Finite Satisfiability of Some Extensions of the Unary Negation Fragment

(Skończona spełnialność pewnych rozszerzeń logiki z unarną negacją)

Daniel Danielski

Praca magisterska

Promotor: dr hab. Emanuel Kieroński

Uniwersytet Wrocławski
Wydział Matematyki i Informatyki
Instytut Informatyki

17 sierpnia 2018

Abstract

We consider two extensions of the unary negation fragment of first-order logic. For the first, in which arbitrarily many binary symbols may be required to be interpreted as equivalence relations, we show the finite model property. More specifically, we prove that every satisfiable formula has a model of at most doubly exponential size. We argue that the finite satisfiability (= satisfiability) problem for this logic is 2-ExpTime-complete. For the second, the unary negation fragment with arbitrary number of transitive relations, we show that its finite satisfiability problem is 2-ExpTime-complete and that every finitely satisfiable formula has a triply-exponential model. Our result actually holds for a more general setting in which one can require that some binary symbols are interpreted as arbitrary transitive relations, some are interpreted as partial orders and some as equivalences. Additionally we consider finite satisfiability of two extensions of both primary logics. The first is a restricted variant of the guarded negation fragment with equivalence or transitive relations. In the second we add inclusions of binary relations.

Rozważamy dwa rozszerzenia logiki z unarną negacją. W pierwszym możemy żądać, aby dowolnie wiele relacji binarnych było interpretowanych jako równoważności. Pokazujemy własność modelu skończonego. Co więcej, dowodzimy, że każda spełnialna formuła ma model wielkości podwójnie wykładniczej. Argumentujemy, że problem skończonej spełnialności, tożsamy w tym wypadku z problemem spełnialności, jest 2-ExpTime-zupełny. W przypadku drugiego rozszerzenia, logiki z unarną negacją i dowolną liczbą relacji przechodnich, pokazujemy, że problem skończonej spełnialności jest 2-ExpTime-zupełny, a każda skończenie spełnialna formuła ma model potrójnie wykładniczej wielkości. Okazuje się, że nasz wynik jest również prawdziwy w nieco ogólniejszym scenariuszu, w którym możemy wymagać, żeby pewne symbole były interpretowane jako dowolne relacje przechodnie, pewne jako porządki częściowe, a pewne jako równoważności. Dodatkowo rozważamy dwa rozszerzenia powyższych logik: do pewnego ograniczonego fragmentu logiki ze strzeżoną negacją oraz o zawierania relacji binarnych.

Contents

Chapter 1

Introduction

Searching for attractive fragments of first-order logic is an important theme in theoretical computer science. Here, by an 'attractive fragment' we mean a fragment with reasonably simple and elegant definition, sufficient expressive power and which, at the same time, has decidable *satisfiability* (a problem of deciding whether a formula has a model). Such properties make a fragment potentially suitable for use in various automated reasoning tasks.

Several seminal decidable fragments of first-order logic were identified. Probably two most important ones are the *two-variable logic*, FO^2, [25], and the *guarded fragment*, GF, [2]. The restriction of FO^2 is that it uses only two variables and thus can speak non-trivially about only relations of arity one or two. GF, on the other hand, requires quantifiers to be appropriately relativized by atomic formulas, called *guards*. Both FO^2 and GF have decidable satisfiability problem. For FO^2 it is NExpTime-complete [13], for GF—2-ExpTime-complete and ExpTime-complete when the number of variables or the arity of relations is bounded [12]. Both have the *finite model property*, that is every satisfiable formula has a finite model, and GF has a *tree-like model property*—every satisfiable formula has a model of tree-like shape. What is also important, both FO^2 and GF contain *modal logic*, a formalism whose variations, including *description logics*, have been successfully applied in many areas of computer science.

Recently another formalism extending modal logic, called the *unary negation fragment*, UNFO, was proposed by ten Cate and Segoufin [27]. In this restriction of first-order logic negation is allowed only in front of subformulas with at most one free variable. UNFO turns out to have many good algorithmic and model theoretic properties, including the finite model property, a tree-like model property and the decidability of the satisfiability problem. The satisfiability problem is 2-ExpTime-complete, this time even if the number of variables is bounded (already three variables suffice for 2-ExpTime-hardness). Moreover UNFO may be interesting for the database community since it can express, e.g., unions of *conjunctive queries*, their negations, and the so called *frontier-one tuple generating dependencies* [3]. This allows us to solve some interesting instances of a classical problem from database theory, called *(finite) open-world query answering*, using (finite) satisfiability procedures for UNFO.

A serious weakness of the expressive power of UNFO, shared with FO^2 and GF, is that it cannot express transitivity of a binary relation, nor related properties, like being an equivalence, a partial order or a linear order. This is a severe limitation, since for example transitive relations can be used in database applications (consider relations like *greater-than* or *part-of*) and equivalences play some role in modal and epistemic logics as well as in XML reasoning [5, 6]. Thus it is natural to think about extensions of FO^2, GF or UNFO, in which some distinguished binary symbols may be explicitly required to be interpreted as transitive (or equivalence) relations. It turns out that such extensions for both FO^2 and GF are undecidable. Moreover, even the intersection of FO^2 and GF becomes undecidable, already in the presence of just two transitive relations [19, 18] or three equivalence relations [22, 18]. Some positive results were obtained for FO^2 and GF only with one transitive relation [24, 19] or two equivalences [21] or when the transitive (equivalence) relations are allowed only as guards [26, 19, 23]

UNFO turns out to be an exception here, since its satisfiability problem remains decidable in the presence of arbitrarily many equivalence or transitive relations. We show this by reducing the satisfiability problem for UNFO with equivalences, UNFO+EQ, to the one for UNFO with transitive relations, UNFO+TR. The latter embeds into both the *guarded negation fragment*, GNFO, with transitive relations restricted to non-guard positions (for more about this logic see Chapter 6) and UNFO with *regular path expressions*. The satisfiability problem for the last two logics was shown decidable and 2-ExpTime-complete, see [1] and [17], respectively.

Both the above mentioned decidability results are obtained by employing tree-like model properties of the logics and then using some automata techniques. Since tree-like unravelings of models are infinite, such approach works only for general satisfiability, and gives little insight into the decidability/complexity of the problem of determining satisfiability of formulas over finite models, *the finite satisfiability problem*. In computer science, the importance of decision procedures for finite satisfiability arises from the fact that most objects about which we may want to reason using logic, e.g., databases and programs, are finite. Thus the ability of solving only the general satisfiability problem may not be fully satisfactory.

In this thesis we show two main results. We prove that UNFO+EQ has the finite model property. It follows that the finite satisfiability and the general satisfiability problems for the considered logic coincide, and, due to the above mentioned reduction to general satisfiability of UNFO+TR, can be solved in 2-ExpTime. The corresponding lower bound can be obtained even for the two-variable version of the logic, in the presence of just two equivalence relations. Furthermore we consider finite satisfiability problem for UNFO+TR. Note that UNFO+TR does not have the finite model property—just look at the following formula with transitive T, $\forall x \exists y T(x,y) \wedge \forall x \neg T(x,x)$, satisfiable only in infinite models. However, we extend some ideas used to obtain the previous result and apply them to show 2-ExpTime-completeness of this problem. En route to this we obtain a triply exponential bound on the size of minimal models of finitely satisfiable UNFO+TR formulas. Actually, our results hold for a more general setting, in which some relations may be required to be equivalences, some as partial orders, and some other just as arbitrary transitive relations. Returning to our database motivation, we get the decidability of the *finite* open-world query answering for unions of conjunctive queries against frontier-one TGDs with equivalences, partial orders and arbitrary transitive relations. For more details about query answering with transitive data see [1].

To get a glimpse of what problems we encounter in the case of extensions of UNFO, let us compare UNFO+EQ with GF^2+EG—another logic that is decidable in the presence of arbitrarily many equivalence relations. The decidability of the satisfiability problem for both GF^2+EG and UNFO+EQ can be shown relatively easily, by exploiting tree-based model properties for both logics. The analysis of the corresponding finite satisfiability problems is more challenging. It turns out that the difficulties arising when considering GF^2+EG and UNFO+EQ are of different nature. The main problem in the case of GF^2+EG is that it allows to restrict some types of elements to appear at most once in every abstraction class. However, one can always construct models in which every pair of elements is connected by at most one equivalence. On the other hand, inequalities in UNFO+EQ trivialize and we do not have any problem with duplicating any elements, but UNFO+EQ allows for a non-trivial interaction among equivalences and this seems to be the main source of obstacles. Differences similar in spirit also appear in the case of transitive relations.

Our solutions employ a novel (up to our knowledge) inductive approach to build finite models, starting from some particular models. In the base of induction we construct some initial fragments in which none of the equivalences (transitive relations) play an important role. Such fragments are then joined into bigger and bigger structures, in which more and more equivalences (transitive relations) become significant.

We further transfer both our main results to the intersection of GNFO, [4], with equivalence relations (transitive relations) on non-guard positions and the one-dimensional fragment [14], BGNFO_1+EQ (BGNFO_1+TR). A formula is *one-dimensional* if its every maximal block of quantifiers leaves at most one variable free. Moving from UNFO to this restricted variant of GNFO significantly increases the expressive power.

Motivated mainly by examples from description logics we also try to strengthen our results by considering some extensions of UNFO+EQ and UNFO+TR covering a combination of *role hierarchies* and *inverse roles* (see, e.g. [15, 10]). In addition to an input formula φ we are given a set of inclusions of relations $\mathcal{H}$ of the form $B_1 \subseteq$

B_2, where B_1 and B_2 are arbitrary (inverses of), possibly equivalence (transitive), binary relations. We argue that the problems of verifying if φ has a finite model respecting $\mathcal{H}$ are decidable and remain 2-ExpTime-complete. (The corresponding general satisfiability problems are decidable as well, in the same complexity class, by much simpler arguments.)

