $\boldsymbol{a}$ if $Y = \widetilde{X}$).

We introduce some additional notation. For $Z \in \{\mathrm{L}, \mathrm{R}\}$, we define $\tau\restriction^{\boldsymbol{b};\boldsymbol{r}}_Z$ to be the Z-mosaic over parameters $\boldsymbol{b}$ and relativizing $\boldsymbol{r}$ that consists of the union of (i) the Z-formulas in τ that only use parameters from $\boldsymbol{b}$, and (ii) the $\widetilde{Z}$-formulas in τ whose parameters are contained in $\boldsymbol{b}$ and are Z-safe in the formulas added in step (i). For mosaics υ, υ' and $Z \in \{\mathrm{L}, \mathrm{R}\}$, we write $\upsilon'(\boldsymbol{a}\boldsymbol{b}) \supseteq_Z \upsilon$ if υ' uses only parameters from $\boldsymbol{a}\boldsymbol{b}$ and is obtained from υ be adding only Z-formulas (so the $\widetilde{Z}$-formulas in υ and υ' are identical). For a Y-mosaic υ', we write $\mathrm{rel}^{\upsilon'}_Y(\boldsymbol{a})$ for the conjunction of all Ua such that $U \in \mathbb{U}$ is an Y-catom and $Y^+ : Ua \in \upsilon'$ (so $Y^+ : Ua$ is an Y-crelativizer in υ').

Assume $Y = X$. Then let

$$\bigvee_{\substack{\upsilon'(\boldsymbol{a}\boldsymbol{b}) \supseteq_X \tau\restriction^{\boldsymbol{b};\boldsymbol{r}}_X \\ \text{s.t. } \tau \to_{X^+:\eta} \upsilon'}} \exists \boldsymbol{x}'.\left(\mathrm{rel}^{\upsilon'}_{\widetilde{X}}(\boldsymbol{a})[\boldsymbol{x}/\boldsymbol{a}] \wedge \bigwedge_{\substack{\tau'(\boldsymbol{a}\boldsymbol{b}) \supseteq_{\widetilde{X}} \upsilon' \\ \text{s.t. } \tau \to_{X^+:\eta} \tau'}} \theta^{\tau'}_X[\boldsymbol{x}/\boldsymbol{a}] \right)$$

where $\boldsymbol{x}' \subseteq \boldsymbol{x}$ is the subset of $\boldsymbol{x}$ appearing in $\mathrm{rel}^{\upsilon'}_X(\boldsymbol{a})[\boldsymbol{x}/\boldsymbol{a}]$.

We now begin the proof of correctness, starting with (Imp). Assume there is a model $\mathfrak{M}$ for τ_X. Since $X^+ : \eta \in \tau$, this means that $\mathfrak{M} \models \exists \boldsymbol{x}.\psi(\boldsymbol{x}\boldsymbol{b})$, so there are elements $\boldsymbol{a}' \in M$ and an expansion $\mathfrak{M}'$ of $\mathfrak{M}$ with the interpretation $\boldsymbol{a}^{\mathfrak{M}'} := \boldsymbol{a}'$ such that $\mathfrak{M}' \models \psi(\boldsymbol{a}\boldsymbol{b})$.

Take υ' to be the union of $\tau\restriction^{\boldsymbol{b};\boldsymbol{r}}_X$ and the set of formulas $X^p : \psi'$ in $\mathrm{cl}_X(\boldsymbol{a}\boldsymbol{b})$ such that $\mathfrak{M}' \models p\psi'$. It is straightforward to check that $\tau \to_{X^+:\eta} \upsilon'$. By construction, $\mathfrak{M}' \models \upsilon'_X$ and $\mathfrak{M}' \models \mathrm{rel}^{\upsilon'}_X(\boldsymbol{a})$.

Moreover, for any choice of $\tau'(\boldsymbol{a}\boldsymbol{b}) \supseteq_{\widetilde{X}} \upsilon'$, $\mathfrak{M}' \models \tau'_X$ (since the X-formulas are identical in υ' and τ'). By the inductive hypothesis, this means that $\mathfrak{M}' \models \theta^{\tau'}_X$. Hence, $\mathfrak{M}' \models \mathrm{rel}^{\upsilon'}_{\widetilde{X}}(\boldsymbol{a}) \wedge \bigwedge_{\substack{\tau'(\boldsymbol{a}\boldsymbol{b}) \supseteq_{\widetilde{X}} \upsilon' \\ \text{s.t. } \tau \to_{X^+:\eta} \tau'}} \theta^{\tau'}_X$ and $\mathfrak{M} \models \theta^\tau_X$.

Now assume that there is a model $\mathfrak{M}$ of θ^τ_X. Then there is some $\upsilon'(\boldsymbol{a}\boldsymbol{b}) \supseteq_X \tau\restriction^{\boldsymbol{b};\boldsymbol{r}}_X$ with $\tau \to_{X^+:\eta} \upsilon'$, elements $\boldsymbol{a}' \in M$, and an expansion $\mathfrak{M}'$ of $\mathfrak{M}$ with the interpretation $\boldsymbol{a}^{\mathfrak{M}'} := \boldsymbol{a}'$, such that $\mathfrak{M}' \models \mathrm{rel}^{\upsilon'}_{\widetilde{X}}(\boldsymbol{a}) \wedge \bigwedge_{\substack{\tau'(\boldsymbol{a}\boldsymbol{b}) \supseteq_{\widetilde{X}} \upsilon' \\ \text{s.t. } \tau \to_{X^+:\eta} \tau'}} \theta^{\tau'}_X$.

Note that for all $\tau'(\boldsymbol{ab}) \supseteq_{\widetilde{X}} \upsilon'$ with $\tau \rightarrow_{X+:\eta} \tau'$, $\mathfrak{M}' \models \theta_X^{\tau'}$ and consequently, the inductive hypothesis implies $\mathfrak{M}' \models \neg\tau'_{\widetilde{X}}$. In particular, consider $\tau' := \upsilon' \cup S' \supseteq_{\widetilde{X}} \upsilon'$ where S' is the set of formulas $\widetilde{X}^p : \psi'$ in $\mathrm{cl}_{\widetilde{X}}(\boldsymbol{ab})$ such that $\mathfrak{M}' \models p\psi'$, $\mathrm{par}(\psi') \cap \boldsymbol{a} \neq \emptyset$, and $\widetilde{X}^p : \psi'$ is X-safe in υ'. It is clear that $\tau'(\boldsymbol{ab}) \supseteq_{\widetilde{X}} \upsilon'$ and $\tau \rightarrow_{X+:\eta} \tau'$, so $\mathfrak{M}' \models \neg\tau'_{\widetilde{X}}$ as observed above.

Since $\mathfrak{M}' \models \neg\tau'_{\widetilde{X}}$, there is some conjunct χ in $\tau'_{\widetilde{X}}$ such that $\mathfrak{M}' \models \neg\chi$. Consider some conjunct $\chi(\boldsymbol{ab})$ in $\tau'_{\widetilde{X}}$ that actually uses some parameters from $\boldsymbol{a}$. Then by choice of S', $\mathfrak{M}' \models \chi(\boldsymbol{ab})$, so this formula cannot witness the fact that $\mathfrak{M}' \models \neg\tau'_{\widetilde{X}}$. This means there must be some conjunct $\chi(\boldsymbol{b})$ in $\tau'_{\widetilde{X}}$ that only uses parameters from $\boldsymbol{b}$ such that $\mathfrak{M}' \models \neg\chi(\boldsymbol{b})$. By definition of $\tau \rightarrow_{X+:\eta} \tau'$, $\chi(\boldsymbol{b})$ must also be a conjunct in $\tau_{\widetilde{X}}$. Hence, $\mathfrak{M} \models \neg\tau_{\widetilde{X}}$.

For (Occ), the desired occurrence of constants and the polarity of relations follows from the inductive hypothesis and the fact that occurrences are preserved in moving to mosaics (Proposition 27).

The last thing to prove is (Par) and (Rel). The inductively defined interpolants $\theta_X^{\tau'}$ only use parameters in τ'. Moreover, any parameters in $\boldsymbol{a} \cup (\boldsymbol{b} \cap \boldsymbol{d})$ are X-crelativized in τ'. By construction, τ' and υ' have the same X-formulas, so these X-crelativizers are also in υ'.

In particular, the X-crelativizers for parameters in $\boldsymbol{a}$ appear as conjuncts in $\mathrm{rel}_X^{\upsilon'}(\boldsymbol{a})$. This means that any parameters from $\boldsymbol{a}$ in $\theta_X^{\tau'}$ are removed using existential quantification, and this quantification is $\mathbb{U}$-relativized by $\mathrm{rel}_X^{\upsilon'}(\boldsymbol{a})$.

The remaining parameters in θ_X^{τ} are from $\boldsymbol{b}$, so $\mathrm{par}(\theta_X^{\tau}) \subseteq \boldsymbol{c}$. Moreover, any X-relativizers for $\boldsymbol{b} \cap \boldsymbol{d}$ in υ' must also be in τ (since $\tau \rightarrow_{X+:\eta} \upsilon'$ and parameters $\boldsymbol{b}$ occur in both τ and υ'). Hence, $\mathrm{par}(\theta_X^{\tau}) \cap \boldsymbol{d}$ is X-crelativized in τ.

Overall, this means that (Par) and (Rel) hold.

Assume $Y = \widetilde{X}$. Then let

$$\mathrm{gdd}_X^{\tau}(\boldsymbol{b}) \wedge \neg\left(\bigvee_{\substack{\upsilon'(\boldsymbol{ab}) \supseteq_{\widetilde{X}} \tau\restriction_{\widetilde{X}}^{\boldsymbol{b};\boldsymbol{r}} \\ \text{s.t. } \tau \rightarrow_{\widetilde{X}+:\eta} \upsilon'}} \exists \boldsymbol{x}'. \left(\mathrm{rel}_{\widetilde{X}}^{\upsilon'}(\boldsymbol{a})[\boldsymbol{x}/\boldsymbol{a}] \wedge \bigwedge_{\substack{\tau'(\boldsymbol{ab}) \supseteq_X \upsilon' \\ \text{s.t. } \tau \rightarrow_{\widetilde{X}+:\eta} \tau'}} \theta_{\widetilde{X}}^{\tau'}[\boldsymbol{x}/\boldsymbol{a}] \right)\right)$$

where $\boldsymbol{x}' \subseteq \boldsymbol{x}$ is the subset of $\boldsymbol{x}$ appearing in $\mathrm{rel}_{\widetilde{X}}^{\upsilon'}(\boldsymbol{a})[\boldsymbol{x}/\boldsymbol{a}]$.

We prove (Imp) by showing $\tau_{\widetilde{X}} \models \neg\theta_X^{\tau}$ and $\neg\theta_X^{\tau} \models \neg\tau_X$.

Assume there is a model $\mathfrak{M}$ for $\tau_{\widetilde{X}}$. Since $\widetilde{X}^+ : \eta \in \tau$, this means that $\mathfrak{M} \models \exists \boldsymbol{x}.\psi(\boldsymbol{xb})$, so there are elements $\boldsymbol{a}' \in M$ and an expansion $\mathfrak{M}'$ of $\mathfrak{M}$ with the interpretation $\boldsymbol{a}^{\mathfrak{M}'} := \boldsymbol{a}'$ such that $\mathfrak{M}' \models \psi(\boldsymbol{ab})$.

Take υ' to be the union of $\tau\restriction_{\widetilde{X}}^{\boldsymbol{b};\boldsymbol{r}}$ and the set of formulas $\widetilde{X}^p : \psi'$ in $\mathrm{cl}_{\widetilde{X}}(\boldsymbol{ab})$ such that $\mathfrak{M}' \models p\psi'$. It is straightforward to check that $\tau \rightarrow_{\widetilde{X}+:\eta} \upsilon'$. By construction, $\mathfrak{M}' \models \upsilon'_{\widetilde{X}}$ and $\mathfrak{M}' \models \mathrm{rel}_{\widetilde{X}}^{\upsilon'}(\boldsymbol{a})$.

Moreover, for any choice of $\tau'(\boldsymbol{ab}) \supseteq_X \upsilon'$, $\mathfrak{M}' \models \tau'_{\widetilde{X}}$ (since the $\widetilde{X}$-formulas are identical in υ' and τ'). By the inductive hypothesis, this means that $\mathfrak{M}' \models \theta_{\widetilde{X}}^{\tau'}$. Hence, $\mathfrak{M}' \models \mathrm{rel}_X^{\upsilon'}(\boldsymbol{a}) \wedge \bigwedge_{\substack{\tau'(\boldsymbol{ab}) \supseteq_X \upsilon' \\ \text{s.t. } \tau \rightarrow_{\widetilde{X}+:\eta} \tau'}} \theta_{\widetilde{X}}^{\tau'}$ so $\mathfrak{M} \models \neg\theta_X^{\tau}$.

Now assume that there is a model $\mathfrak{M}$ of $\neg\theta_X^{\tau}$.

If $\mathfrak{M} \models \neg\,\mathrm{gdd}_X^{\tau}(\boldsymbol{b})$, then it must be the case that $|\boldsymbol{b}| > 1$ (since $\neg\top$ and $\neg(b = b)$ cannot hold in $\mathfrak{M}$). Thus, $\mathfrak{M} \models \neg\tau_X$ since $X^+ : \mathrm{gdd}_X^{\tau}(\boldsymbol{b}) \in \tau$ (guaranteed since $\widetilde{X}^+ : \eta(\boldsymbol{b}) \in \tau$).

Otherwise, there is some $\upsilon'(\boldsymbol{ab}) \supseteq_{\widetilde{X}} \tau\restriction_{\widetilde{X}}^{\boldsymbol{b};\boldsymbol{r}}$ with $\tau \rightarrow_{\widetilde{X}+:\eta} \upsilon'$, elements $\boldsymbol{a}' \in M$, and an expansion $\mathfrak{M}'$ of $\mathfrak{M}$ with the interpretation $\boldsymbol{a}^{\mathfrak{M}'} := \boldsymbol{a}'$, such that $\mathfrak{M}' \models \mathrm{rel}_{\widetilde{X}}^{\upsilon'}(\boldsymbol{a}) \wedge \bigwedge_{\substack{\tau'(\boldsymbol{ab}) \supseteq_X \upsilon' \\ \text{s.t. } \tau \rightarrow_{\widetilde{X}+:\eta} \tau'}} \theta_{\widetilde{X}}^{\tau'}$.

Note that for all $\tau'(\boldsymbol{ab}) \supseteq_X \upsilon'$ with $\tau \rightarrow_{\widetilde{X}+:\eta} \tau'$, $\mathfrak{M}' \models \theta_{\widetilde{X}}^{\tau'}$ and consequently, the inductive hypothesis implies $\mathfrak{M}' \models \neg\tau'_X$. In particular, consider $\tau' := \upsilon' \cup S' \supseteq_X \upsilon'$ where S' is the set of formulas $X^p : \psi'$ in $\mathrm{cl}_X(\boldsymbol{ab})$ such that $\mathfrak{M}' \models p\psi'$, $\mathrm{par}(\psi') \cap \boldsymbol{a} \neq \emptyset$, and $X^p : \psi'$ is $\widetilde{X}$-safe in υ'. It is clear that $\tau'(\boldsymbol{ab}) \supseteq_X \upsilon'$ and $\tau \rightarrow_{\widetilde{X}+:\eta} \tau'$, so $\mathfrak{M}' \models \neg\tau'_X$ as observed above.

Since $\mathfrak{M}' \models \neg\tau'_X$, there is some conjunct χ in τ'_X such that $\mathfrak{M}' \models \neg\chi$. Consider some conjunct $\chi(\boldsymbol{ab})$ in τ'_X that actually uses some parameters from $\boldsymbol{a}$. Then by choice of S', $\mathfrak{M}' \models \chi(\boldsymbol{ab})$, so this formula cannot witness the fact that $\mathfrak{M}' \models \neg\tau'_X$. This means there must be some conjunct $\chi(\boldsymbol{b})$ in τ'_X that only uses parameters from $\boldsymbol{b}$ such that $\mathfrak{M}' \models \neg\chi(\boldsymbol{b})$. By definition of $\tau \rightarrow_{\widetilde{X}+:\eta} \tau'$, $\chi(\boldsymbol{b})$ must also be a conjunct in τ_X. Hence, $\mathfrak{M} \models \neg\tau_X$.

For (Occ), the desired occurrence of constants and the polarity of relations follows from the inductive hypothesis and the fact that occurrences are preserved in moving to mosaics (Proposition 27).

As an example, we give the proof if $X = \mathrm{R}$ and $\mathrm{gdd}_X^{\tau}(\boldsymbol{b})$ corresponds to an atomic formula. Let $(\alpha, p) \in \mathrm{occ}(\neg\theta_X^{\tau})$ (we abuse notation slightly and write this to mean that the relation S used in α satisfies $(S, p) \in \mathrm{occ}(\neg\theta_X^{\tau})$). One possibility is that (α, p) is $(\mathrm{gdd}_X^{\tau}(\boldsymbol{b}), -)$, since $\mathrm{gdd}_X^{\tau}(\boldsymbol{b})$ occurs negatively in $\neg\theta_X^{\tau}$. By definition of an X-catom for $X = \mathrm{R}$, $(\mathrm{gdd}_X^{\tau}(\boldsymbol{b}), -) \in \mathrm{occ}(\varphi_{\mathrm{L}}) \cap \mathrm{occ}(\varphi_{\mathrm{R}})$. Another possibility is that (α, p) is $(U, +)$ for a conjunct Ua in $\mathrm{rel}_{\widetilde{X}}^{\upsilon'}(\boldsymbol{a})$. By definition of a $\widetilde{X}$-crelativizer for $X = \mathrm{R}$, we must have $(U, +) \in \mathrm{occ}(\varphi_{\mathrm{L}}) \cap \mathrm{occ}(\varphi_{\mathrm{R}})$. Otherwise, (α, p) comes from $\mathrm{occ}(\theta_{\widetilde{X}}^{\tau'})$ for some τ' such that $\tau \rightarrow_{\widetilde{X}+:\eta} \tau'$. By the inductive hypothesis, this means $(\alpha, p) \in \mathrm{occ}(\varphi_{\mathrm{L}}) \cap \mathrm{occ}(\varphi_{\mathrm{R}})$. Overall, this means that $\mathrm{occ}(\neg\theta_X^{\tau}) \subseteq \mathrm{occ}(\varphi_{\mathrm{L}}) \cap \mathrm{occ}(\varphi_{\mathrm{R}})$ so (Occ) holds for θ_X^{τ}.

The last thing to prove is (Par) and (Rel). The inductively defined interpolants $\theta_{\widetilde{X}}^{\tau'}$ only use parameters in τ'. Moreover, any parameters in $\theta_{\widetilde{X}}^{\tau'}$ from $\boldsymbol{a}$ are $\widetilde{X}$-crelativized in τ'. By construction, τ' and υ' have the same $\widetilde{X}$-formulas, so these $\widetilde{X}$-crelativizers are also in υ'.

In particular, the $\widetilde{X}$-crelativizers for parameters in $\boldsymbol{a}$ appear as conjuncts in $\mathrm{rel}^{\upsilon'}_{\widetilde{X}}(\boldsymbol{a})$. This means that any parameters from $\boldsymbol{a}$ in $\theta^{\tau'}_X$ are removed using existential quantification, and this quantification is $\mathbb{U}$-relativized by $\mathrm{rel}^{\upsilon'}_{\widetilde{X}}(\boldsymbol{a})$.

The remaining parameters in the negated subformula of θ^τ_X are from $\boldsymbol{b}$, Because $\widetilde{X}^+ : \eta$ is a $\widetilde{X}^+$-formula using $\boldsymbol{b}$ in an X-mosaic, we know that any parameters in $\boldsymbol{b} \cap \boldsymbol{d}$ are X-crelativized in τ and $\boldsymbol{b}$ is X-cguarded in τ. This means that $\mathrm{gdd}^\tau_X(\boldsymbol{b})$ is defined and can guard the parameters $\boldsymbol{b}$ in the negated subformula.

Overall, this means (Par) and (Rel) hold.

Size of interpolants. Finally, we seek to bound the size of the shared DAG representation Θ_i of all θ^τ_X coming from $\tau \in \mathcal{N}_i$, such that $|\Theta_i| \leq i \cdot 2^{q(n)\cdot 2^{q(hw)}}$ for some polynomial function q independent of φ_{L} and φ_{R}.

Let $n = |\varphi_{\mathrm{L}}| + |\varphi_{\mathrm{R}}|$, $w = \mathrm{width}(\varphi_{\mathrm{L}} \wedge \neg\varphi_{\mathrm{R}})$, and $h = \mathrm{rank}_{\mathrm{CQ}}(\varphi_{\mathrm{L}} \wedge \neg\varphi_{\mathrm{R}})$. Let $f(n,h,w) = 2^{p(n)\cdot 2^{p(hw)}}$ be the bound on the total number of mosaics for φ_{L} and φ_{R} in $\mathcal{M}$ given by Proposition 35.

Let Θ_0 be the empty graph. We claim that at each stage $i > 0$ in the interpolant construction, the number of nodes and edges added to Θ_{i-1} in order to represent the new formulas added in stage i is polynomial in $f(n,h,w)$. The desired bound for the size of the DAG representations follows from this claim.

First consider $i = 1$. To build Θ_1, we start with the graph consisting of nodes labelled by atomic formulas (without provenance labels) from $\mathrm{cl}_{\mathrm{L}}(P) \cup \mathrm{cl}_{\mathrm{R}}(P)$, as well as nodes labelled with $\top$ and $\bot$. The number of nodes in this graph is at most $f(n,h,w)$. Constructing a formula θ^τ_X for $\tau \in \mathcal{N}_1$ results in a formula with at most $w+2$ non-atomic subformulas (for a formula of the form $\exists \boldsymbol{x}.\beta(\boldsymbol{xb}) \wedge \neg\alpha(\boldsymbol{b})$), so $w+2$ new nodes may be needed in the DAG representation in addition to the nodes for atomic formulas. Since $|\mathcal{N}_1| \leq f(n,h,w)$, this means the total size of Θ_1 is quadratic in $f(n,h,w)$.

Now consider $i > 1$. For each $\tau \in \mathcal{N}_i$, the formula θ^τ_X results in formulas with (at most) the following new subformulas:

- w subformulas from existentially quantified variables in gdd,
- three subformulas from the guarded negation,
- $f(n,h,w)$ subformulas from the disjunction over υ', each with w existential subformulas, w subformulas from rel, and $f(n,h,w)$ subformulas from the conjunction over τ'.

Since $|\mathcal{N}_i| \leq f(n,h,w)$, this means that the number of new nodes and edges added to the DAG representation of Θ_{i-1} in order to construct the DAG representation of Θ_i is polynomial (cubic) in $f(n,h,w)$. ☐

C.2 Constructive interpolation (Proof of Theorem 9 for equality-free formulas)

We now prove Theorem 9, with the additional restriction that the input formulas do not use equality.