In the world of description logics a lot of formalisms involving transitivity were investigated. The basic description logic allowing one to express transitivity is the logic S. The research on description logics led to the identification of several expressive examples with decidable satisfiability and finite satisfiability. However, description logics are essentially two-variable logics and speak about at most binary relations. Closer to our setting is the problem of (finite) ontology mediated query answering (F)OMQA, which is a counterpart of (finite) open-world query answering in the world of description logics. In this problem we are given a (multi-variable) conjunctive query (or a union of conjunctive queries) and a knowledge base specified in a DL. We need to answer whether the query holds in every model of the knowledge base. Not much is known, however, about the finite version of this problem. In particular, for DLs with transitive roles (S) the only positive results we are aware of are the ones obtained recently in [11], where the decidability and 2-ExpTime-completeness of FOMQA for the logics SOI, SIF and SOF is shown. This is orthogonal to our results, since UNFO+TR does not capture neither nominals (O) nor functional roles (F). On the other hand, we are able to express any positive boolean combinations of roles, including their intersection ($\sqcap$). This allows us to embed the logic SHI$^\sqcap$ (see, e.g., [10]), which equips S with role hierarchies, inverse roles and role conjunctions, in the obtained formalism equipping UNFO+TR with inclusions of relations. The (finite) satisfiability problem for SHI$^\sqcap$ is known to be 2-ExpTime-complete [9], exactly as for our logic. As a corollary, we obtain that FOMQA for SHI$^\sqcap$ is decidable and 2-ExpTime-complete. Up to out knowledge, this is the first decidability result for FOMQA in the case of a DL equipped with both transitive roles and role hierarchies. However, extending UNFO with possibly the simplest realization of the idea of capturing logics like SHIQ or SHOIQ (see, e.g., [28, 16]), that is adding *counting quantifiers*, immediately gives undecidability. Actually, even weaker *functional restrictions* suffice. This is implicit in [27].

Finally, we consider both types of extensions together in the cases of UNFO+EQ and UNFO+TR. Interestingly, one can express in BGNFO$_1$+EQ (BGNFO$_1$+TR) inclusions of the form $B_1 \subseteq B_2$, where B_1 is not an equivalence (transitive) relation, and our constructions additionally respect the inclusions of the form $B_1 \subseteq B_2$ for both B_1 and B_2 being equivalence (transitive) relations. We show that both the extensions of BGNFO$_1$+EQ and BGNFO$_1$+TR by inclusions of the form $T \subseteq B$ for equivalence (transitive) T and a non-equivalence (non-transitive) B become undecidable.

The thesis is organized as follows. Chapter 2 contains definitions, basic facts and high-level description of both of the main proofs. In Chapter 3 we work with tree-like models. These chapters are common for both the part of this paper involving equivalences and the one involving transitive relations. In Chapters 4 and 5 we build 'small' finite models for UNFO+EQ and UNFO+TR respectively. Chapter 4 is carried out in more sketchy way to familiarize the reader with the idea, while Chapter 5, containing a proof similar to the one in Chapter 4, but expanded to cover a more complicated case, contains more details. (In fact, the fragments concerning the finite model property for UNFO+EQ could have been removed and covered by a slight modification of the statements—not the proof methods—of some of the remaining theorems and lemmas; however allowing this kind of description is not a purpose of this thesis.) Finally, in Chapter 6 we consider the above mentioned further extensions of UNFO.

The thesis is largely based on two papers coauthored by the author (of this thesis) and the advisor, *The Unary Negation Fragment with Equivalence Has the Finite Model Property*, which was accepted for LICS 2018, and *Finite Satisfiability of Unary Negation Fragment with Transitivity*, which has just been submitted to a conference, each of them containing one of the two main results of this thesis. The material presented here consists of those (slightly rewritten towards a better and more unified presentation) parts of the above mentioned papers to which the author had non-trivial contribution. (According to the advisor 'The role of the author in deriving these results was dominant. This concerns, on the one hand, some conceptual work and devising appropriate technical tools, and, on the other hand, writing up the proofs.') In particular, the fragment of the first paper concerning the two-variable version of the logic, which was written mainly by the advisor, was omitted here (even though, according to the advisor, 'the ideas of the author were crucial also in this fragment'). The results are presented from a viewpoint slightly leaning towards the one of the author.

Chapter 2

Preliminaries

2.1. Logics, structures, types and functions

We employ standard terminology and notation from model theory. In particular, we refer to structures using Fraktur capital letters, and their domains using the corresponding Roman capitals. For a structure $\mathfrak{A}$ and $A' \subseteq A$ we use $\mathfrak{A}\restriction A'$ or $\mathfrak{A}'$ to denote the restriction of $\mathfrak{A}$ to A'.

The *unary negation fragment of first-order logic*, UNFO is defined by the following grammar [27]:

$$\varphi = B(\bar{x}) \mid x = y \mid \varphi \wedge \varphi \mid \varphi \vee \varphi \mid \exists x \varphi \mid \neg\varphi(x)$$

where, in the first clause, B represents any relational symbol, and, in the last clause, φ has no free variables besides (at most) x. An example formula not expressible in UNFO is $x \neq y$. We formally do not have universal quantification. However we allow ourselves to use $\forall\bar{x}\neg\varphi$ as an abbreviation for $\neg\exists\bar{x}\varphi$, for an UNFO formula φ. Note that $\forall xy\neg P(x, y)$ is in UNFO but $\forall xyP(x, y)$ is not.

We work with purely relational signatures $\sigma = \sigma_{\text{base}} \cup \sigma_{\text{dist}} \cup \sigma_{\text{aux}}$ where σ_{base} is the *base signature*, σ_{dist} is the *distinguished signature*, and σ_{aux} is the *auxiliary signature*. Assume that $\sigma_{\text{dist}} = \{E_1, \ldots, E_k\}$ and σ_{aux} is empty. The *unary negation fragment with equivalences*, UNFO+EQ, is defined by the same grammar as UNFO. When satisfiability of its formulas is considered, we restrict the class of admissible models to those that interpret all symbols from σ_{dist} as equivalence relations.
Now assume that $\sigma_{\text{dist}} = \{T_1, \ldots, T_k\}$, and $\sigma_{\text{aux}} = \{T_1^{-1}, \ldots, T_k^{-1}, E_1, \ldots, E_k\}$; all symbols in $\sigma_{\text{dist}} \cup \sigma_{\text{aux}}$ are binary. The *unary negation fragment with transitive relations*, UNFO+TR, is defined by the same grammar as UNFO, however when satisfiability of its formulas is considered, we restrict the class of admissible models to those that interpret all symbols from σ_{dist} as transitive relations, each T_u^{-1} as the inverse of T_u ($T_u^{-1}ab$ holds iff T_uba holds), and each E_u as the equivalence relation induced by T_u (E_uab holds iff $a = b$ or $T_uab \wedge T_uba$ holds). The intended role of the symbols from σ_{aux} is to simplify the presentation of our constructions, however

admitting them in formulas is not harmful since they can be easily, just by applying their definition, replaced by formulas using only $\sigma_{\rm dist}$ symbols so that the resulting formula is still in UNFO+TR. This way we can now have equivalence relations (just use some E_u without using corresponding T_u nor T_u^{-1}) and UNFO+EQ may be treated as a fragment of UNFO+TR. For simplicity of arguments and, in fact, without loss of expressive power, we consider for the rest of this thesis only formulas not using $\sigma_{\rm aux}$ symbols. For convenience we jointly enumerate all the symbols T_u and T_u^{-1} as $\overline{T}_1, \ldots, \overline{T}_{2k}$ assuming that $\overline{T}_{2u-1} = T_u$ and $\overline{T}_{2u} = T_u^{-1}$. In our constructions we sometimes consider some auxiliary structures in which symbols from $\sigma_{\rm dist}$ are not necessarily interpreted as transitive relations. Then, when we transitively close such relations we always assume that we take care of the proper interpretation of the symbols from $\sigma_{\rm aux}$ as well.

An *atomic* 1-*type* (or, shortly, a 1-*type*) over a signature σ is a maximal satisfiable set of literals (atoms and negated atoms) over σ with just one variable x. We sometimes identify a 1-type with the conjunction of its elements. Given a σ-structure $\mathfrak{A}$ and and an element $a \in A$ we denote by $\mathrm{atp}^{\mathfrak{A}}(a)$ the atomic 1-type *realized* by a, that is the unique 1-type $\alpha(x)$ such that $\mathfrak{A} \models \alpha(a)$.

We use various functions in this thesis. Given a function $f : A \to B$ we denote by $\mathrm{Rng} f$ its range, by $\mathrm{Dom} f$ its domain, and by $f \restriction A_0$ the restriction of f to $A_0 \subseteq A$.

2.2. Normal form, witnesses and basic facts

We say that an UNFO+EQ (resp. UNFO+TR) formula is in Scott-normal form if it is of the shape

$$\forall x_1, \ldots, x_t \neg\varphi_0(\bar{x}) \wedge \bigwedge_{i=1}^{m} \forall x \exists \bar{y} \varphi_i(x, \bar{y}) \tag{2.1}$$

where each φ_i is an UNFO+EQ (resp. UNFO+TR) quantifier-free formula and φ_0 is in negation normal form (NNF). A similar normal form for UNFO was introduced in the bachelor's thesis [8].

Lemma 2.1. *For every UNFO+EQ (resp. UNFO+TR) formula φ one can compute in polynomial time a normal form UNFO+EQ (resp. UNFO+TR) formula φ' over signature extended by some fresh unary symbols, such that any model of φ' is a model of φ and any model of φ can be expanded to a model of φ' by an appropriate interpretation of the additional unary symbols.*

Lemma 2.1 allows us, when dealing with decidability and complexity issues for UNFO+EQ or UNFO+TR, or when considering the size of minimal finite models of formulas, to restrict attention to normal form formulas.

Given a structure $\mathfrak{A}$, a normal form formula φ as in (2.1) and elements $a, \bar{b}$ of A such that $\mathfrak{A} \models \varphi_i(a, \bar{b})$ we say that the elements of $\bar{b}$ are *witnesses* for a and φ_i and that $\mathfrak{A}\restriction\{a,\bar{b}\}$ is a *witness structure* for a and φ_i. For an element a and every conjunct φ_i choose a witness structure $\mathfrak{W}_i$. Then the structure $\mathfrak{W} = \mathfrak{A}\restriction\{W_1 \cup \ldots \cup W_m\}$ is called a *φ-witness structure* for a.

In Chapters 3, 4 and 5 we present a construction which given some model of a normal form UNFO+EQ or UNFO+TR formula φ builds a finite model of φ of a bounded size. The construction has several stages. To argue that after some stages we still have a model of φ we use the following basic observation. The result below can be also used to prove inexpressibility of equivalence and transitive relations in pure UNFO. We leave it as a simple exercise for the reader.

Lemma 2.2. *Let $\mathfrak{A}$ be a model of a normal form UNFO+EQ or UNFO+TR formula φ. Let $\mathfrak{A}'$ be a structure in which all symbols from $\sigma_{\text{dist}} \cup \sigma_{\text{aux}}$ are interpreted as required, such that*

(h1) for every $a' \in A'$ there is a φ-witness structure for a' in $\mathfrak{A}'$,

(h2) for every tuple $a'_1, \ldots, a'_t$ (recall that t is the number of variables of the $\forall$-conjunct of φ) of elements of A' there is a homomorphism $\mathfrak{h} : \mathfrak{A}'\restriction\{a'_1, \ldots, a'_t\} \to \mathfrak{A}$ which preserves 1-types of elements.