Recall the statement:

> Let φ'_{L} and φ'_{R} be GNF (respectively, UNF) formulas without equality over signatures σ'_{L} and σ'_{R}, respectively. If $\varphi'_{\mathrm{L}} \models \varphi'_{\mathrm{R}}$ and $|\varphi'_{\mathrm{L}}| + |\varphi'_{\mathrm{R}}| = n$, then we can construct a DAG representation of a GNF (respectively, UNF) interpolant θ such that
>
> - $\varphi'_{\mathrm{L}} \models \theta$ and $\theta \models \varphi'_{\mathrm{R}}$;
> - $\mathrm{occ}(\theta) \subseteq \mathrm{occ}(\varphi'_{\mathrm{L}}) \cap \mathrm{occ}(\varphi'_{\mathrm{R}})$;
> - $\mathrm{free}(\theta) \subseteq \mathrm{free}(\varphi_{\mathrm{L}}) \cap \mathrm{free}(\varphi_{\mathrm{R}})$;
> - $\mathrm{con}(\theta) \subseteq \mathrm{con}(\varphi'_{\mathrm{L}}) \cup \mathrm{con}(\varphi'_{\mathrm{R}})$;
> - if φ'_{L} and φ'_{R} are $\mathbb{U}$-relativized, for $\mathbb{U}$ a distinguished set of unary relations from $\sigma'_{\mathrm{L}} \cup \sigma'_{\mathrm{R}}$, then θ is $\mathbb{U}$-relativized (when treating $\top$ and $\bot$ as atomic formulas);
> - the DAG representation of θ is of size at most $2^{2^{p(n)}}$ for some polynomial function p independent of φ'_{L} and φ'_{R} (and this can be improved to a size of at most $2^{p(n)}$ when φ'_{L} and φ'_{R} are in GF without equality and the bound on the arity of relations is fixed).

We consider various cases, depending on the form of the input formulas φ'_{L} and φ'_{R}. Let $n = |\varphi'_{\mathrm{L}}| + |\varphi'_{\mathrm{R}}|$.

We can always assume that $\mathrm{free}(\varphi_{\mathrm{L}}) = \mathrm{free}(\varphi_{\mathrm{R}})$ (because, we can existentially quantify out the free variables of φ_{L} that are not free variables of φ_{R}, and we can universally quantify out the free variables of φ_{R} that are not free variables of φ_{L}, introducing a fresh guard relation if necessary).

Relativized GNF input in weak GN-normal form. Assume φ'_{L} and φ'_{R} are in $\mathbb{U}$-relativized GNF formulas without equality, and are already in weak GN-normal form. Moreover, assume $\mathrm{free}(\varphi_{\mathrm{L}}) = \mathrm{free}(\varphi_{\mathrm{R}}) = \boldsymbol{z}$.

Let $\varphi_{\mathrm{L}} := \varphi'_{\mathrm{L}}$ and $\varphi_{\mathrm{R}} := \varphi'_{\mathrm{R}}$.

Use the mosaic elimination procedure over the set $\mathcal{M}$ of relativized GNF mosaics for $\varphi_{\mathrm{L}} \wedge \neg\varphi_{\mathrm{R}}$. Then construct mosaic interpolants for the eliminated mosaics $\mathcal{N}' \subseteq \mathcal{M}$ using Lemma 36, and let

$$\theta := \bigvee_{\tau(\boldsymbol{z}) \supseteq_{\mathrm{L}} \{\mathrm{L}^+ : \varphi_{\mathrm{L}}\}} \ \bigwedge_{\substack{\tau'(\boldsymbol{z}) \supseteq_{\mathrm{R}} \tau \\ \text{s.t. } \tau' \in \mathcal{N}'}} \theta^{\tau'}_{\mathrm{L}}$$

where $\mathcal{N}'$ is the set of eliminated mosaics and we write $\upsilon'(\boldsymbol{z}) \supseteq_Z \upsilon$ if υ' is over parameters $\boldsymbol{z}$ and relativizing $\emptyset$, and is obtained from υ be adding only Z-formulas.

We claim θ is a GNF interpolant for $\varphi'_{\mathrm{L}} \models \varphi'_{\mathrm{R}}$ satisfying the properties in Theorem 9.

We first prove $\varphi'_{\mathrm{L}} \models \theta \models \varphi'_{\mathrm{R}}$.

Assume there is a structure $\mathfrak{M}$ such that $\mathfrak{M} \models \varphi_{\mathrm{L}}$. Let $\tau(\boldsymbol{z}) \supseteq_{\mathrm{L}} \{\mathrm{L}^+ : \varphi_{\mathrm{L}}\}$ be the set of formulas from $\mathrm{cl}_{\mathrm{L}}(\boldsymbol{z})$ that are true in $\mathfrak{M}$. For any $\tau'(\boldsymbol{z}) \supseteq_{\mathrm{R}} \tau$, $\mathfrak{M} \models \tau'_{\mathrm{L}}$. By Lemma 36, this means that $\mathfrak{M} \models \theta^{\tau'}_{\mathrm{L}}$. Overall, $\mathfrak{M} \models \theta$.

Now assume there is a structure $\mathfrak{M}$ such that $\mathfrak{M} \models \theta$. Then there is some $\tau(\boldsymbol{z}) \supseteq_{\mathrm{L}} \{\mathrm{L}^+ : \varphi_{\mathrm{L}}\}$ such that $\mathfrak{M} \models \bigwedge_{\substack{\tau'(\boldsymbol{z}) \supseteq_{\mathrm{R}} \tau \\ \text{s.t. } \tau' \in \mathcal{N}'}} \theta^{\tau'}_{\mathrm{L}}$. Let S' be the set of formulas $\mathrm{R}^p : \psi$ in $\mathrm{cl}_{\mathrm{R}}(\boldsymbol{z})$ such that $\mathfrak{M} \models p\psi$, and $\mathrm{R}^p : \psi$ is L-safe in τ. Consider the L-mosaic over $\boldsymbol{z}$ (relativizing $\emptyset$) such that

$\tau' := \tau \cup S' \cup \{\mathrm{R}^- : \varphi_\mathrm{R}\}$. By completeness of the mosaic system (Proposition 34, which is part of Theorem 28), every mosaic containing $\mathrm{L}^+ : \varphi_\mathrm{L}$ and $\mathrm{R}^- : \varphi_\mathrm{R}$ must be in $\mathcal{N}'$. Hence, $\tau'(\boldsymbol{z}) \supseteq_\mathrm{R} \tau$ and $\tau' \in \mathcal{N}'$, so $\mathfrak{M} \models \theta_\mathrm{L}^{\tau'}$. By Lemma 36, this implies that $\mathfrak{M} \models \neg\tau'_\mathrm{R}$. But every formula $\mathrm{R}^p : \psi$ in S' was chosen such that $\mathfrak{M} \models p\psi$. The only way that $\mathfrak{M} \models \neg\tau'_\mathrm{R}$ is if $\mathfrak{M} \models \varphi_\mathrm{R}$.

Since each $\theta_\mathrm{L}^{\tau'}$ is in $\mathbb{U}$-relativized GNF by Lemma 36, this formula is in $\mathbb{U}$-relativized GNF, with $\top$ and $\bot$ Replacing $\top$ and $\bot$ by $\exists x.x = x$ and $\exists x.(x = x \wedge \neg(x = x))$, respectively, keeps the formula in GNF, but the resulting formula is technically no longer $\mathbb{U}$-relativized. Note that it is not necessarily in GN-normal form or weak GN-normal form.

The other properties about the polarity of relations and constants follow from Lemma 36.

Let if $n = |\varphi_\mathrm{L}| + |\varphi_\mathrm{R}|$, $w = \mathrm{width}(\varphi_\mathrm{L} \wedge \neg\varphi_\mathrm{R})$, and $h = \mathrm{rank}_\mathrm{CQ}(\varphi_\mathrm{L} \wedge \neg\varphi_\mathrm{R})$. Since the combined size of the DAG representation of all $\theta_\mathrm{L}^{\tau'}$ used in θ is at most $2^{q(n)\cdot 2^{q(hw)}}$ (for polynomial function q given by Lemma 36 that are independent of the input formulas), and there are at most $\left(2^{q(n)\cdot 2^{q(hw)}}\right)^2$ new nodes that need to be added to this DAG representation to construct θ, the overall size of the DAG representation of θ is at most $2^{r(n)\cdot 2^{r(hw)}}$ for some polynomial functions r independent of the input formulas. Since $w, h \le n$, this means that overall, the size of this DAG representation for θ is at most doubly exponential in n.

Relativized GNF input. Convert the formulas above to weak GN-normal form using Lemma 23, and then argue as in the previous case for these $\mathbb{U} \cup \{U\}$-relativized formulas.

Because the conversion to weak GN-normal form keeps the width and CQ-rank bounded by n and the overall size exponential in n, the size of the DAG representation of the interpolant is still at most doubly exponential in n.

GNF input with ordinary (unrelativized) quantifiers. Assume φ'_L and φ'_R are in GNF and have ordinary (unrelativized) quantifiers. We introduce a new unary relation U that is not in σ_L or σ_R. Let φ_X^U be the result of relativizing all quantifiers in φ'_X to U. Consider the validity

$$\bigwedge_{z \in \boldsymbol{z}} Uz \wedge \bigwedge_{e \in \boldsymbol{e}} Ue \wedge \varphi_\mathrm{L}^U \models \varphi_\mathrm{R}^U$$

where $\boldsymbol{z} := \mathrm{par}(\varphi'_\mathrm{L}) = \mathrm{par}(\varphi'_\mathrm{R})$ and $\boldsymbol{e} := \mathrm{con}(\varphi'_\mathrm{L}) \cup \mathrm{con}(\varphi'_\mathrm{R})$. Then apply the interpolation result for this relativized GNF input, which yields a $\mathbb{U}$-relativized interpolant θ^U for $\mathbb{U} = \{U\}$. Replacing all occurrences of Ux in θ^U with $\top$ results in a GNF interpolant θ for the original validity.

UNF input. Assume φ'_L and φ'_R are in UNF (and for simplicity that they are already in weak GN-normal form). Observe that in the interpolation construction in Lemma 36, the only guards for negation that are used for an interpolant θ_X^τ are equality guards of the form $x = x$ or X-cguards in τ. By definition, X-cguards must have appeared as guards in one of the original formulas φ'_L and φ'_R. Since the only guards in the original formulas were unary guards, this means that θ_X^τ uses only unary negation.

For UNF input that is not in weak GN-normal form, the conversion into weak GN-normal form using Lemma 23 keeps the formula in UNF and does not change the guards, so the same argument applies.

GF input. Assume φ'_L and φ'_R are in GF. We can convert these formulas to weak GN-normal form formulas φ_L and φ_R as described in Lemma 24, and then argue as in the previous cases. Because the conversion to weak GN-normal form keeps the width the same, CQ-rank 1, and size polynomial in n, this means that the size of the DAG representation for the interpolant is doubly exponential in n in general, but singly exponential in n when the width is considered fixed.

D. Equality

In this section, we describe a way to extend our interpolation results to handle input formulas with equality.

The first thing we show is that any GNF formula with equality can be converted to a form with a very limited use of equality. Let us say that a GNF formula ϕ is *equality-normalized* if

(i) every occurrence of a relational atomic formula $R\,\boldsymbol{t}$ in ϕ appears in conjunction with

$$\text{all-distinct}(\boldsymbol{t}) := \bigwedge_{t \in \boldsymbol{t}} (t = t) \wedge \bigwedge_{t,t' \in \boldsymbol{t}, t \neq t'} \neg(t = t'),$$

(ii) whenever equalities are used as guards for negations, then these equality guards are of the form $x = x$, and

(iii) there are no occurrences of equality in ϕ other than those described in (i) and (ii).