Then $\mathfrak{A}' \models \varphi$.

Proof. Due to (h1) all elements of $\mathfrak{A}'$ have the required witness structures for all $\forall\exists$-conjuncts. It remains to see that the $\forall$-conjunct is not violated. But since $\mathfrak{A} \models \neg\varphi_0(\mathfrak{h}(a_1), \ldots, \mathfrak{h}(a_t))$ and φ_0 is a quantifier-free formula in which only unary atoms may be negated, it is straightforward, using (h2). □

2.3. Plan of small model constructions

Now we explain the general ideas behind both of the constructions. To prove the finite model property for UNFO+EQ, we show that the following are equivalent.

1. φ has a model
2. φ has a tree-like model
3. φ has a regular tree-like model with doubly exponentially many subtrees (up to isomorphism)
4. φ has a finite (doubly exponential) model

When considering the finite satisfiability of UNFO+TR, we show that the following are equivalent

1. φ has a finite model
2. φ has a tree-like model with paths of at most doubly exponential rank (to be defined later)

3. φ has a model as above that is regular with doubly exponentially many subtrees (up to isomorphism)
4. φ has a finite, triply exponential, model

The proofs of the implications (1) $\Rightarrow$ (2) $\Rightarrow$ (3) may be used to devise a 2-ExpTime algorithm checking whether the second condition holds. The following part may appear very informal or too high-level, however in the opinion of the author, it covers the essence of all the methods in the whole paper. In both proofs of (3) $\Rightarrow$ (4) we are going to use Lemma 2.2. The implications (1) $\Rightarrow$ (2) $\Rightarrow$ (3) may be treated as a preparation for such a usage, that is they create good pattern models, and mainly consist in cut-and-rearrange arguments. The proofs of the implications (3) $\Rightarrow$ (4) build structures that have a simple arrangement of paths (each of them locally resembles a tree) and the pattern structures should have lots of (partial) automorphisms, since this allows us to manipulate with some homomorphisms (by codomain switching) to join them into bigger and bigger ones. Therefore a natural candidate for a pattern model is a regular tree. The focal points of the whole paper are Claim 4.3 and its counterpart Claim 5.4. Such a division of the proofs also allows for simpler, modular arguments. Now we give some more details. We outline the latter construction and compare it with the former.

We are going to show that given an arbitrary finite model $\mathfrak{A}$ of a normal form UNFO+TR formula φ one can construct a model $\mathfrak{A}'$ of size bounded triply exponentially in the size of φ. The construction has several stages.

Stage 1: tree-like models. In this first step we simply unravel $\mathfrak{A}$ into an (infinite) tree-like model $\mathfrak{A}_1$. (See Lemma 3.1.)

*Stage 2: bounded transitive paths.*We introduce a measure associating with transitive paths in tree-like structures their *ranks*. A $\overline{T}_u$-path is a branch of a tree whose consecutive nodes are joined by $\overline{T}_u$ (and possibly by its inverse). Its $\overline{T}_u$-rank is the number of its one-way $\overline{T}_u$-edges. We show that, performing some surgery on $\mathfrak{A}_1$ we can build a tree-like model $\mathfrak{A}_2$ in which every transitive path has rank bounded doubly exponentially in $|\varphi|$. (See Lemma 3.5.) Note that a $\overline{T}_u$-path with a bounded $\overline{T}_u$-rank can still be arbitrarily long, as its two-way edges do not count to the rank.

*Stage 3: regular trees.*We then build another tree-like model $\mathfrak{A}_3$ which in addition to the properties of $\mathfrak{A}_2$ has only at most doubly exponentially many non-isomorphic subtrees. (See Lemma 3.8.)

*Stage 4: building finite model.*Finally, in this most complicated step, we extract from $\mathfrak{A}_3$ some *components* (their construction is inductive), and arrange some number of their copies to eventually form $\mathfrak{A}'$.

First of all, if a given formula φ belongs to UNFO+EQ then we can start our constructions leading to a finite model of φ from its arbitrary model, while if φ is in UNFO+TR we start from a *finite* model of φ. A very simple *Stage 1* in both arguments is, essentially, identical. The counterpart of *Stage 3* in the case of equivalences is slightly simpler, but the main differences lie in *Stage 2* and *Stage 4*. *Stage 2*, clearly, is not present at all in the construction for UNFO+EQ. While the general idea in this step is quite standard, as we just use a kind of tree pruning, the details are rather delicate due to possible interactions among different transitive relations—when decreasing the $\overline{T}_u$-rank of some path we often need to avoid increasing accidentally its $\overline{T}_v$-rank for some $v \neq u$. Regarding *Stage 4*, in the main part of the construction for UNFO+EQ we build bigger and bigger substructures in which some equivalence relations are total. The induction goes, roughly speaking, by the number of non-total equivalences in the substructure. In the construction for UNFO+TR we extend this approach to handle one-way transitive connections. Its single inductive step has two phases: building the so-called *components* and then arranging them into a bigger structure. It is this first phase in which we have now to work harder in the construction for UNFO+TR than in one for UNFO+EQ. Having components prepared, the scheme of joining them is similar in both constructions.

The whole construction may appear quite complex. It may be not a bad idea to look first at some simpler scenarios. A reader who wants such a warm-up is advised to consult paper [7] in which we explain how to construct finite models for satisfiable UNFO+EQ formulas with a simplification that only two variables are present. It may be helpful to look at this simplified fragment to get used to our idea of building finite models in an inductive way.

Chapter 3

Tree-like models

The arguments in this chapter are similar for both main problems considered in this thesis and can be carried out simultaneously. We will mark the differences, which exist mostly due to some issues with one-directional transitive connections. In this chapter we assume, for uniformity, that the distinguished signature, σ_{dist}, consists of the symbols $T_1, \ldots, T_k$, even in the case of UNFO+EQ formulas.

We use a standard notion of a (finite or infinite) rooted tree and related terminology. Additionally, any set consisting of a node and all its children is called a *family*. Any node b, except for the root and the leaves, belongs to two families: the one containing its parent is called the *upward family of* b and the one containing its children—the *downward family of* b.

We say that $\mathfrak{A}$ is a *tree-like structure* (*of degree bounded by* d) if its nodes can be arranged into a tree (of degree bounded by d) in such a way that if $\mathfrak{A} \models B(\bar{a})$ for some relation symbol B and $\bar{a} \subseteq A$ then either $\bar{a}$ is contained in a family, or B is a transitive relation and $B(\bar{a})$ follows from transitively closing the B-connections inside some families. Given a tree-like structure $\mathfrak{A}$ and $a \in A$ we denote by A_a the set of all nodes in the subtree rooted at a and by $\mathfrak{A}_a$ the corresponding substructure.

Given an arbitrary model $\mathfrak{A}$ of a normal form UNFO+EQ or UNFO+TR formula φ we can simply construct its tree-like model of degree bounded by $|\varphi|$. We define a φ*-tree-like unraveling* $\mathfrak{A}'$ of $\mathfrak{A}$ and a function $\mathfrak{h} : A' \to A$ in the following way. $\mathfrak{A}'$ is divided into levels $L_0, L_1, \ldots$. Choose an arbitrary element $a \in A$ and put to level L_0 of A' an element a' such that $\text{atp}^{\mathfrak{A}'}(a') = \text{atp}^{\mathfrak{A}}(a)$; set $\mathfrak{h}(a') = a$. This only element of L_0 will become the root of $\mathfrak{A}'$. Having defined L_i repeat the following for every $a' \in L_i$. Choose in $\mathfrak{A}$ a φ-witness structure for $\mathfrak{h}(a')$. Assume it consists of $\mathfrak{h}(a'), a_1, \ldots, a_s$. Add a fresh copy a'_j of every a_j to L_{i+1}, make $\mathfrak{A}' \restriction \{a', a'_1, \ldots, a'_s\}$ isomorphic to $\mathfrak{A} \restriction \{\mathfrak{h}(a'), a_1, \ldots, a_s\}$ and set $\mathfrak{h}(a'_j) = a_j$. Complete the definition of $\mathfrak{A}'$ transitively closing all relations from σ_{dist}.

Lemma 3.1. *Let $\mathfrak{A}$ be a model of a normal form UNFO+EQ or UNFO+TR formula φ. Let $\mathfrak{A}'$ be a φ-tree-like unraveling of $\mathfrak{A}$. Then $\mathfrak{A}' \models \varphi$ and $\mathfrak{A}'$ is a tree-like structure* *of degree bounded by* $|\varphi|$*.*

PROOF. It is readily verified that $\mathfrak{A}'$ meets the properties required by Lemma 2.2. In particular $\mathfrak{h}$ is the required homomorphism. That $\mathfrak{A}'$ is tree-like and has an appropriately bounded degree is also straightforward. □

We often work with tree-like models $\mathfrak{A}$ of normal form φ in which the downward family of every element a forms a φ-witness structure. In such case we call this downward family *the* φ-witness structure for a even if some other φ-witness structures for a exist in $\mathfrak{A}$.

3.1. Declarations

We now introduce an apparatus of *declarations* that allows us to perform some surgery on tree-like models of normal form formulas. Its main purpose is dealing with their universal conjuncts $\forall \bar{x} \neg \varphi_0(\bar{x})$.

For a normal form φ, a φ-declaration is a description of some patterns of connections taking into account the literals of φ_0 (and, for technical reasons, some additional transitive atoms, equalities and inequalities). In particular, it may describe some dangerous patterns, leading to a violation of φ_0. Let us give a formal definition.

Definition 3.2. Let φ be a normal form UNFO+EQ or UNFO+TR formula. Recalling that $\varphi_0 = \varphi_0(x_1, \ldots, x_t)$ is in NNF, and that k is the number of σ_{dist} relations, let $\mathcal{R}$ be the set consisting of all non-transitive literals (atoms or negated atoms) that appear in φ_0 (recall that only atoms which have at most one variable may be negated), $\mathcal{T} = \{1, \ldots, k\} \times \{1, \ldots, t\}^2$ and $\mathcal{Q} = \{1, \ldots, t\}$. A *$\varphi$-declaration* is a set consisting of some triples (R, T, Q) such that $R \subseteq \mathcal{R}$, $T \subseteq \mathcal{T}$ and $Q \subseteq \mathcal{Q}$.

A triple $d = (R, T, Q)$ may be alternatively viewed as a formula describing a pattern of connections on a tuple consisting of $t + 1$ (not necessarily distinct) elements: $\psi_d(x_1, \ldots, x_t, y) = \bigwedge_{r \in R} r(\bar{x}) \wedge \bigwedge_{(u,j,j') \in T} T_u x_j x_{j'} \wedge \bigwedge_{i \in Q} x_i = y \wedge \bigwedge_{i \in \mathcal{Q} \setminus Q} x_i \neq y$. Clearly, the notions of formulas and triples are equivalent and we use them interchangeably.