Let $\phi(\boldsymbol{x})$ be a GNF formula containing constants $\boldsymbol{e}$, and let $\equiv$ be any equivalence relation over $\{\boldsymbol{x}, \boldsymbol{e}\}$. We denote by $\xi_\equiv$ the first-order formula

$$\bigwedge_{s \equiv t} (s = t) \wedge \bigwedge_{s \not\equiv t} \neg(s = t),$$

that is, the conjunction of all equalities and inequalities corresponding to $\equiv$. Note that this formula does not necessarily belong to GNF as it may contain unguarded inequalities.

Lemma 37. *Let $\phi(\boldsymbol{x})$ be a* GNF *formula in* GN*-normal form containing constants* $\boldsymbol{e}$*, and let $\equiv$ be any equivalence relation over $\{\boldsymbol{x}, \boldsymbol{e}\}$. We can construct an equality-normalized $\phi'_\equiv(\boldsymbol{x})$ in* GN*-normal form such that*

- $\models \forall \boldsymbol{x}\big(\xi_\equiv(\boldsymbol{x}) \to (\phi(\boldsymbol{x}) \leftrightarrow \phi'_\equiv(\boldsymbol{x}))\big)$.
- $|\phi'_\equiv| \le p(|\phi|^k)$*, where k is the quantifier rank of ϕ and p is a polynomial function independent of ϕ.*
- $\mathrm{width}(\phi'_\equiv) \le \mathrm{width}(\phi)$*;*
- $\mathrm{rank}_\mathrm{CQ}(\phi'_\equiv) \le \mathrm{rank}_\mathrm{CQ}(\phi) \cdot (\mathrm{width}(\phi) + |\{\boldsymbol{e}\}|)^2$.

Furthermore, $\mathrm{occ}(\phi'_\equiv) \subseteq \mathrm{occ}(\phi)$*, and, if ϕ is $\mathbb{U}$-relativized, then $\phi'_\equiv$ is as well.*

Here, by *quantifier rank* we mean the maximal nesting depth of quantifiers in the formula.

Proof. We may assume that every equality guard in ϕ is of the form $x = x$ (if an equality guard is of the form $x = y$, we can freely replace all occurrences of y in the subformula by x). This takes care of condition (ii) in the definition of equality-normalized formulas.

For each $\equiv$-equivalence class, we fix an arbitrary representative — a constant whenever possible. We replace all occurrences of each variables and constant in $\boldsymbol{x}, \boldsymbol{e}$ by the representative of its equivalence class. Next, we replace subformulas of the form $s = t$, with $s, t \in \{\boldsymbol{x}, \boldsymbol{e}\}$, by $\top$ if $s = t$ and by $\bot$ otherwise; finally, we conjoin every relational atom containing distinct $s, t \in \{\boldsymbol{x}, \boldsymbol{e}\}$ with $\neg(s = t)$.

At this point, the conditions (i) and (iii) in the definition of equality-normalized formulas are satisfied *for equalities that have a free variables or constants on both sides*. It remains only to take care of the equalities that involve a quantified variable.

Consider any subformula of the form

$$\psi(\boldsymbol{z}) = \exists \boldsymbol{y}.\chi(\boldsymbol{y}, \boldsymbol{z})$$

We will essentially do a case distinction, for each quantified variable $y_i \in \boldsymbol{y}$, of the possible values that y_i may take. More precisely, we replace $\psi(\boldsymbol{z})$ by the disjunction, for each map $f : \{\boldsymbol{y}\} \to \{\boldsymbol{y}, \boldsymbol{z}, \boldsymbol{e}\}$, of the formula ψ_f obtained from ψ by (i) replacing each $y_i \in \boldsymbol{y}$ by $f(y_i)$, (ii) replacing $y_i = y_j$ by $\top$ if $f(y_i) = f(y_j)$ and by $\bot$ otherwise, (iii) replacing $y_i = t$ or or $t = y_i$ for $t \in \{\boldsymbol{z}, \boldsymbol{e}\}$ by $\top$ if $f(y_i) = t$ and by $\bot$ otherwise (and dropping the quantifiers corresponding to quantified variables that no longer occur in the formula). Finally, we conjoin every atom $R(\boldsymbol{t})$ with $\bigwedge_{t \in \boldsymbol{t}}(t = t) \wedge \bigwedge_{t,t' \in \boldsymbol{t}, t \neq t'} \neg(t = t')$. This clearly preserves the semantics of the formula over structures satisfying $\xi_\equiv$, and after performing this rewriting on each subformula $\psi(\boldsymbol{z})$ of the above form we obtain the desired equality-normalized $\phi'_\equiv$ in GN-normal form. There is an blowup involved in this procedure, but this blowup is at most exponential in the quantifier rank of ϕ (as can be seen by performing the rewriting step in a bottom-up fashion starting with the innermost quantifiers). Furthermore, even though we introduce disjunctions, the resulting formula is easily seen to be still in GN-normal form. The bounds on the width and CQ-rank of $\phi'_\equiv$ are immediate from the construction (note that the number of distinct terms occurring in any atomic formula in ϕ is at most $\mathrm{width}(\phi) + |\{\boldsymbol{e}\}|$). ☐

Combining Lemma 37 with Lemma 23 (and using also the fact that the formula transformation in the proof of Lemma 23 preserves quantifier rank) we obtain:

Proposition 38. *Let $\phi(\boldsymbol{x})$ be a* GNF *formula (not necessarily in* GN*-normal form) containing constants* $\boldsymbol{e}$*, and let $\equiv$ be any equivalence relation over $\{\boldsymbol{x}, \boldsymbol{e}\}$. We can construct an equality-normalized $\phi'_\equiv(\boldsymbol{x})$ in* GN*-normal form such that*

- $\models \forall \boldsymbol{x}\big(\xi_\equiv(\boldsymbol{x}) \to (\phi(\boldsymbol{x}) \leftrightarrow \phi'_\equiv(\boldsymbol{x}))\big)$.
- $|\phi'_\equiv|$ *is exponential in* $|\phi|$*;*
- $\mathrm{width}(\phi'_\equiv) \leq |\phi|$*;*
- $\mathrm{rank}_{\mathrm{CQ}}(\phi'_\equiv) \leq p(|\phi|)$ *for a polynomial function p independent of ϕ*

Furthermore, $\mathrm{occ}(\phi'_\equiv) \subseteq \mathrm{occ}(\phi)$*, and, if ϕ is $\mathbb{U}$-relativized, then $\phi'_\equiv$ is as well.*

Let φ_{L} and φ_{R} be GNF formulas with equality, and let $\boldsymbol{x}$ be the free variables and $\boldsymbol{e}$ the constants occurring in these formulas. Applying the above lemma, we obtain that $\varphi_{\mathrm{L}} \models \varphi_{\mathrm{R}}$ holds if and only if, for each equivalence relation $\equiv$ on $\{\boldsymbol{x}, \boldsymbol{e}\}$,

$$\xi_\equiv \wedge (\varphi_{\mathrm{L}})'_\equiv \models (\varphi_{\mathrm{R}})'_\equiv. \tag{2}$$

Moreover, a GNF-interpolant for $\varphi_{\mathrm{L}} \models \varphi_{\mathrm{R}}$ can be constructed by taking the disjunction of GNF-interpolants for (2) for all equivalence relations $\equiv$ (but recall that $\xi_\equiv$ may not be a GNF-formula). Towards constructing the latter, we can simplify the entailment (2) even further:

Lemma 39. *The entailment $\xi_\equiv \wedge (\varphi_{\mathrm{L}})'_\equiv \models (\varphi_{\mathrm{R}})'_\equiv$, is valid if and only if $(\varphi_{\mathrm{L}})'_\equiv \models (\varphi_{\mathrm{R}})'_\equiv$, is valid.*

Proof. One direction is trivial. For the other direction, we proceed by contraposition: let $\mathfrak{M} \models (\varphi_{\mathrm{L}})'_\equiv \wedge \neg(\varphi_{\mathrm{R}})'_\equiv$. Recall that $(\varphi_{\mathrm{L}})'_\equiv$ and $(\varphi_{\mathrm{R}})'_\equiv$ contain only the representative of each $\equiv$-equivalence class. We can freely expand $\mathfrak{M}$ so that all other variables and constants denote the same element that is denoted by the representative of their equivalence class. The only potential source of difficulty is that $\mathfrak{M}$ may interpret two constants or free variables, s and t, as the same element a, even though they are representatives of different $\equiv$-equivalence classes. However, in this case, we can simply "pull s and t apart", that is, we can replace a by two copies, a_1 and a_2, such that a_1 and a_2 participate in all the same facts that a did. This modification of the structure does not affect the truth of $(\varphi_{\mathrm{L}})'_\equiv$ and the falsity of $(\varphi_{\mathrm{R}})'_\equiv$, due to the fact that these formulas are equality normalized (as can be shown by a straightforward formula induction). ☐

We have thus reduced the problem at hand to obtaining interpolants for entailments of the form $(\varphi_{\mathrm{L}})'_\equiv \models (\varphi_{\mathrm{R}})'_\equiv$, where both formulas involved are equality-normalized. We can do this by applying the mosaic method, as in the previous section, treating equality as an ordinary relation symbol. We say that an X-mosaic (with $X \in \{\mathrm{L}, \mathrm{R}\}$) is *equality-trivial* if (i) $Y^+ : t = t' \in \tau$ only if t and t' are the same term (parameter or constant), and (ii) $Y^- : t = t'$ only if t and t' are distinct terms (parameters or constants). A close inspection of the soundness and completeness proofs for the mosaic method show that, for equality-normalized input formulas ψ_{L} and ψ_{R}, we can restrict attention to equality-trivial mosaics. Furthermore, in the completeness proof, the model that is constructed is such that distinct terms are realized by distinct elements of the model. Consequently, whenever an equality-trivial mosaic contains an equality or inequality, then this equality or inequality is guaranteed to be realized by every corresponding pair of elements of the model. We can conclude that, for equality-normalized input sentences, the mosaic method constitutes a sound and complete decision procedure.

Let $\mathcal{M}$ be the set of such equality-normalized GNF mosaics for $\psi_{\mathrm{L}} \wedge \neg\psi_{\mathrm{R}}$ over some set of parameters of size twice the width of $\psi_{\mathrm{L}} \wedge \neg\psi_{\mathrm{R}}$.

Theorem 40. *For equality-normalized* GNF*-formulas ψ_{L} and ψ_{R}, $\psi_{\mathrm{L}} \models \psi_{\mathrm{R}}$ holds if and only if there is a saturated set of internally consistent equality-trivial* GNF *mosaics from $\mathcal{M}$ that contains a mosaic τ such that $\mathrm{L}^+ : \psi_{\mathrm{L}}, \mathrm{R}^- : \psi_{\mathrm{R}} \in \tau$.*

Proposition 38 ensures that the complexity for deciding satisfiability using these mosaics remains in 2EXPTIME, and the DAG size of the resulting interpolants is still at most doubly exponential in the size of the original input. This is the case, even when we consider bounded arity GF input.

Interpolation follows from this as before: we prove a mosaic interpolation lemma like Lemma 36, but where the construction now ranges only over equality-trivial mosaics. For the purposes of the construction, equality is treated like any other relation.

We remark that, using this method, we can make no claims about the polarity of occurrences of equality in the interpolant – and this is why no such claims are made in our interpolation result, Theorem 9. The conversion to equality-normalized GN-normal form may change the polarity of occurrences of equality and equality may be introduced as unary guards in both polarities in the interpolant. This implies that the polarity of occurrences of equality in the interpolant are unrelated to the polarity of equality in the original formulas (in contrast to the other relations in the common signature, which can only appear in the interpolant in polarity p if they occurred in both φ_{L} and φ_{R} in polarity p). Indeed, even if the original formulas did not use equality, equality may appear in both polarities in the resulting interpolant due to the introduction of guards for unary negation.

E. Lower Bounds

In this section, we provide details for the lower bound argument described in Section 6.

E.1 Defining indexed AND/OR trees in UNF

Recall from Section 6 that we are interested in defining in UNF indexed AND/OR trees of depth 2^{2^n} where each node x in the AND/OR tree is the root of an index tree, with the index equal to the depth of x in the AND/OR tree.

Let σ_n consist of

- unary predicates IsOr, $\mathrm{InputOne}$, $\mathrm{AOValOne}$, $\mathrm{IndValOne}$, $\mathrm{IndDepth}_0, \ldots, \mathrm{IndDepth}_n$,
- binary predicate $\mathrm{AOChild}$, and
- ternary predicate $\mathrm{IndChild}$.

Signature σ'_n is the same as σ_n, except predicate $\mathrm{AOValOne}$ is replaced by a different unary predicate $\mathrm{AOValOne}'$. We write ρ_n for the signature $\sigma_n \cap \sigma'_n$.

The relations $\mathrm{IndValOne}, \mathrm{IndDepth}_0, \ldots, \mathrm{IndDepth}_n, \mathrm{IndChild}$ are used to describe index trees. We say a set of nodes is an *index tree* if the nodes form a tree (considering the relation $\mathrm{IndChild}$, and viewing $\mathrm{IndChild}\, xx_{\mathbf{l}}x_{\mathbf{r}}$ as asserting that $x_{\mathbf{l}}$ and $x_{\mathbf{r}}$ are, respectively, left and right children of x), the labels $\mathrm{IndDepth}_i$ correctly identify the depth i of each node in the tree (with the root satisfying $\mathrm{IndDepth}_0$), and any two nodes that have a corresponding position in the tree (based on the left/right branching used to reach it) agree on the relation $\mathrm{IndValOne}$. The index of such a tree is the value when viewing the leaves as 2^n bits of a binary number (where the order of the bits is determined by the branching, and only one bit from each branching type is used).

The other relations are used for the AND/OR tree. $\mathrm{InputOne}$ is viewed as a labelling of the AND/OR tree with either 0 or 1 (we will only be interested in these labels at the leaves). If $\mathrm{IsOr}\, x$ holds (respectively, $\neg\mathrm{IsOr}\, x$ holds), then we view this as a labelling of x with OR (respectively, AND). $\mathrm{AOValOne}$ and $\mathrm{AOValOne}'$ will be used to describe the intermediate calculations of the value of the AND/OR tree.

We will be interested in structures where the nodes in the AND/OR tree are roots of index trees describing the depth.

We first define the following auxiliary UNF formulas (technically, these should all be indexed by n). We will use the convention that built-in predicates like $\mathrm{IndChild}$ will be written $\mathrm{IndChild}\, xx_lx_r$, whereas our derived auxiliary formulas like $\mathrm{IndChild}_d$ that have free variables x and x_d will be written $\mathrm{IndChild}_d(x, x_d)$.

$\mathrm{IndChild}_d(x, x_d)$**:** x_d is the d-child of x for $d \in \{\mathbf{l}, \mathbf{r}\}$.

$$\begin{aligned} &\exists x_{\mathbf{r}}.\mathrm{IndChild}\, xx_dx_{\mathbf{r}} \quad \text{if } d = \mathbf{l} \\ &\exists x_{\mathbf{l}}.\mathrm{IndChild}\, xx_{\mathbf{l}}x_d \quad \text{if } d = \mathbf{r} \end{aligned}$$

$\mathrm{Link}(x_n, x'_n)$**:** x_n and x'_n are leaves in corresponding positions in the same index tree.

$$\exists x_0 \cdots x_{n-1}x'_1 \cdots x'_{n-1}.\Bigg(\mathrm{IndDepth}_0\, x_0 \wedge (\mathrm{IndChild}_{\mathbf{l}}(x_0, x_1) \wedge \mathrm{IndChild}_{\mathbf{l}}(x_0, x'_1)) \vee (\mathrm{IndChild}_{\mathbf{r}}(x_0, x_1) \wedge \mathrm{IndChild}_{\mathbf{r}}(x_0, x'_1)) \wedge \bigwedge_{i=1}^{n-1}\Big((\mathrm{IndChild}_{\mathbf{l}}(x_i, x_{i+1}) \wedge \mathrm{IndChild}_{\mathbf{l}}(x'_i, x'_{i+1})) \vee (\mathrm{IndChild}_{\mathbf{r}}(x_i, x_{i+1}) \wedge \mathrm{IndChild}_{\mathbf{r}}(x'_i, x'_{i+1}))\Big)\Bigg)$$

$\mathrm{LinkNext}(x_n, y_n)$**:** x_n and y_n are leaves in corresponding positions in neighboring index trees.

$$\exists x_0 \cdots x_{n-1}y_0 \cdots y_{n-1}.\Bigg(\mathrm{AOChild}\, x_0y_0 \wedge \mathrm{IndDepth}_0\, x_0 \wedge \mathrm{IndDepth}_0\, y_0 \quad \wedge \bigwedge_{i=0}^{n-1}\Big((\mathrm{IndChild}_{\mathbf{l}}(x_i, x_{i+1}) \wedge \mathrm{IndChild}_{\mathbf{l}}(y_i, y_{i+1})) \vee (\mathrm{IndChild}_{\mathbf{r}}(x_i, x_{i+1}) \wedge \mathrm{IndChild}_{\mathbf{r}}(y_i, y_{i+1}))\Big)\Bigg)$$

$\text{Before}(x_n, x'_n)$**:** x_n and x'_n are leaves in the same index tree, with x_n to the left of x'_n.

$$\exists x_0 \cdots x_{n-1} x'_1 \cdots x'_{n-1}. \bigvee_{j=0}^{n-1} \Bigg(\text{IndDepth}_0 x_0 \wedge \text{IndChild}_{\mathbf{l}}(x_j, x_{j+1}) \wedge \text{IndChild}_{\mathbf{r}}(x_j, x'_{j+1}) \quad \wedge$$
$$\bigwedge_{i=0}^{j-1} \Big((\text{IndChild}_{\mathbf{l}}(x_i, x_{i+1}) \vee \text{IndChild}_{\mathbf{r}}(x_i, x_{i+1})) \Big) \wedge$$
$$\bigwedge_{i=j+1}^{n-1} \Big((\text{IndChild}_{\mathbf{l}}(x_i, x_{i+1}) \vee \text{IndChild}_{\mathbf{r}}(x_i, x_{i+1})) \wedge (\text{IndChild}_{\mathbf{l}}(x'_i, x'_{i+1}) \vee \text{IndChild}_{\mathbf{r}}(x'_i, x'_{i+1})) \Big) \Bigg)$$

IndexedAOTrees**:** Every node satisfies exactly one IndDepth_i, AOChild only connects IndDepth_0-nodes, and every IndDepth_0-node is the root of an index tree. Take the conjunction of the following sentences.

$$\neg\exists x. \left(\bigwedge_{i=0}^{n} \Big(\neg\text{IndDepth}_i\, x \Big) \quad \vee \bigvee_{\substack{i,j \in \{0,\ldots,n\} \\ i \neq j}} \Big(\text{IndDepth}_i\, x \wedge \text{IndDepth}_j\, x \Big) \right)$$

$$\neg\exists xy. \Big(\text{AOChild}\, xy \wedge (\neg\text{IndDepth}_0\, x \vee \neg\text{IndDepth}_0\, y) \Big)$$

$$\neg\exists x x_{\mathbf{l}} x_{\mathbf{r}}. \Big(\text{IndDepth}_n\, x \wedge \text{IndChild}\, x x_{\mathbf{l}} x_{\mathbf{r}} \Big)$$

$$\neg\exists x x_{\mathbf{l}} x_{\mathbf{r}}. \bigvee_{i=0}^{n-1} \Big(\text{IndDepth}_i\, x \wedge \text{IndChild}\, x x_{\mathbf{l}} x_{\mathbf{r}} \wedge (\neg\text{IndDepth}_{i+1}\, x_{\mathbf{l}} \vee \neg\text{IndDepth}_{i+1}\, x_{\mathbf{r}}) \Big)$$

$$\neg\exists x. \bigvee_{i=0}^{n-1} \Big(\text{IndDepth}_i\, x \wedge \neg\exists x_{\mathbf{l}} x_{\mathbf{r}}. \text{IndChild}\, x x_{\mathbf{l}} x_{\mathbf{r}} \Big)$$

$$\neg\exists x_n x'_n. \big(\text{Link}(x_n, x'_n) \wedge ((\text{IndValOne}\, x_n \wedge \neg\text{IndValOne}\, x'_n) \vee (\neg\text{IndValOne}\, x_n \wedge \text{IndValOne}\, x'_n)) \big)$$

$\text{IndMin}(x_0)$**:** The index tree rooted at x_0 is the minimum (all 0's).

$$\text{IndDepth}_0\, x_0 \wedge \neg\exists x_1 \cdots x_n. \left(\bigvee_{i=0}^{n-1} \Big(\text{IndChild}_{\mathbf{l}}(x_i, x_{i+1}) \vee \text{IndChild}_{\mathbf{r}}(x_i, x_{i+1}) \Big) \wedge \text{IndValOne}\, x_n \right)$$

$\text{IndMax}(x_0)$**:** The index tree rooted at x_0 is the maximum (all 1's).

$$\text{IndDepth}_0\, x_0 \wedge \neg\exists x_1 \cdots x_n. \left(\bigvee_{i=0}^{n-1} \Big(\text{IndChild}_{\mathbf{l}}(x_i, x_{i+1}) \vee \text{IndChild}_{\mathbf{r}}(x_i, x_{i+1}) \Big) \wedge \neg\text{IndValOne}\, x_n \right)$$

$\text{AORoot}(x_0)$**:** x_0 is the root of an index tree with no incoming AOChild-edges.

$$\text{IndDepth}_0\, x_0 \wedge \neg\exists x'. \text{AOChild}\, x' x_0$$

$\text{AOLeaf}(x_0)$**:** x_0 is the root of an index tree with no outgoing AOChild-edges.

$$\text{IndDepth}_0\, x_0 \wedge \neg\exists x'. \text{AOChild}\, x_0 x'$$

IndexIsDepth**:** Every root of an index tree with no incoming AOChild-edges has minimal index, and for every index tree with non-maximal value l, every AOChild-neighboring tree has index $l + 1$.

$$\neg\exists x. (\text{AORoot}(x) \wedge \neg\text{IndMin}(x)) \quad \wedge$$
$$\neg\exists x. \text{IndDepth}_n\, x \wedge \neg\text{IndValOne}\, x \wedge \neg\exists x'. \big(\text{Before}(x, x') \wedge \neg\text{IndValOne}\, x' \big) \quad \wedge$$
$$\left(\begin{array}{l} \exists y. \big(\text{LinkNext}(x, y) \wedge \neg\text{IndValOne}\, y \big) \quad \vee \\ \exists x' y'. \big(\text{LinkNext}(x', y') \wedge \text{Before}(x, x') \wedge (\neg\text{IndValOne}\, x' \vee \text{IndValOne}\, y') \big) \quad \vee \\ \exists x' y'. \big(\text{LinkNext}(x', y') \wedge \text{Before}(x', x) \wedge ((\text{IndValOne}\, x' \wedge \neg\text{IndValOne}\, y') \vee (\neg\text{IndValOne}\, x' \wedge \text{IndValOne}\, y')) \big) \end{array} \right)$$