Let $\mathfrak{A}$ be a tree-like structure and $a \in A$. We say that a *respects* a φ-declaration $\mathfrak{d}$ if for each $\psi \in \mathfrak{d}$ we have $\mathfrak{A}_a \models \neg \exists \bar{x} \psi(\bar{x}, a)$. Given an element a in $\mathfrak{A}$, we denote by $\text{dec}^{\mathfrak{A}}_{\varphi}(a)$ the (unique) maximal declaration respected by a. Note that if a_0 is the root of $\mathfrak{A}$ then, knowing $\text{dec}^{\mathfrak{A}}_{\varphi}(a_0)$, we can determine if $\mathfrak{A} \models \forall \bar{x} \neg \varphi_0(\bar{x})$.

Let us now give some intuitions and describe how we are going to use declarations. We work with tree-like structures with φ-declarations assigned to all its nodes. Assigning $\mathfrak{d}$ to a node a of a structure $\mathfrak{A}$ may be treated as *a promise that none of the patterns described by* $\mathfrak{d}$ *appear in* $\mathfrak{A}_a$. Note that we do not require that $\mathfrak{d}$ equals

$\mathrm{dec}^{\mathfrak{A}}_{\varphi}(a)$. Given a system of declarations assigned to all nodes of $\mathfrak{A}$ we formulate some natural local conditions such that their violation at a node a breaks the promise of a (i.e., some forbidden pattern occurs), and, the other way round, they are sufficient to guarantee that for every node a the declaration $\mathfrak{d}$ assigned to a is a subset of $\mathrm{dec}^{\mathfrak{A}}_{\varphi}(a)$, which means that a respects $\mathfrak{d}$, that is, fulfills its promise. This allows us to proceed as follows: Take a tree-like model $\mathfrak{A} \models \varphi$, perform on it some surgery, obtaining a new tree-like structure $\mathfrak{A}'$. Assign to the nodes of $\mathfrak{A}'$ a system of φ-declarations in such a way that (i) the root of $\mathfrak{A}'$ gets the declaration $\mathrm{dec}^{\mathfrak{A}}_{\varphi}(a_0)$ where a_0 is the root of $\mathfrak{A}$, and (ii) for a node a' its downward family $F' = \{a', a'_1, \ldots, a'_s\}$ gets the declarations $\mathrm{dec}^{\mathfrak{A}}_{\varphi}(a), \mathrm{dec}^{\mathfrak{A}}_{\varphi}(a_1), \ldots, \mathrm{dec}^{\mathfrak{A}}_{\varphi}(a_s)$, where $F = \{a, a_1, \ldots, a_s\}$ is the downward family of some node a from $\mathfrak{A}$, and the structures on F and F' are isomorphic. This guarantees that the system of declarations satisfies the local conditions and thus that its promises are fulfilled. Due to the declaration of the root we have that $\mathfrak{A}'$ satisfies the universal conjunct of φ.

We are now ready for the details. Let $F = \{a, a_1, \ldots, a_s\}$ be the downward family of a node a. We say that a function $\mathfrak{f} : \{1, \ldots, t\} \to \{a, a_1, A^-_{a_1}, \ldots, a_s, A^-_{a_s}\}$ is a *fitting* (to F). We think that a fitting describes a distribution of elements of a t-element tuple of nodes of $\mathfrak{A}_a$ among the downward family of a and the subtrees rooted at the children of a. With a fitting $\mathfrak{f}$ we associate a function $\bar{\mathfrak{f}} : \{1, \ldots, t\} \to \{a, a_1, \ldots, a_s\}$ defined as follows: $\bar{\mathfrak{f}}(i) = a$ iff $\mathfrak{f}(i) = a$ and $\bar{\mathfrak{f}}(i) = a_j$ iff $\mathfrak{f}(i) \in \{a_j, A^-_{a_j}\}$. Let $b_1, \ldots, b_t = \bar{b} \subseteq A_a$. The fitting $\mathfrak{f}$ *induced by* $\bar{b}$ is defined naturally: $\mathfrak{f}(i) = a$ iff $b_i = a$, $\mathfrak{f}(i) = a_j$ iff $b_i = a_j$ and $\mathfrak{f}(i) = A^-_{a_j}$ iff $b_i \in A_{a_j} \setminus \{a_j\}$.

Let $\mathfrak{f}$ be a fitting, (R, T, Q) a tuple belonging to some declaration and $F = \{a, a_1, \ldots, a_s\}$ the downward family of some a. If r is a literal from R (resp. a tuple $(u, j, j') \in T$) then we denote by $\mathrm{Var}r$ the set of the indices of variables of r (resp. the set $\{j, j'\}$). If r is a literal from R or a literal $T_u x_j x_{j'}$ then we say that r is *fully fitted* (to F) if $\mathfrak{f}(\mathrm{Var}r) \subseteq F$.

Now we define the *local consistency conditions (LCCs)* for a system of declarations. Consider declarations $\mathfrak{d}, \mathfrak{d}_1, \ldots, \mathfrak{d}_s$ assigned to the elements of some family $F = \{a, a_1, \ldots, a_s\}$. We say that they *satisfy LCCs at a* if for each fitting $\mathfrak{f}$ and $\psi_{(R,T,Q)} \in \mathfrak{d}$ at least one of the following conditions holds.

(l1) Some R-conjunct $r \in R$ is not fully fitted, and $a \in \mathfrak{f}(\mathrm{Var}r)$ or $a_j, a_{j'} \in \mathfrak{f}(\mathrm{Var}r)$ for some $j \neq j'$.

(l2) Some R-conjunct $r \in R$ is fully fitted but $\mathfrak{A} \models \neg r(\mathfrak{f}(1), \ldots, \mathfrak{f}(t))$.

(l3) Some T-conjunct $T_u x_j x_{j'}$ is fully fitted but $\mathfrak{A} \models \neg T_u \mathfrak{f}(j)\mathfrak{f}(j')$.

(l4) For some Q-conjunct $x_i = a$ we have $\mathfrak{f}(i) \neq a$.

(l5) For some $(\boldsymbol{Q} \backslash Q)$-conjunct $x_i \neq a$ we have $\mathfrak{f}(i) = a$.

(l6) Some R-conjunct $r \in R$ is not fully fitted and is 'distributed over several subtrees', that is $|\bar{\mathfrak{f}}(\mathrm{Var}r) \cap \{a_1, \ldots, a_s\}| \geq 2$.

(l7) Two elements of two different subtrees cannot be transitively joined due to the structure on F, that is for some T-conjunct $T_u x_j x_{j'}$ we have $\bar{\mathfrak{f}}(j) \neq \bar{\mathfrak{f}}(j')$ but $\mathfrak{A} \models \neg T_u \bar{\mathfrak{f}}(j)\bar{\mathfrak{f}}(j')$.

(l8) All elements are fitted to a single subtree, that is for some i and all j we have $\bar{\mathfrak{f}}(j) = a_i$, and the promise is propagated to this subtree: $(R, T, \mathfrak{f}^{-1}(a_i)) \in \mathfrak{d}_i$.

(l9) There exists i such that $a_i \in \mathrm{Rng}\bar{\mathfrak{f}}$ and $\mathfrak{d}_i$ contains (R', T', Q') defined as follows: fix some $h \in \{1, \ldots, t\} \setminus \bar{\mathfrak{f}}^{-1}(a_i)$ and (i) $R' := \{r \in R : \mathrm{Var}r \subseteq \bar{\mathfrak{f}}^{-1}(a_i)\}$, (ii) T' is the minimal set such that for $(u, j, j') \in T$ if $\bar{\mathfrak{f}}(j) = \bar{\mathfrak{f}}(j') = a_i$ then $(u, j, j') \in T'$ and if $\bar{\mathfrak{f}}(j) = a_i$ (resp. $\bar{\mathfrak{f}}(j) \neq a_i$) and $\bar{\mathfrak{f}}(j') \neq a_i$ (resp. $\bar{\mathfrak{f}}(j') = a_i$) then $(u, j, h) \in T'$ (resp. $(u, h, j') \in T'$), and (iii) $Q' = \mathfrak{f}^{-1}(a_i) \cup \{h\} \cup (\boldsymbol{Q} \setminus \bar{\mathfrak{f}}^{-1}(a_i))$.

Note that the above conditions are of two sorts. Conditions (l1)–(l7) describe situations in which we immediately, just looking at the structure on F, observe that any tuple of elements corresponding to the given fitting does not break the promise of a. Conditions (l8)–(l9), on the other hand, describe situations in which, intuitively speaking, we need to relegate such observation to one of the children of a. Conditions (l1)–(l7) may appear a little redundant and partially cover each other. It follows from a further, slightly informal, division: (l1)–(l5) describe a situation where the promise is not broken just by looking only at the structure of F and (l6)–(l7) consider F in the scope of the entire structure on the subtree rooted at a.

Given a structure $\mathfrak{A}$ we say that a system of declarations $(\mathfrak{d}_a)_{a \in A}$ is *locally consistent* if it satisfies LCCs at each $a \in A$ and is *globally consistent* if $\mathfrak{d}_a \subseteq \mathrm{dec}^{\mathfrak{A}}_{\varphi}(a)$ for each $a \in A$. Note that the global consistency means that the promises of all nodes are fulfilled. Conditions (l1)–(l9) are tailored so that local and global consistency play along in the following sense.

Lemma 3.3. *Let $\mathfrak{A}$ be a tree-like structure. Then (i) if a system of declarations $(\mathfrak{d}_a)_{a \in A}$ is locally consistent then it is globally consistent (ii) the canonical system of declarations $(\mathrm{dec}^{\mathfrak{A}}_{\varphi}(a))_{a \in A}$ is locally consistent.*

Proof. (i) A contrario. Suppose that there exist $a \in A$, $\psi \in \mathfrak{d}_a$ and $\bar{b} \subseteq A_a$ such that $\mathfrak{A}_a \models \psi(\bar{b}, a)$. Take the fitting $\mathfrak{f}$ to the downward family $F = \{a, a_1, \ldots, a_s\}$ of a induced by $\bar{b}$. By the choice of $\bar{b}$, none of (l1)–(l7) holds. Thus $\bar{b} \not\subseteq F$ and there exist some $a_i \in F \setminus \{a\}$ and $\psi' \in \mathfrak{d}_{a_i}$ such that $\mathfrak{A}_{a_i} \models \psi'(\bar{b}', a_i)$ where $\bar{b}' = b'_1, \ldots, b'_t$ is defined as follows: $b'_j = b_j$ if $\bar{\mathfrak{f}}(j) = a_i$, otherwise $b'_j = a_i$. Denote by $\mathrm{depth}(\bar{b})$ the maximal level of $\mathfrak{A}$ inhabited by an element of $\bar{b}$. Obviously $\mathrm{depth}(\bar{b}') \leq \mathrm{depth}(\bar{b})$. Thus after finitely many steps we get a^*, $\psi^* \in \mathfrak{d}_{a^*}$ and $\bar{b}^*$ contained in the downward family F^* of a^* such that $\mathfrak{A}_{a^*} \models \psi(\bar{b}^*, a^*)$. But this cannot happen, since neither (l8) nor (l9) can hold for the fitting to F^* induced by $\bar{b}^*$.

(ii) Follows from a careful inspection of the definition of LCCs. Basically, if for some a, $\mathfrak{f}$ and $\psi \in \mathrm{dec}^{\mathfrak{A}}_{\varphi}(a)$ none of (l1)–(l9) holds then we can find $\bar{b} \subseteq A_a$ such that $\mathfrak{A} \models \psi(\bar{b}, a)$. Indeed, use non-satisfaction of (l8) and (l9) to find fragments of $\bar{b}$ belonging to the respective subtrees $(\bar{b} \cap \mathfrak{A}_{a_i})$ (i.e., to find appropriate $b_j \in A_{a_i}$

for $j \in \bar{\mathfrak{f}}^{-1}(a_i)$); non-satisfaction of (l1)–(l7) implies that they are connected so that $\psi(\bar{b}, a)$ holds. □

3.2. Shortening transitive paths

This section is featured only in the construction for UNFO+TR. The reader may safely omit it if interested only in the case of UNFO+EQ. However some part of the proof of Lemma 3.8, which is a part of the next section, refers to some part of the proof of Lemma 3.5, which is contained in this section. We mark it using a gray vertical line along the left side of the text.

Definition 3.4. Let $\mathfrak{A}$ be a tree-like structure. A sequence of nodes $a_1, \ldots, a_N \in A$ is a *path* in $\mathfrak{A}$ if for each i a_{i+1} is a child of a_i. A $\overline{T}_u$*-path* is a path such that for each i we have that $\mathfrak{A} \models \overline{T}_u a_i a_{i+1}$. The $\overline{T}_u$*-rank* of a $\overline{T}_u$-path $\vec{a}$, $\mathfrak{r}_u(\vec{a})$, is the cardinality of the set $\{i : \mathfrak{A} \models \neg\overline{T}_u a_{i+1} a_i\}$. The $\overline{T}_u$*-rank* of an element $a \in A$ is defined as $\mathfrak{r}_u(a) = \sup\{\mathfrak{r}_u(\vec{a}) : \vec{a} = a, a_2, \ldots, a_N; \vec{a} \text{ is a } \overline{T}_u\text{–path}\}$. For an integer M, we say that $\mathfrak{A}$ has $\overline{T}_u$-paths *bounded by* M when for all $a \in A$ we have $\mathfrak{r}_u(a) \leq M$, and that $\mathfrak{A}$ has *transitive paths bounded by* M if it has $\overline{T}_u$-paths bounded by M for all u.

Lemma 3.5. *If a normal form UNFO+TR formula φ has a finite model then it has a tree-like model of degree bounded linearly and transitive paths bounded doubly exponentially in $|\varphi|$.*

Proof. Let $\mathfrak{A}_0$ be a finite model of φ, $\mathfrak{A}$ its φ-tree-like unraveling, and $\mathfrak{h} : A \to A_0$ the function associated with this unraveling. By Lemma 3.1 $\mathfrak{A} \models \varphi$. It is easy to see that $\mathfrak{A}$ has transitive paths bounded by the size M_0 of $\mathfrak{A}_0$. Indeed, if not, then there exist u, a $\overline{T}_u$-path $(a_i)_{i=0}^N$ and indices $i_0, \ldots, i_{M_0}$ such that $\mathfrak{A} \models \overline{T}_u a_{i_j} a_{i_j+1} \wedge \neg\overline{T}_u a_{i_j+1} a_{i_j}$. Since $\mathfrak{h}$ preserves the structure on downward families we have $\mathfrak{A}_0 \models \overline{T}_u \mathfrak{h}(a_{i_j})\mathfrak{h}(a_{i_j+1}) \wedge \neg\overline{T}_u \mathfrak{h}(a_{i_j+1})\mathfrak{h}(a_{i_j})$. By the pigeonhole principle there exist $x < x'$ such that $\mathfrak{h}(a_{i_x}) = \mathfrak{h}(a_{i_{x'}})$. This gives, by transitivity, that $\mathfrak{A}_0 \models \overline{T}_u \mathfrak{h}(a_{i_x+1})\mathfrak{h}(a_{i_x})$. Contradiction.

Let $M_\varphi := |\boldsymbol{\alpha}| \cdot |D_\varphi| + 2$, where $\boldsymbol{\alpha}$ is the set of atomic 1-types over the signature of φ and D_φ is the set of φ-declarations. Clearly, M_φ is bounded doubly exponentially in $|\varphi|$.

Consider a mapping: $A \ni a \mapsto_{\mathfrak{g}} (\mathrm{atp}^{\mathfrak{A}}(a), \mathrm{dec}^{\mathfrak{A}}_{\varphi}(a))$. Observe that $|\mathrm{Rng}\mathfrak{g}| \leq M_\varphi - 2$. We construct a tree-like model $\mathfrak{A}'$ having levels $L'_0, L'_1, \ldots$. During our construction we maintain a function $\mathfrak{s} : A' \to \mathrm{Sym}(\{1, \ldots, 2k\})$ whose purpose is to define some order of shortening paths at a given node. Intuitively, for $\mathfrak{s}(a) = \sigma$, if $v < v'$ then we prefer to shorten $\overline{T}_{\sigma(v)}$ over $\overline{T}_{\sigma(v')}$.

Let L'_0 consist of a'_0—a copy of the root a_0 of $\mathfrak{A}$ (i.e. $\mathrm{atp}^{\mathfrak{A}'}(a'_0) = \mathrm{atp}^{\mathfrak{A}}(a_0)$). Put $\mathfrak{p}(a'_0) = a_0$ and set $\mathfrak{s}(a'_0)$ arbitrarily. Suppose that we have defined L'_i. For each $a' \in L'_i$ let $\{\mathfrak{p}(a'), a_1, \ldots, a_s\}$ be the downward family of $\mathfrak{p}(a')$ in $\mathfrak{A}$ and let

$\mathfrak{s}(a') = \sigma$. Take fresh copies a'_j of a_j and make $\mathfrak{A}'\restriction\{a', a'_1, \ldots, a'_s\}$ isomorphic to $\mathfrak{A}\restriction\{\mathfrak{p}(a'), a_1, \ldots, a_s\}$.

Presently we set the $\mathfrak{p}(a'_j)$ and $\mathfrak{s}(a'_j)$. Let $K^{a'_j} = K = \{v : \mathfrak{A} \models \neg\overline{T}_{\sigma(v)}\mathfrak{p}(a')a_j\}$ (the *killed* $\overline{T}_{\sigma(v)}$), $S^{a'_j} = S = \{v : \mathfrak{A} \models \overline{T}_{\sigma(v)}\mathfrak{p}(a')a_j \wedge \overline{T}_{\sigma(v)}a_j\mathfrak{p}(a')\}$ (the *sustained* $\overline{T}_{\sigma(v)}$) and $D^{a'_j} = D = \{v : \mathfrak{A} \models \overline{T}_{\sigma(v)}\mathfrak{p}(a')a_j \wedge \neg\overline{T}_{\sigma(v)}a_j\mathfrak{p}(a')\}$ (the *diminished* $\overline{T}_{\sigma(v)}$). If $D \neq \emptyset$ then let $v_D^{a'_j} = v_D = \min D$ and take as $\mathfrak{p}(a'_j)$ a $b_j \in A_{a_j}$ such that (i) $\mathfrak{g}(b_j) = \mathfrak{g}(a_j)$ (ii) for all $v < v_D$, $v \in S$: $\mathfrak{r}_{\sigma(v)}(b_j) \leq \mathfrak{r}_{\sigma(v)}(\mathfrak{p}(a'))$ (iii) $\mathfrak{r}_{\sigma(v_D)}(b_j)$ is the lowest possible. Note that such an element exists and $\mathfrak{r}_{\sigma(v_D)}(b_j) < \mathfrak{r}_{\sigma(v_D)}(\mathfrak{p}(a'))$. If $D = \emptyset$ then let $\mathfrak{p}(a'_j) = a_j$. If $K \neq \emptyset$ then let $v_K = \min K$ and set $\mathfrak{s}(a'_j) := \sigma'$ where σ' is defined as follows: for $v < v_K$ let $\sigma'(v) = \sigma(v)$, for $v_K \leq v < 2k$ let $\sigma'(v) = \sigma(v+1)$ and let $\sigma'(2k) = \sigma(v_K)$. Otherwise put $\mathfrak{s}(a'_j) = \sigma$. To finish the construction, transitively close all the appropriate relations in $\mathfrak{A}'$.

We claim that $\mathfrak{A}'$ constructed as above is a model of φ and has the desired properties.

$\forall\exists$-conjuncts are satisfied since for all $a' \in A'$ the structure on the downward family of a' in $\mathfrak{A}'$ is isomorphic to the structure on the downward family of $\mathfrak{p}(a')$ in $\mathfrak{A}$ and the latter is the φ-witness structure for $\mathfrak{p}(a')$.

For the universal conjunct of φ consider the canonical system of declarations on $\mathfrak{A}$ transported by $\mathfrak{p}$ to $\mathfrak{A}'$: $(\mathrm{dec}^{\mathfrak{A}}_{\varphi}(\mathfrak{p}(a')))_{a' \in A'}$. Note that in this system the declarations on the downward family of any node a' in A' are copies of the declarations on the downward family of $\mathfrak{p}(a')$ in $\mathfrak{A}$ in the canonical system of declarations on $\mathfrak{A}$. This canonical system on $\mathfrak{A}$ is locally consistent by part (ii) of Lemma 3.3. This in turn gives that the system we have defined on $\mathfrak{A}'$ is also locally consistent. By part (i) of Lemma 3.3 this system is also globally consistent. In particular, since φ_0 is equivalent to $\varphi_0^1 \vee \ldots \vee \varphi_0^s$ where the φ_0^j are conjunctions of some $\mathcal{R}$ and $\mathcal{T}$ formulas, for the root a'_0 of $\mathfrak{A}'$, for each j and $Q \subseteq \boldsymbol{Q}$ we have $\mathfrak{A}' = \mathfrak{A}'_{a'_0} \models \neg\exists\bar{x}(\varphi_0^j(\bar{x}) \wedge \bigwedge_{i \in Q} x_i = a_0 \wedge \bigwedge_{i \in \boldsymbol{Q}\setminus Q} x_i \neq a_0)$, so $\mathfrak{A}' \models \forall\bar{x}\neg\varphi_0(\bar{x})$.