$\text{ConsistentAOValues}$**:** For all leaves in an AND/OR tree, AOValOne matches InputOne, and for all other nodes in an AND/OR tree, AOValOne matches the evaluation of the AND/OR tree rooted at that node.

$$\neg\exists x. (\text{AOLeaf}(x) \wedge ((\text{InputOne}\, x \wedge \neg\text{AOValOne}\, x) \vee (\neg\text{InputOne}\, x \wedge \text{AOValOne}\, x))) \wedge$$
$$\neg\exists x. (\text{IndDepth}_0\, x \wedge \text{IsOr}\, x \wedge \text{AOValOne}\, x \wedge \neg\exists y. (\text{AOChild}\, xy \wedge \text{AOValOne}\, y)) \wedge$$
$$\neg\exists x. (\text{IndDepth}_0\, x \wedge \neg\text{IsOr}\, x \wedge \text{AOValOne}\, x \wedge \exists y. (\text{AOChild}\, xy \wedge \neg\text{AOValOne}\, y)) \wedge$$
$$\neg\exists x. (\text{IndDepth}_0\, x \wedge \text{IsOr}\, x \wedge \neg\text{AOValOne}\, x \wedge \exists y. (\text{AOChild}\, xy \wedge \text{AOValOne}\, y)) \wedge$$
$$\neg\exists x. (\text{IndDepth}_0\, x \wedge \neg\text{IsOr}\, x \wedge \neg\text{AOValOne}\, x \wedge \neg\exists y. (\text{AOChild}\, xy \wedge \neg\text{AOValOne}\, y))$$

Note that some unary negations in these UNF formulas are not explicitly guarded by an atomic relation, but it is straightforward to modify the formulas to explicitly guard the negations with relations from ρ_n.

Let χ_n, χ'_n be UNF sentences over signatures σ_n and σ'_n, such that

$$\begin{aligned}\chi_n := {} & \mathrm{IndexedAOTrees} \wedge \mathrm{IndexIsDepth} \wedge \mathrm{ConsistentAOValues} \wedge \\ & \exists x.(\mathrm{AORoot}(x)) \wedge \\ & \neg\exists x.(\mathrm{IndDepth}_0\, x \wedge \neg\mathrm{IndMax}(x) \wedge \neg\exists y.(\mathrm{IndDepth}_0\, y \wedge \mathrm{AOChild}\, xy)) \wedge \\ & \neg\exists x.(\mathrm{AORoot}(x) \wedge \neg\mathrm{AOValOne}\, x)\end{aligned}$$

$$\begin{aligned}\chi'_n := {} & \neg[\mathrm{ConsistentAOValues}'] \vee \\ & \neg\exists x.(\mathrm{AORoot}(x) \wedge \neg\mathrm{AOValOne}'\, x)\end{aligned}$$

Note that in UNF we cannot enforce that a structure has an actual tree-shaped part corresponding to an indexed AND/OR tree; we can only pick out parts of the structure that can be viewed as indexed AND/OR trees (more precisely, parts of the structure that are GN bisimilar to such trees). Thus, χ_n expresses that (up to GN-bisimulation) there is an indexed AND/OR tree in the structure and every indexed AND/OR tree has depth 2^{2^n}. Moreover, $\mathrm{AOValOne}$ describes the internal value computations of every AND/OR tree, and $\mathrm{AOValOne}$ holds at the root of every such tree. Likewise, χ'_n expresses that $\mathrm{AOValOne}'$ describes the internal value computations of every AND/OR tree in the structure, and $\mathrm{AOValOne}$ holds at the root of every such tree.

We have $\chi_n \models \chi'_n$, and it is straightforward to check that $|\chi_n| + |\chi'_n|$ is polynomial in n.

E.2 GN bisimulation game and DAG size (Proof of Proposition 11)

Let $\mathfrak{A}$ and $\mathfrak{B}$ be structures over signatures containing σ, and let $\boldsymbol{a}$ and $\boldsymbol{b}$ be elements within $\mathfrak{A}$ and $\mathfrak{B}$, respectively. We write $\mathfrak{A}, \boldsymbol{a} \to^{k,m}_{\mathrm{gn}[\sigma]} \mathfrak{B}, \boldsymbol{b}$ if Duplicator has a winning strategy in the k-width m-round GN bisimulation game relative to signature σ starting from a partial rigid homomorphism $g : \boldsymbol{a} \mapsto \boldsymbol{b}$. We write $\mathfrak{A}, \boldsymbol{a} \sim^{k,m}_{\mathrm{gn}[\sigma]} \mathfrak{B}, \boldsymbol{b}$ if $\mathfrak{A}, \boldsymbol{a} \to^{k,m}_{\mathrm{gn}[\sigma]} \mathfrak{B}, \boldsymbol{b}$ and $\mathfrak{B}, \boldsymbol{b} \to^{k,m}_{\mathrm{gn}[\sigma]} \mathfrak{A}, \boldsymbol{a}$.

Recall the statement of Proposition 11:

> Assume $\mathfrak{A}, \boldsymbol{a} \to^{k,m}_{\mathrm{gn}[\sigma]} \mathfrak{B}, \boldsymbol{b}$ and $g : \boldsymbol{a} \mapsto \boldsymbol{b}$ is a witnessing partial rigid homomorphism from $\mathfrak{A}$ to $\mathfrak{B}$. Let $\varphi(\boldsymbol{x})$ be a DAG-represented GNF formula over signature σ such that $|\varphi| \leq m$, $\mathrm{width}(\varphi) \leq k$, and $\mathrm{free}(\varphi) = \boldsymbol{x}$. Then $\mathfrak{A} \models \varphi(\boldsymbol{a})$ implies $\mathfrak{B} \models \varphi(g(\boldsymbol{a}))$.

Assume that $\mathfrak{A}, \boldsymbol{a} \to^{k,m}_{\mathrm{gn}[\sigma]} \mathfrak{B}, \boldsymbol{b}$, and $g : \boldsymbol{a} \mapsto \boldsymbol{b}$ is a witnessing partial rigid homomorphism from $\mathfrak{A}$ to $\mathfrak{B}$.

We proceed by induction on m. The base case is trivial.

Let $m > 0$ and consider some DAG-represented GNF formula $\varphi(\boldsymbol{x})$ such that $|\varphi| \leq m$, $\mathrm{width}(\varphi) \leq k$, $\mathrm{free}(\varphi) = \boldsymbol{x}$ and $\mathfrak{A} \models \varphi(\boldsymbol{a})$. We must show that $\mathfrak{B} \models \varphi(\boldsymbol{b})$.

- Assume φ is of the form $\alpha(\boldsymbol{x})$. The result follows since g is a partial homomorphism.
- Assume φ is of the form $\psi_1 \wedge \psi_2$. For each $i \in \{1, 2\}$, we can consider the partial rigid homomorphism after Spoiler restricts to $\boldsymbol{a}_i \subseteq \boldsymbol{a}$ corresponding to $\mathrm{free}(\psi_i) \subseteq \boldsymbol{x}$. Since this is a valid move for Spoiler, Duplicator has a winning strategy in the $m-1$ round game from $g' := g \upharpoonright_{\boldsymbol{a}_i}$. Since $|\psi_i| \leq m - 1$ and $\mathrm{width}(\psi_i) \leq k$, the inductive hypothesis implies that $\mathfrak{B} \models \psi_i(g'(\boldsymbol{a}_i))$ for $i \in \{1, 2\}$, so $\mathfrak{B} \models \varphi(\boldsymbol{b})$.
  The proof is similar for φ of the form $\psi_1 \vee \psi_2$.
- Assume φ is of the form $\alpha(\boldsymbol{x}) \wedge \neg\psi(\boldsymbol{x}')$ where $\boldsymbol{x}' \subseteq \boldsymbol{x}$. Since $\mathfrak{A} \models \varphi(\boldsymbol{a})$, $\mathfrak{A} \models \alpha(\boldsymbol{a})$ and hence $\boldsymbol{a}$ is guarded in $\mathfrak{A}$. We also know that $\mathfrak{B} \models \alpha(\boldsymbol{b})$ since g is a partial homomorphism from $\mathfrak{A}$ to $\mathfrak{B}$. Let $\boldsymbol{a}'$ be the restriction of $\boldsymbol{a}$ to $\boldsymbol{x}'$. Since $\boldsymbol{a}$ is guarded, $\boldsymbol{a}'$ is guarded in $\mathfrak{A}$. Consider the play where Spoiler restricts to the guarded tuple $\boldsymbol{a}'$, and then switches structures. Since these are valid moves for Spoiler in the game, Duplicator has a winning strategy in the $m-2$ round game starting from $g' := (g \upharpoonright_{\boldsymbol{a}'})^{-1}$. Let $\boldsymbol{b}' = g(\boldsymbol{a}')$.
  Assume for the sake of contradiction that $\mathfrak{B} \models \psi(\boldsymbol{b}')$. Since $|\psi| \leq m-2$, the inductive hypothesis ensures that $\mathfrak{B} \models \psi(\boldsymbol{b}')$ implies $\mathfrak{A} \models \psi(g'(\boldsymbol{b}'))$, so we must have $\mathfrak{A} \models \psi(\boldsymbol{a}')$. But this contradicts the assumption that $\mathfrak{A} \models \varphi(\boldsymbol{a})$.
  Therefore, we must have $\mathfrak{B} \models \neg\psi(\boldsymbol{b}')$, so $\mathfrak{B} \models \varphi(\boldsymbol{b})$.
- Assume φ is the form $\exists y.\psi(\boldsymbol{x}y)$. Since $\mathfrak{A} \models \varphi(\boldsymbol{a})$, there must be some c such that $\mathfrak{A} \models \psi(\boldsymbol{a}c)$.
  If $c \in \boldsymbol{a}$, then let $g' := g$. Otherwise, let g' be Duplicator's choice of partial rigid homomorphism after Spoiler adds c. This is a valid move in the game, since $\mathrm{width}(\varphi) \leq k$ implies $|\boldsymbol{x}| < k$ and hence $|\boldsymbol{a}| < k$.
  In both cases, we must have $\mathfrak{A}, \boldsymbol{a}c \to^{k,m-1}_{\mathrm{gn}[\sigma]} \mathfrak{B}, \boldsymbol{b}d$ for $d := g'(c)$. Since $|\psi| \leq m-1$, the inductive hypothesis implies that $\mathfrak{B} \models \psi(\boldsymbol{b}d)$, so $\mathfrak{B} \models \varphi(\boldsymbol{b})$ as desired.

E.3 Counterexample using indexed AND/OR trees

We now seek to exhibit indexed AND/OR trees that cannot be distinguished by GNF formulas of certain sizes. We define inductively a sequence of indexed AND/OR trees $\mathfrak{A}_i$ and $\mathfrak{B}_i$ of depth $2i$ over signature ρ_n (technically, these structures should be indexed by n).

The basic structure of these AND/OR tree are defined recursively as follows, where we write a to denote a single a-labelled node, and $a(T_1, T_2)$ to denote a tree where the root is labelled a, the left child is the root of some tree T_1, and the right child is the root of some tree T_2:

- Let $\mathfrak{A}_0 := 1$ and $\mathfrak{B}_0 := 0$.
- Let $\mathfrak{A}_i := \mathrm{AND}(\mathrm{OR}(\mathfrak{B}_{i-1}, \mathfrak{A}_{i-1}), (\mathrm{OR}(\mathfrak{B}_{i-1}, \mathfrak{A}_{i-1}))$.
- Let $\mathfrak{B}_i := \mathrm{AND}(\mathrm{OR}(\mathfrak{B}_{i-1}, \mathfrak{A}_{i-1}), (\mathrm{OR}(\mathfrak{B}_{i-1}, \mathfrak{B}_{i-1}))$.

Evaluating $\mathfrak{A}_i$ (respectively, $\mathfrak{B}_i$) gives value 1 (respectively, 0).

A separate index tree (binary tree of depth n) is then attached to each node in the trees described above, where the index correctly describes the depth of this node in the AND/OR tree.

We view these indexed AND/OR trees as structures over the signature $\rho_n = \sigma_n \cap \sigma'_n$. In particular, this means that there are no internal computations using $\mathrm{AOValOne}$ or $\mathrm{AOValOne}'$. Although the children of nodes in the AND/OR tree are unordered, we will refer to the left (direction **l**) and right (direction **r**) child (based on the order of the children as shown above). This is just to aid in the description below.

We write $\mathfrak{B}_i(b)$ for the restriction of $\mathfrak{B}_i$ to the subtree rooted at b. For $i > 0$, let C_i (respectively, D_i) denote the set of nodes in $\mathfrak{A}_i$ (respectively, $\mathfrak{B}_i$) at depth 2 (these are the roots of the $\mathfrak{A}_{i-1}$ and $\mathfrak{B}_{i-1}$ subtrees of $\mathfrak{A}_i$ and $\mathfrak{B}_i$). We write $\boldsymbol{b}\restriction_b$ for the restriction of $\boldsymbol{b}$ to those elements appearing in $\mathfrak{B}_i(b)$.

Let $h : \boldsymbol{b} \mapsto \boldsymbol{a}$ be a position (partial rigid homomorphism) in the k-width GN bisimulation game between $\mathfrak{B}_i$ and $\mathfrak{A}_i$. For ease of presentation, we assume that h is restricted to elements in the AND/OR tree (rather than the elements in the index trees). This is without loss of generality since, in our construction below, h will always preserve the depth of positions, and the local structure of the index trees is identical in $\mathfrak{A}_i$ and $\mathfrak{B}_i$.

We define inductively what it means for h to be (i, k, m)-*safe*. We say a strategy in the m-round game is (i, k, m)-safe if it uses only (i, k, m')-safe positions where m' is the number of remaining moves in the game.

We say $h : \boldsymbol{b} \mapsto \boldsymbol{a}$ is (i, k, m)-safe if it satisfies the following properties:

(S1) for all $b \in \boldsymbol{b}$, $\mathrm{depth}(b) = \mathrm{depth}(h(b))$;

(S2) if $i > 0$, then there exists a mapping $f : D_i \to C_i$ such that for all $b, b' \in D_i$,

- if b, b' are siblings, then $f(b), f(b')$ are siblings;
- if $d \in \boldsymbol{b}$ is the parent of b, then $h(d)$ is the parent of $f(b)$;
- if $d \in \boldsymbol{b}$ appears in $\mathfrak{B}_i(b)$, then $h(d)$ appears in $\mathfrak{A}_i(f(b))$;
- $\mathfrak{B}_i(b), \boldsymbol{b}\restriction_b \to^{k,m}_{\mathrm{gn}[\rho_n]} \mathfrak{A}_i(f(b)), h(\boldsymbol{b}\restriction_b)$ via a strategy for Duplicator that is $(i-1, k, m)$-safe.

Likewise, we say a position $h : \boldsymbol{a} \mapsto \boldsymbol{b}$ is (i, k, m)-safe if the conditions (S1) and (S2) hold with $\mathfrak{B}_i$, $\boldsymbol{b}$, D_i exchanged with $\mathfrak{A}_i$, $\boldsymbol{a}$, C_i, respectively.

Observe that the empty partial rigid homomorphism h from $\mathfrak{B}_i$ to $\mathfrak{A}_i$ is (i, k, m)-safe: (S1) vacuously holds, and (S2) holds since (if $i > 0$) we can choose f such that sibling relationships are preserved, and $\mathfrak{A}_{i-1}$ (respectively, $\mathfrak{B}_{i-1}$) subtrees in $\mathfrak{B}_i$ are mapped to $\mathfrak{A}_{i-1}$ (respectively, $\mathfrak{B}_{i-1}$) subtrees in $\mathfrak{A}_i$. Similarly for the empty partial rigid homomorphism from $\mathfrak{A}_i$ to $\mathfrak{B}_i$.

More importantly, we can show that Duplicator has a winning (i, k, m)-safe strategy in certain games when starting from an (i, k, m)-safe position.

Lemma 41. *For all k, $i \leq 2^{2^n-1}$, and $m \leq i$,*

- *if $\boldsymbol{b} \mapsto \boldsymbol{a}$ is (i, k, m)-safe then $\mathfrak{B}_i, \boldsymbol{b} \to^{k,m}_{\mathrm{gn}[\rho_n]} \mathfrak{A}_i, \boldsymbol{a}$ via a strategy for Duplicator that is (i, k, m)-safe;*
- *if $\boldsymbol{a} \mapsto \boldsymbol{b}$ is (i, k, m)-safe then $\mathfrak{A}_i, \boldsymbol{a} \to^{k,m}_{\mathrm{gn}[\rho_n]} \mathfrak{B}_i, \boldsymbol{b}$ via a strategy for Duplicator that is (i, k, m)-safe.*

Proof. We fix k, and proceed by induction on i and $m \leq i$.

The base cases are covered (for $i = 0$ and $m = 0$, or $i > 0$ and $m = 0$), since Duplicator wins by default when $m = 0$.

Assume $i > 0$ and $m > 0$ and the starting position h is (i, k, m)-safe. We assume that the active structure is $\mathfrak{B}_i$, the current position is $h : \boldsymbol{b} \to \boldsymbol{a}$, and (S2) is witnessed by a map $f : D_i \to C_i$. For each possible move of Spoiler, we describe Duplicator's move, leading to an $(i, k, m-1)$-safe position h'.

By the inductive hypothesis (on m), this is enough to ensure that we can extend this to a full winning strategy for Duplicator.

- If Spoiler restricts to elements $\boldsymbol{b}' \subseteq \boldsymbol{b}$, then the new position $h' := h\restriction_{\boldsymbol{b}'}$ is still (i, k, m)-safe, and hence $(i, k, m-1)$-safe.
- If Spoiler switches structures to the new position $h' := \boldsymbol{a} \mapsto \boldsymbol{b}$ then $\boldsymbol{b}$ must be a single position or neighboring positions in the AND/OR tree (since $\boldsymbol{b}$ must be guarded). We must show h' is $(i, k, m-1)$-safe. (S1) trivially holds since no new elements have been added. We must define $f' : C_i \to D_i$ satisfying (S2) (for i, k, and $m-1$).

  Assume $\boldsymbol{b}$ corresponds to a single element b.

  If b is at depth at least 2, then let $d \in D_i$ be the node at depth 2 such that $\mathfrak{B}_i(d)$ contains b.
  Let $a := h(b)$ and $c := f(d) \in C_i$. We define f' such that $f'(c) := d$. Then $\mathfrak{A}_i(c)$ contains a and $\mathfrak{B}_i(d)$ contains b, and we can use the $(i-1, k, m)$-safe strategy from $\mathfrak{B}_i(d), b \to^{k,m}_{\mathrm{gn}[\rho_n]} \mathfrak{A}_i(f(d)), h(b)$ guaranteed by (S2) for h, to see that $\mathfrak{A}_i(c), a \to^{k,m-1}_{\mathrm{gn}[\rho_n]} \mathfrak{B}_i(f'(c)), h'(a)$ as desired.
  The sibling c' of c can be mapped by f' to the sibling d' of d (in fact, it would also work to map to d). If the sibling of c corresponds to an $\mathfrak{A}_{i-1}$ subtree, then $\mathfrak{A}_i(c')$ may correspond to an $\mathfrak{A}_{i-1}$ subtree and $\mathfrak{B}_i(d')$ to a $\mathfrak{B}_{i-1}$ subtree. Because $\mathrm{dom}(h')$ does not include any elements from these mismatched subtrees, the inductive hypothesis (on i) ensures that $\mathfrak{A}_{i-1}, \emptyset \to^{k,m-1}_{\mathrm{gn}[\rho_n]} \mathfrak{B}_{i-1}, \emptyset$ using an $(i-1, k, m-1)$-safe strategy, so $\mathfrak{A}_i(c'), \emptyset \to^{k,m-1}_{\mathrm{gn}[\rho_n]} \mathfrak{B}_i(d'), \emptyset$ using an $(i-1, k, m-1)$-safe strategy as desired.
  The remaining elements $c' \in C_i$ can be mapped to elements in positions **ll** and **lr** in $\mathfrak{B}_i$ such that if c' is the root of an $\mathfrak{A}_{i-1}$ (respectively, $\mathfrak{B}_{i-1}$) subtree in $\mathfrak{A}_i$, then $f'(c')$ is the root of an $\mathfrak{A}_{i-1}$ (respectively, $\mathfrak{B}_{i-1}$) subtree in $\mathfrak{B}_i$. In these subgames, Duplicator clearly has an $(i-1, k, m-1)$-safe winning strategy, since the arenas are isomorphic.

  If b is at depth 1, then $h(b)$ is also at depth 1 by (S1). Let c_1 and c_2 be the children of $h(b)$ in $\mathfrak{A}_i$. Set $f'(c_1)$ and $f'(c_2)$ to be the corresponding children of b in $\mathfrak{B}_i$. As in the previous case, this may result in mapping an $\mathfrak{A}_{i-1}$ subtree in $\mathfrak{A}_i$ to a $\mathfrak{B}_{i-1}$ subtree in $\mathfrak{B}_i$, but the inductive hypothesis (on i) ensures that (S2) still holds for these positions.
  The remaining elements $c' \in C_i$ can be mapped to elements in positions **ll** and **lr** in $\mathfrak{B}_i$ such that if c' is the root of an $\mathfrak{A}_{i-1}$ (respectively, $\mathfrak{B}_{i-1}$) subtree in $\mathfrak{A}_i$, then $f'(c')$ is the root of an $\mathfrak{A}_{i-1}$ (respectively, $\mathfrak{B}_{i-1}$) subtree in $\mathfrak{B}_i$. In these subgames, Duplicator clearly has an $(i-1, k, m-1)$-safe winning strategy, since the arenas are isomorphic.

If b is the root in $\mathfrak{B}_i$, then $h(b)$ must be the root in $\mathfrak{A}_i$. Let f' be the map **ll** $\mapsto$ **ll**, **lr** $\mapsto$ **lr**, **rl** $\mapsto$ **ll**, **rr** $\mapsto$ **lr**. Notice that we are 'cheating' for nodes in the right subtree of $\mathfrak{A}_i$ by mapping them to the left subtree in $\mathfrak{B}_i$. This clearly satisfies (S2), since we are mapping $\mathfrak{A}_{i-1}$ (respectively, $\mathfrak{B}_{i-1}$) subtrees in $\mathfrak{A}_i$ to $\mathfrak{A}_{i-1}$ (respectively, $\mathfrak{B}_{i-1}$) subtrees in $\mathfrak{B}_i$.

Now assume $\boldsymbol{b}$ is a pair of neighboring elements.
If $\boldsymbol{b}$ are neighboring positions both of depth at least 2, then the argument is similar to the depth 2 case above.
If $\boldsymbol{b}$ spans the root and depth 1, then the argument is similar to the depth 1 case above.
If $\boldsymbol{b}$ spans depth 1 and depth 2, then the argument is similar to the depth 2 case above.

- If Spoiler adds an element b in $\mathfrak{B}_i$, then we must select a in $\mathfrak{A}_i$ such that $h' := h[b \mapsto a]$ is $(i, k, m-1)$-safe.