That the degree of nodes in $\mathfrak{A}'$ is bounded linearly in φ follows from the fact that it was so bounded in $\mathfrak{A}$.

It remains to show that the transitive paths in $\mathfrak{A}'$ are doubly exponentially bounded. Let us first make an auxiliary estimation.

Claim 3.6. *Let v_0 and a $\overline{T}_u$-path $\vec{a} = (a_i)_{i=1}^N$ in $\mathfrak{A}'$ be such that for all i, i' and $v \leq v_0$ we have $\mathfrak{s}(a_i)(v) = \mathfrak{s}(a_{i'})(v)$ (in this case, slightly abusing notation, we write $\mathfrak{s}(\vec{a})(v) = \mathfrak{s}(a_i)(v)$) and $\mathfrak{s}(\vec{a})(v_0) = u$. Then $\mathfrak{r}_u(\vec{a}) \leq (M_\varphi+1)(\sum_{v<v_0} \mathfrak{r}_{\mathfrak{s}(\vec{a})(v)}(\vec{a})) + M_\varphi$.*

Proof. Induction on $v_0 = 1, \ldots, 2k$. Assume to the contrary that there is a $\overline{T}_u$-path $\vec{a}$ in $\mathfrak{A}$ meeting the required conditions such that $\mathfrak{r}_u(\vec{a}) > (M_\varphi+1)(\sum_{v<v_0} \mathfrak{r}_{\mathfrak{s}(\vec{a})(v)}(\vec{a})) + M_\varphi$. Then there are more than $M_\varphi(1+\sum_{v<v_0} \mathfrak{r}_{\mathfrak{s}(\vec{a})(v)}(\vec{a}))$ indices i such that $v_D^{a_i} = v_0$. So there exist indices $i_1, \ldots, i_{M_\varphi}$ such that for all j $v_D^{a_{i_j}} = v_0$ and for all $i_1 \leq i \leq i_{M_\varphi}$

and $v < v_0$ we have $v \notin D^{a_i}$ (and thus $v \in S^{a_i}$). It follows that for all $v \leq v_0$ the function $i \mapsto \mathfrak{r}_{\mathfrak{s}(\vec{a})(v)}(a_i)$ is non-increasing and the function $j \mapsto \mathfrak{r}_{\mathfrak{s}(\vec{a})(v_0)}(a_{i_j+1})$ is strictly decreasing. By the pigeonhole principle there exist $x < x' < M_\varphi$ satisfying $\mathfrak{g}(\mathfrak{p}(a_{i_x+1})) = \mathfrak{g}(\mathfrak{p}(a_{i_{x'}+1}))$. This contradicts the choice of $\mathfrak{p}(a_{i_x+1})$. □

The above claim allows us in particular to compute a (uniform) doubly exponential bound on $\mathfrak{r}_u(\vec{a})$ for all v_0, u and $\vec{a}$ as in assumption. Denote this bound by $\overline{M}_\varphi$.

Consider now any $\overline{T}_u$-path $\vec{a} = (a_i)_{i=1}^N$ in $\mathfrak{A}'$. For each node a_i from $\vec{a}$ let $v_u(a_i)$ be such that $\mathfrak{s}(a_i)(v_u) = u$. Due to the strategy that we use to define $\mathfrak{s}$ the value of v_u is non-increasing along $\vec{a}$. Indeed, when moving from a_i to $a_i + 1$ the value of v_u is either unchanged or decreases by 1; the only chance of increasing it would be to change it to $2k$ but this happens only when $\overline{T}_u$ is killed. Let us divide $\vec{a}$ into fragments $\vec{a}_1, \vec{a}_2, \dots$ on which v_u is constant. The number of such fragments is obviously bounded by $2k$. On each of such fragments $\vec{a}_i$ for all $v \leq v_u$ we have that $\mathfrak{s}(\vec{a}_i)(v)$ is constant. So we can apply Claim 3.6 to bound $\mathfrak{r}_u(\vec{a}_i)$ by $\overline{M}_\varphi$. This gives the desired doubly exponential bound $\mathbf{M}_\varphi = 2k\overline{M}_\varphi$ on $\mathfrak{r}_u(\bar{a})$. This finishes the proof of Lemma 3.5. □

3.3. Regular tree-like models

In this section we mark some parts of the text with solid (resp. dashed) vertical lines to indicate that they belong only to the argument concerning the finite model property for UNFO+EQ (resp. finite satisfiability of UNFO+TR). Additionally, a text between parentheses beginning with ♣ concerns only finite satisfiability of UNFO+TR.

We conclude this section by showing that for finitely satisfiable UNFO+TR formulas and satisfiable UNFO+EQ formulas we can always construct *regular* tree-like models (♣ with bounded transitive paths).

Let us introduce a tool, which allows us to verify the property of having bounded transitive paths looking only at some local conditions.

Definition 3.7. Let $\mathfrak{A}$ be a tree-like structure with root a_0. Then a function $\mathcal{S} : A \to \{0, \dots, M\}$ is a $(\overline{T}_u, M)$-*stopwatch labeling* if: $\mathcal{S}(a_0) = 0$; for every $a \in A$ and its child b: (i) if $\mathfrak{A} \models \overline{T}_u ab \wedge \overline{T}_u ba$ then $\mathcal{S}(b) = \mathcal{S}(a)$, (ii) if $\mathfrak{A} \models \overline{T}_u ab \wedge \neg \overline{T}_u ba$ then $\mathcal{S}(b) = \mathcal{S}(a) + 1$ (in particular $\mathcal{S}(a) < M$) (iii) if $\mathfrak{A} \models \neg \overline{T}_u ab$ then $\mathcal{S}(b) = 0$.

It is easy to see that $(\overline{T}_u, M)$-stopwatch labeling exists iff the structure has $\overline{T}_u$-paths bounded by M.

Recall that by Lemma 3.1 a satisfiable UNFO+EQ formula has a tree-like model—the unraveling of its model.

Lemma 3.8. *If φ has a tree-like model (♣ with doubly exponentially bounded transitive paths) with linearly bounded degree then it has a regular model having linearly bounded degree with doubly exponentially many non-isomorphic subtrees (♣ and transitive paths bounded doubly exponentially).*

Proof.

Let $\mathfrak{A}$ be a tree-like model of φ. Consider a mapping $A \ni a \mapsto_{\mathfrak{g}} (\text{atp}^{\mathfrak{A}}(a), \text{dec}_{\varphi}^{\mathfrak{A}}(a))$.

Let $\mathfrak{A}$ be a model of φ with doubly exponentially bounded paths (denote the bound $\mathbf{M}_{\varphi}$). For each $\overline{T}_u$ take a $(\overline{T}_u, \mathbf{M}_{\varphi})$-stopwatch labeling $\mathcal{S}_u$ of $\mathfrak{A}$. Consider the mapping $A \ni a \mapsto_{\mathfrak{g}} (\text{atp}^{\mathfrak{A}}(a), \text{dec}_{\varphi}^{\mathfrak{A}}(a), (\mathcal{S}_u(a))_{u=1}^{2k})$.

Note that $|\text{Rng}\mathfrak{g}|$ is bounded doubly exponentially in $|\varphi|$. We rebuild $\mathfrak{A}$ into a regular model $\mathfrak{A}'$.

For each $p \in \text{Rng}\mathfrak{g}$ choose a representative $\mathfrak{c}(p) \in \mathfrak{g}^{-1}(p)$. Put to L'_0 an element a'_0 such that $\text{atp}^{\mathfrak{A}'}(a'_0) = \text{atp}^{\mathfrak{A}}(a_0)$, where a_0 is the root of $\mathfrak{A}$, and let $\mathfrak{p}(a'_0) = \mathfrak{c}(\mathfrak{g}(a_0))$. Having defined L'_i, for $i \geq 0$, repeat the following for all $a' \in L'_i$. Denoting $\mathfrak{p}(a'), a_1, \ldots, a_s$ the downward family of $\mathfrak{p}(a')$ in $\mathfrak{A}$, add a fresh copy a'_i of each a_i to L'_{i+1} and make $\mathfrak{A}' \restriction \{a', a'_1, \ldots, a'_s\}$ isomorphic to $\mathfrak{A} \restriction \{\mathfrak{p}(a'), a_1, \ldots, a_s\}$. Set $\mathfrak{p}(a'_i) := \mathfrak{c}(\mathfrak{g}(a_i))$. Finally, transitively close all transitive relations in $\mathfrak{A}'$.

The proof that $\mathfrak{A}' \models \varphi$ is similar to the corresponding proof in Lemma 3.5, which is marked with a vertical line along the left side of the text. We observe that all elements have appropriate φ-witness structures copied from $\mathfrak{A}$ and then use the apparatus of declarations to argue that the $\forall$-conjunct of φ is respected. By construction $\mathfrak{A}'$ is a regular tree-like model with the number of different subtrees bounded by $|\text{Rng}g|$.

To see that $\mathfrak{A}'$ has paths bounded by $\mathbf{M}_{\varphi}$ create stopwatch labelings for $\mathfrak{A}'$ just by transferring them from $\mathfrak{A}$ using $\mathfrak{p}$. It is not difficult to see that they meet the conditions from the definition of stopwatch labelings.

□

Chapter 4

Small model theorem for UNFO+EQ

In this chapter we show the following small model theorem for UNFO+EQ.

Theorem 4.1. *Every satisfiable UNFO+EQ formula φ has a finite model of size bounded doubly exponentially in $|\varphi|$.*

Let us fix a satisfiable normal form UNFO+EQ formula φ, and the finite relational signature $\sigma = \sigma_{\text{base}} \cup \sigma_{\text{dist}}$ consisting of all symbols appearing in φ. Enumerate the equivalences as $\sigma_{\text{dist}} = \{E_1, \ldots, E_k\}$. Fix a regular tree-like σ-structure $\mathfrak{A} \models \varphi$ with at most doubly exponentially many non-isomorphic subtrees, which exists due to Lemma 3.1 and Lemma 3.8. We show how to build a finite model of φ.

Generally, we will work in an expected way, starting from copies of some elements of $\mathfrak{A}$, adding for them fresh witnesses (using some patterns of connections extracted from $\mathfrak{A}$), then providing fresh witnesses for the previous witnesses, and so on. At some point, instead of producing new witnesses, we need a strategy of using only a finite number of them. It is perhaps worth explaining what are the main difficulties in such a kind of construction. A naive approach would be to unravel $\mathfrak{A}$ into a tree-like structure, like in Lemma 3.1, then try to cut each branch of the tree at some point a and look for witnesses for a among earlier elements. The problem is when we try to reuse an element b as a witness for a, and b is already connected to a by some equivalence relations. Then, if a needs a connection to b by some other equivalences, adding them may eventually result in inconsistency with $\neg\varphi_0$. Another danger, similar in spirit, is that some b may be needed as a witness for several elements, $a_1, \ldots, a_s$. Then some of the a_i may become connected by some equivalences which, again, may be forbidden.

It seems to be a non-trivial task to find a safe strategy of providing witnesses using only finitely many elements and avoiding conflicts described above. This is why we employ a rather intricate inductive approach. We will produce substructures of the desired finite model in which some number of equivalences are total, using patterns extracted from the corresponding substructures of the original model. Intuitively, knowing that an equivalence is total, we can forget about it in our construction.

Roughly speaking, our induction goes on the number of equivalence relations that are not total in the given substructures. The constructed substructures will later become fragments of bigger and bigger substructures, which will eventually form the whole model. To enable composing bigger substructures from smaller ones in our inductive process we will additionally keep some information about the intended generalized types of elements in form of a pattern function pointing them to elements in the original model $\mathfrak{A}$. Here, the role of such a generalized type is played by the isomorphism type of the subtree (of $\mathfrak{A}$) rooted in the pattern element. Since we build substructures with some equivalences total on them, some more care is needed when the φ-witness structures are copied. They are not replicated entirely at once. Rather than that, we provide for each element only its *partial φ-witness structure*. This partial φ-witness structure is an isomorphic copy of the restriction of the φ-witness structure of the pattern element (a substructure of $\mathfrak{A}$) to the equivalence class of the intersection of the current total equivalences. The remaining part is completed in the further steps of construction.

An important property of the substructures created during our inductive process is that they admit some partial homomorphisms to the pattern tree-like model $\mathfrak{A}$ which restricted to (partial) witness structures act as isomorphisms into the corresponding parts of the φ-witness structures in $\mathfrak{A}$. Instead of introducing some technical notions concerning preserving such correspondence, we impose that every homomorphism respects the required condition using directly the structure of $\mathfrak{A}$. To this end we introduce further fresh (non-equivalence) binary symbols W^i whose purpose is to relate elements to their witnesses. We number the elements of the φ-witness structures in $\mathfrak{A}$ arbitrarily (recall that each element is a member of its own φ-witness structure) and interpret W^i in $\mathfrak{A}$ so that for each $a, b \in A$, $\mathfrak{A} \models W^i ab$ iff b is the i-th element of the φ-witness structure for a (from now, for short, we refer to the element b satisfying $W^i ab$ as the i-th witness for a). We do this in such a way that if two subtrees of $\mathfrak{A}$ were isomorphic before interpreting the W^i then they still are after such expansion. Now, if we mark b as the i-th witness for a during the construction (that is set $\mathfrak{A}' \models W^i ab$), then for any homomorphism $\mathfrak{h}$ we have $\mathfrak{A} \models W^i \mathfrak{h}(a)\mathfrak{h}(b)$.

Let us state the main inductive lemma.

Lemma 4.2. *Let $\mathcal{E}_0 \subseteq \sigma_{\text{dist}}$, $\mathcal{E}_{tot} = \sigma_{\text{dist}} \setminus \mathcal{E}_0$, $E^* = \bigcap_{E_u \in \mathcal{E}_{tot}} E_u$, $a_0 \in A$, $\mathfrak{A}_0$ be the induced substructure of $\mathfrak{A}$ on $A_{a_0} \cap [a_0]_{E^*}$. Then there exists a finite structure $\mathfrak{A}_0'$, an element (called the origin of $\mathfrak{A}_0'$) $a_0' \in A_0'$ and a function $\mathfrak{p} : A_0' \to A_0$ such that:*

- *(e1) E^* is total on $\mathfrak{A}_0'$.*
- *(e2) $\mathfrak{p}(a_0') = a_0$.*
- *(e3) For each $a' \in A_0'$ and each i, if the i-th witness for $\mathfrak{p}(a')$ lies in A_0 (that is $\mathfrak{A}_0 \models \exists y\, W^i \mathfrak{p}(a')y$) then there exists a unique element $b' \in A_0'$ such that $\mathfrak{A}_0' \models W^i a'b'$. Otherwise there exists no such element. Denote $W_{a'} = \{b' : \exists i\ \mathfrak{A}_0' \models W^i a'b'\}$ and for a tuple $\bar{a}$ let $W_{\bar{a}} = \bigcup_{a \in \bar{a}} W_a$.*
- *(e4) For each $\bar{a} \subseteq A_0'$ satisfying $|\bar{a}| \leq t$ there exists a homomorphism $\mathfrak{h} : \mathfrak{W}_{\bar{a}} \to \mathfrak{A}_0$*

such that for each $a \in \bar{a}$ we have $\mathfrak{A}_{\mathfrak{p}(a)} \cong \mathfrak{A}_{\mathfrak{h}(a)}$ and $\mathfrak{h}\restriction W_a$ is an isomorphism (onto its image). Moreover, if $a_0' \in \bar{a}$ then we can choose $\mathfrak{h}$ so that $\mathfrak{h}(a_0') = a_0$.

(e5) *For each $a \in A_0'$ we have $\mathfrak{W}_a \cong \mathfrak{W}\restriction A_0$ where $\mathfrak{W}$ is the φ-witness structure for $\mathfrak{p}(a)$. (Note that, by the definition of the W^i, each such isomorphism sends a to $\mathfrak{p}(a)$.)*

$\mathfrak{A}_0'$ may be seen as a small counterpart of $\mathfrak{A}_0$ in which each element has an appropriate fragment of the φ-witness structure. A reason for considering in (e4) a tuple $\bar{a}$ *together with* $\mathfrak{W}_{\bar{a}}$ is that a connection in a φ-witness structure may not include the element for which this witness structure is created.

The proof of Lemma 4.2 goes by induction on $l = |\mathcal{E}_0|$. Consider the base of induction, $l = 0$. In this case all equivalences in $\mathfrak{A}_0$ are total. Assume without loss of generality that $|A_0| = 1$ (if this is not the case, simply add an artificial equivalence relation E_{k+1} and interpret it as the identity on $\mathfrak{A}$). Take $\mathfrak{A}_0' := \mathfrak{A}_0$ and $\mathfrak{p}(a) = a$ for the only $a \in A_0$. Conditions (e1)–(e5) are obviously satisfied.

For the inductive step, suppose that lemma holds for $l - 1$. We show that it holds for l. Without loss of generality let $\mathcal{E}_0 = \{E_1, \ldots, E_l\}$. The rest of the proof is presented in Sections 4.1–4.3.

4.1. Pattern components

We create one pattern building block, called a *pattern component*, for each isomorphism type of a subtree rooted at a node of $\mathfrak{A}_0$. We denote by $\boldsymbol{\gamma}[A_0]$ the set of such isomorphism types. Let γ_{a_0} be the type of $\mathfrak{A}_{a_0}$. In the next section we take some number of copies of pattern components and join them forming the structure $\mathfrak{A}_0'$.

Take $\gamma \in \boldsymbol{\gamma}[A_0]$ and the root $a \in A_0$ of a subtree of type γ. If $\gamma = \gamma_{a_0}$, take $a = a_0$. We explain how to construct a pattern component $\mathfrak{C}^\gamma$. The component is divided into $l(2t+1)+1$ *layers* $L_1, \ldots, L_{l(2t+1)+1}$. The first $l(2t+1)$ of them are called the *inner layers* while the last one is called the *interface layer*. We start the construction of an inner layer L_i by defining its initial part, L_i^{init}, and then expand it to a full layer. The interface layer $L_{l(2t+1)+1}$ has no internal division but, for convenience, is sometimes referred to as $L_{l(2t+1)+1}^{init}$. The elements of $L_{l(2t+1)}$ are called *leaves* and the elements of $L_{l(2t+1)+1}$ are called *interface elements*.

$\mathfrak{C}^\gamma$ will have a shape resembling a tree, with structures obtained by the inductive assumption as nodes. All elements of the inner layers of $\mathfrak{C}^\gamma$ will have appropriate partial φ-witness structures provided. See Fig. 4.1.

We remark that during the process of building a pattern component we do not yet apply the transitive closure to the equivalence relations. Taking the transitive closures would not affect the correctness of the construction, but not doing this at this point will allow us for a simpler presentation of the correctness proof. Given a

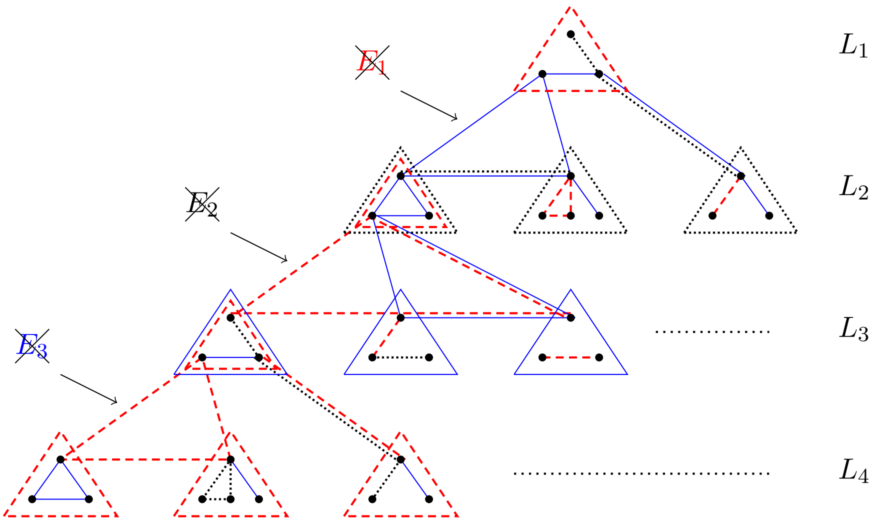

Figure 4.1: Top of a component for $l = 3$. Triangles correspond to subcomponents. Dashed lines represent E_1, dotted are used for E_2 and solid for E_3. L_i and L_{i+1} are not joined by E_i.

pattern component $\mathfrak{C}$ we will sometimes denote by $\mathfrak{C}_+$ the structure obtained from $\mathfrak{C}$ by applying the appropriate transitive closures. The crucial property we want to enforce is that the root of $\mathfrak{C}^\gamma$ will be *far* from its leaves in the following sense. Denote by $G_l(\mathfrak{S})$, for a σ-structure $\mathfrak{S}$, the Gaifman graph of the structure obtained by removing from $\mathfrak{S}$ the equivalences $E_{l+1}, \ldots, E_k$. Then there will be no connected induced subgraph of $G_l(\mathfrak{C}^\gamma_+)$ of size t containing an element of one of the first l layers and, simultaneously, an element of one of the last l inner layers of $\mathfrak{C}^\gamma$.