  If b is at depth at least 2, then let b_2 be the node at depth 2 such that b is in $\mathfrak{B}_i(b_2)$. By (S2), $\mathfrak{B}_i(b_2), \boldsymbol{b} \upharpoonright_{b_2} \rightarrow^{k,m}_{\mathrm{gn}[\rho_n]} \mathfrak{A}_i(f(b_2)), h(\boldsymbol{b} \upharpoonright_{b_2})$ using an $(i-1, k, m)$-safe strategy. Select a based on Duplicator's move according to this strategy when Spoiler plays b in $\mathfrak{B}_i(b_2)$ (the strategy will give a position in $\mathfrak{A}_i(f(b_2))$, so a should be the element in the corresponding position in $\mathfrak{A}_i$). We know that $h[b \mapsto a]$ is $(i-1, k, m-1)$-safe restricted to $\mathfrak{B}_i(b_2)$ and $\mathfrak{A}_i(f(b_2))$. This ensures that (S1) holds. We set $f' := f$.
  We must check that $h' := h[b \mapsto a]$ is a partial rigid homomorphism from $\mathfrak{B}_i$ to $\mathfrak{A}_i$ (not just restricted to the $(i-1)$-subtrees).
  The delicate situation is if $b = b_2$, and the parent d of b is in $\boldsymbol{b}$, since this means $\mathrm{AOChild}\, db$ holds (recall that AOChild is the successor relation in the AND/OR tree). In that case, (S2) for h ensures that the parent of $f(b)$ is $h(d)$. But in order for h' to be $(i-1, k, m-1)$-safe restricted to $\mathfrak{B}_i(b_2)$ and $\mathfrak{A}_i(f(b_2))$, a must be $f(b)$. Since the parent of a is $h(d) = h'(d)$, this ensures that $\mathrm{AOChild}\, h'(d)h'(b)$ holds, so h' is still a partial rigid homomorphism.
  The conditions of (S2) for $(i, k, m-1)$-safety follow since h was (i, k, m)-safe using f, and $f' := f$.
  Thus, $h[b \mapsto a]$ is $(i, k, m-1)$-safe.

  If b is at depth 1, then let d_1 and d_2 be the nodes at depth 2 that are children of b. By (S2), it must be the case that $f(d_1)$ and $f(d_2)$ are children of a single node at depth 1 in $\mathfrak{A}_i$ (though $f(d_1)$ could be equal to $f(d_2)$). Select a to be the parent of $f(d_1)$ and $f(d_2)$. This ensures $h[b \mapsto a]$ is a partial rigid homomorphism (using the fact that if $d' \in \boldsymbol{b}$ is in $\mathfrak{B}_i(d_j)$ then $h(d')$ is in $\mathfrak{A}_i(f(d_j))$). (S1) still holds since a is a node at depth 1. (S2) still holds with $f' := f$ since no elements in depth at least 2 have changed, and we have respected siblings. Hence, $h[b \mapsto a]$ is $(i, k, m-1)$-safe.

  If b is at the root, then select a to be the root in $\mathfrak{A}_i$. This is a partial rigid homomorphism since depths of elements are preserved by h. (S1) clearly still holds. (S2) holds with $f' := f$. Thus, $h[b \mapsto a]$ is $(i, k, m-1)$-safe.

$\square$

E.4 Lowerbound on size of GNF interpolants for UNF (Proof of Theorem 10)

We are now ready to prove Theorem 10. The statement here is slightly different than in the body due to the fact that the AND/OR trees $\mathfrak{A}_i$ and $\mathfrak{B}_i$ defined in the previous subsection have depth $2i$ (so technically, the bound has to be adjusted by a factor of 2); of course, this has no impact on the lower bound.

There is a polynomial function p and a family of UNF sentences $\chi_n \models \chi'_n$ such that $|\chi_n| + |\chi'_n| \leq p(n)$, and there is no GNF interpolant θ_n for $\chi_n \models \chi'_n$ of size at most 2^{2^n-1}, even when θ_n is represented via a DAG.

Proof. Let $i = 2^{2^n-1}$. Then the indexed AND/OR trees $\mathfrak{A}_i$ and $\mathfrak{B}_i$ defined in the previous subsection are of depth $2i = 2^{2^n}$. In particular, note that $\mathfrak{A}_i \models \chi_n$ (when viewed as a σ_n structure), but $\mathfrak{B}_i \not\models \chi'_n$ (when viewed as a σ'_n structure).

By Lemma 41 (and the observation that the empty partial rigid homomorphism is safe), $\mathfrak{A}_i \sim^{i,i}_{\mathrm{gn}[\rho_n]} \mathfrak{B}_i$. Hence, by Proposition 11, $\mathfrak{A}_i$ and $\mathfrak{B}_i$ must agree on all DAG-represented GNF sentences φ over the signature ρ_n such that $|\varphi| \leq i$ and $\mathrm{width}(\varphi) \leq i$.

Since $\chi_n \models \chi'_n$, there is a GNF sentence θ_n over the signature ρ_n such that $\chi_n \models \theta_n$ and $\theta_n \models \chi'_n$.

Assume for the sake of contradiction that θ_n is of size at most i (which implies that there is also a DAG representation of θ_n of size at most i). Since $\mathfrak{A}_i \models \chi_n$, it must be the case that $\mathfrak{A}_i \models \theta_n$. Since a DAG representation of θ_n is at most size i and $\mathrm{width}(\theta_n) \leq i$, $\mathfrak{A}_i$ and $\mathfrak{B}_i$ must agree on θ_n, so $\mathfrak{B}_i \models \theta_n$. But $\mathfrak{B}_i \models \theta_n$ implies $\mathfrak{B}_i \models \chi'_n$, which is a contradiction since the value at the root of $\mathfrak{B}_i$ based on the internal calculations using $\mathrm{AOValOne}'$ is 0, not 1 as asserted by χ'_n. $\square$

E.5 Beth definability

We now show that there are UNF sentences ϕ_n over a signature $\sigma_n \cup \rho_n \cup \{R\}$ such that ϕ_n implicitly defines R over σ_n, but any sequence of DAG representations of GNF ϕ'_n over σ that provide explicit definitions for R with respect to ϕ must grow doubly-exponentially in the size of ϕ_n.

Let σ_n be as in the previous lower bound, $\rho_n = \sigma_n$, and $R = \mathrm{AOValOne}$. We can implicitly define AOValOne by:

there is AND/OR tree of depth 2^{2^n}, and
every AND/OR tree has depth 2^{2^n}, and
$\mathrm{ConsistentAOValues}$.

Assume for the sake of contradiction that there is a GNF formula $\psi(x)$ explicitly defining AOValOne over signature ρ_n, and such that $|\psi(x)| \leq 2^{2^n-1}$ (so there is a DAG representation with this size bound as well). Let $i = 2^{2^n-1}$ and let a, b be the roots of the structures $\mathfrak{A}_i$ and $\mathfrak{B}_i$ over signature σ_n, and satisfying the implicit definition above (i.e. such that AOValOne correctly labels intermediate calculations of the value of the AND/OR tree). We have $\mathfrak{A}_i \models \mathrm{AOValOne}\, a$, so $\mathfrak{A}_i \models \psi(a)$. But $\mathfrak{A}_i, a \sim^{k,m}_{\mathrm{gn}[\rho_n]} \mathfrak{B}_i, b$ for $i = k = m = 2^{2^n-1}$ by Lemma 41 (and observing that $a \mapsto b$ and $b \mapsto a$ are (i, k, m)-safe since a and b are the roots). This means $\mathfrak{B}_i \models \psi(b)$ by Proposition 11, so $\mathfrak{B}_i \models \mathrm{AOValOne}\, b$, a contradiction.

Since an explicit definition for $\mathrm{AOValOne}$ corresponds exactly to an interpolant, the lower bound follows from the lower bound for interpolation.

E.6 Lowerbound on size of GNF interpolants for GF (Proof sketch for Theorem 13)

In [12], a family of GF sentences φ_n using equality and using relations of unbounded arity are defined such that every model $\mathfrak{A}$ of φ_n contains a binary tree of depth 2^{2^n}. This is described in detail in the proof of Theorem 4.4 in [12] (where it is used to show that satisfiability is 2EXPTIME-hard for GF).

We can adapt this construction to define indexed AND/OR trees with small (polynomial in n) size GF formulas. A technical difference in this approach is that a node in the AND/OR tree is specified by a pair of elements in the structure rather than a single element, and equality is used in the GF formulas.

As before, this allows us to prove a doubly exponential lower bound on the size of GNF interpolants (even represented as a DAG) for GF validities, but it relies on relations of unbounded arity.

For the single exponential lower bound for relations of bounded arity, we can modify the construction in a different way. Instead of attaching index trees to each node in the AND/OR tree, we can introduce new unary relations $\mathrm{IndBit}_0, \ldots, \mathrm{IndBit}_n$, and require that for every node in the AND/OR tree, the index value (when we view these new predicates as n bits of a binary number) corresponds to the depth of this node in the AND/OR tree. This allows us to define indexed AND/OR trees of depth 2^n in GF, which can be used to prove an exponential lower bound on the size of GNF interpolants for GF validities with relations of bounded arity.

F. Preservation Theorems

Details of effective Łoś-Tarski for GF (Theorem 16)

Recall the statement:

> Let $\varphi(\boldsymbol{x})$ be a GF formula over signature σ with $|\varphi| = n$. If φ is preserved under extensions, then we can construct an equivalent existential GF formula $\varphi'(\boldsymbol{x})$ such that $|\varphi'| \leq 2^{2^{2^{2^{p(n)}}}}$ for a polynomial function p independent of φ.

For simplicity, we focus on the case of boolean queries. By Corollary 14, we can find a DAG representation of some GNF existential χ' that is equivalent to φ and is of doubly-exponential size. We can convert χ' to χ'' that is a union of χ''_i, where each χ''_i is a "GNCQ" – a conjunctive query with atomic negation, where every negated atom is guarded. Further, each GNCQ χ''_i within χ'' is at most exponential in the size of χ', and the overall size of χ'' (as a formula) is at most doubly exponential in the size of χ'. Let A_i be the set of queries in GF of the form $\exists \boldsymbol{y}. \bigwedge_j A_{ij}$, where $\bigwedge A_{ij}$ is obtained by identifying variables in χ''_i and then adding on at most $2^{|\varphi|}$ additional positive atoms. We claim that χ' is equivalent to $\bigvee_i \bigvee_{Q \in A_i} Q$. Since the set of indices i is doubly-exponential in the size of χ' and each A_i has size doubly-exponential in χ', this gives the desired bound. Clearly each query in A_i implies χ''_i and hence implies φ. On the other hand, consider a model M satisfying φ. Given any model M there is another structure M^* agreeing with M on all GF sentences that has a guarded tree decomposition. Indeed, the guarded unravelling of M (see, e.g., [2]) gives such an M^*. M^* must satisfy φ and therefore must satisfy some χ''_i via a homomorphism h. Let C be the image of h and Q_C be the conjunctive query that describes the restriction of M^* to the bags containing C – such a query can be written in GF by constructing it inductively from the leaves of the tree up. The size of Q_C is at most the size of χ''_i times the maximal size of bags – the latter being a factor which is at most exponential in the size of the schema. Note that every element $c \in C$ corresponds to exactly one variable y_c of Q_C. We let Q'_C extend Q_C be adding, for every negated atom in χ''_i, the result of replacing each variable x_i by $y_{h(x_i)}$. Then Q'_C also holds in M^* and $Q'_C \in$ GF. Hence Q'_C holds in M and is in A_i, which completes the argument.