We set $L_1^{init} = \{a'\}$ to consist of a copy of element a, i.e., we set $\text{atp}^{\mathfrak{C}^\gamma}(a') := \text{atp}^{\mathfrak{A}_0}(a)$. Put $\mathfrak{p}(a') = a$. We call a' the *root* of $\mathfrak{C}^\gamma$.

Construction of a layer. Suppose we have defined levels $L_1, \ldots, L_{i-1}$ and L_i^{init}, $1 \leq i \leq l(2t+1)$, and the structure and the values of $\mathfrak{p}$ on $L_1 \cup \ldots \cup L_{i-1} \cup L_i^{init}$. We now explain how to define L_i and L_{i+1}^{init}. Let $v = 1 + (i - 1 \bmod l)$.

*Step 1: Subcomponents.*Take any element $c \in L_i^{init}$. From the inductive assumption we have a structure $\mathfrak{B}_0$ with $E^* \cap E_v$ total on it, its origin $b_0 \in B_0$ and a function $\mathfrak{p}_c : B_0 \to A_{\mathfrak{p}(c)} \cap [\mathfrak{p}(c)]_{E^* \cap E_v} \subseteq A_0$ with $\mathfrak{p}_c(b_0) = \mathfrak{p}(c)$. The substructures obtained owing to the inductive assumption are called *subcomponents*. We identify b_0 with c, add isomorphically $\mathfrak{B}_0$ to L_i, and extend function $\mathfrak{p}$ so that $\mathfrak{p} \restriction B_0 = \mathfrak{p}_c$. We do this independently for all $c \in L_i^{init}$.

Step 2: Providing witnesses. For $i < l(2t+1)+1$ we now define L_{i+1}^{init}. Take $c \in L_i$. Let $\mathfrak{W}$ be the φ-witness structure for $\mathfrak{p}(c)$ in $\mathfrak{A}$. Let $\mathfrak{F}$ be the restriction of $\mathfrak{W}$ to $[\mathfrak{p}(c)]_{E^* \cap E_v}$. Let $\mathfrak{F}'$ be the isomorphic copy of $\mathfrak{F}$ created for c in the subcomponent $\mathfrak{B}_0$ built in Step 1 that contains c ($\mathfrak{F}'$ exists due to (e5)). Let $\mathfrak{E} = \mathfrak{W} \restriction [\mathfrak{p}(c)]_{E^*}$. We add F''—a copy of $E \setminus F$ to L_{i+1}^{init}, and isomorphically copy the structure of $\mathfrak{E}$ to $F' \cup F''$ identifying F' with F. See Fig. 4.2. Note that this operation is consistent with the

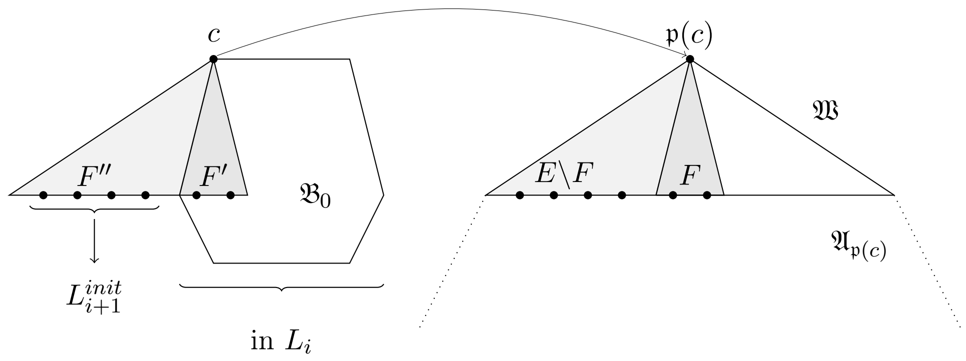

Figure 4.2: Providing witnesses.

previously defined structure on $\mathfrak{F}'$. The structure on $F' \cup F''$ will be the structure $\mathfrak{W}_c$ in $\mathfrak{C}^\gamma$ and then in $\mathfrak{A}_0'$. We define $\mathfrak{p}{\restriction}F''$ in a natural way, for each element $b \in F''$ choosing as the value of $\mathfrak{p}(b)$ the isomorphic counterpart of b in $E \setminus F$. We repeat this step independently for all for all $c \in L_i$.

When the interface layer, $L^{init}_{l(2t+1)+1}$ $(= L_{l(2t+1)+1})$, is created the construction of $\mathfrak{C}^\gamma$ is completed.

4.2. Joining the components

In this step we are going to arrange a number of copies of our pattern components to obtain the desired structure $\mathfrak{A}_0'$. We explicitly connect leaves of components with the roots of other components. This is done by identifying the interface elements of all components with some roots. We do it carefully, avoiding modifications to the internal structure of components, which could potentially result from transitivity of relations from σ_{dist}. In particular, a pair of elements that are not connected by an equivalence $E_u \in \mathcal{E}_0$ in $\mathfrak{C}_+$ will not become connected by a chain of E_u-connections external to $\mathfrak{C}_+$. To deal with the additional 'moreover' part of Condition (e4) we will simply define a_0' in such a way that it will not be used as a witness for any leaf.

As promised above we create pattern components for all types from $\boldsymbol{\gamma}[A_0]$. Let max be the maximal number of interface elements over all pattern components. For each $\mathfrak{C}^\gamma$ we number its interface elements. We create components $\mathfrak{C}^{\gamma,g}_{i,\gamma'}$ for all $\gamma, \gamma' \in \boldsymbol{\gamma}[A_0]$, $g \in \{0,1\}$ (g is often called a *color*), $1 \leq i \leq max$, as isomorphic copies of $\mathfrak{C}^\gamma$. We also create an additional component $\mathfrak{C}^{\gamma_{a_0},0}_{\perp,\perp}$ as a copy of $\mathfrak{C}^{\gamma_{a_0}}$, and define a_0' to be its root.

For each γ, g consider components of the form $\mathfrak{C}^{\gamma,g}_{\cdot,\cdot}$. Perform the following procedure for each i—the number of an interface element. Let b be the i-th interface element of any such component, let γ' be the type of $\mathfrak{A}_{\mathfrak{p}(b)}$. Identify the i-th interface elements of all $\mathfrak{C}^{\gamma,g}_{\cdot,\cdot}$ with the root c_0 of $\mathfrak{C}^{\gamma',1-g}_{i,\gamma}$. Note that the values of $\mathfrak{p}(c_0)$ and $\mathfrak{p}(b)$ (the latter equals to the value of $\mathfrak{p}$ on the i-th interface element in all the $\mathfrak{C}^{\gamma,g}_{\cdot,\cdot}$)

may differ. However, by construction, $\mathfrak{A}_{\mathfrak{p}(b)} \cong \mathfrak{A}_{\mathfrak{p}(c_0)}$ (in particular, the 1-types of b and c_0 match). For the element c^* obtained in this identification step we define $\mathfrak{p}(c^*) = \mathfrak{p}(c_0)$.

Finally, we take as $\mathfrak{A}_0^0$ the structure restricted to the components accessible in the graph of components from $\mathfrak{C}_{\perp,\perp}^{\gamma_{a_0},0}$. The graph of components G^{comp} is formed by joining a pair of components iff we identified the root of one of them with an interface element of the other.

We now define $\mathfrak{A}_0'$ as $\mathfrak{A}_0^0$ with transitively closed equivalences and set the root of $\mathfrak{C}_{\perp,\perp}^{\gamma_{a_0},0}$ to be its origin. Recall that in the structure $\mathfrak{A}_0^0$ we, exceptionally, do not transitively close σ_{dist}-connections, and thus allow the interpretations of the symbols from σ_{dist} not to be transitive (we will keep using superscript 0 for auxiliary structures of this kind).

4.3. Correctness of the construction

Now we proceed to the proof that $\mathfrak{A}_0'$ satisfies Conditions (e1)–(e5).

(e1) After taking the transitive closures, E^* is total on each pattern component. Thus, by the definition of the graph of components G^{comp}, E^* is total on $\mathfrak{A}_0'$.

(e2) Follows directly from the definition of L_1^{init} in $\mathfrak{C}_{\perp,\perp}^{\gamma_{a_0}}$ and the fact that $C_{\perp,\perp}^{\gamma_{a_0}} \subseteq A_0'$.

(e3) The interpretations of the W^i are defined in the step of providing witnesses where, implicitly, we take care of this condition for every element a' of the inner layers by extending the fragment of the partial φ-witness structure for a' created on the previous level of induction by a copy of a further fragment of the same pattern φ-witness structure. The identifications of elements during the step of joining the components do not spoil the required property and cause that it holds for all elements of $\mathfrak{A}_0'$.

(e4) This is the key part of our argumentation. For simplicity, let us ignore the 'moreover' part of this condition for some time. We will explain how to take care of it near the end of this proof. Now we find a homomorphism $\mathfrak{h}$ such that $\mathfrak{A}_{\mathfrak{p}(a)} \cong \mathfrak{A}_{\mathfrak{h}(a)}$ for all $a \in \bar{a}$ (we say that such a homomorphism has *the subtree isomorphism property*). Later we will show that its restrictions to the substructures $\mathfrak{W}_a$ are indeed isomorphisms. The proof consists of several homomorphic reductions performed in order to show that we can restrict attention to a structure built as a component but twice as high.

Reduction 0. Take $\bar{a} \subseteq A_0'$, $|\bar{a}| \leq t$. Observe that for each $a \in \bar{a}$ the structure $\mathfrak{W}_a$ is connected in $G_l(\mathfrak{A}_0' \restriction W_{\bar{a}})$ (recall the definition of Gaifman graph $G_l(\mathfrak{S})$ and the interpretation of the symbols W^i). Let $\mathfrak{W}_{\bar{a}_1}, \ldots, \mathfrak{W}_{\bar{a}_K}$ be the connected components of $\mathfrak{W}_{\bar{a}}$ in $G_l(\mathfrak{A}_0' \restriction W_{\bar{a}})$. If we have homomorphisms $\mathfrak{h}_i : \mathfrak{W}_{\bar{a}_i} \to \mathfrak{A}_0$, it is sufficient to

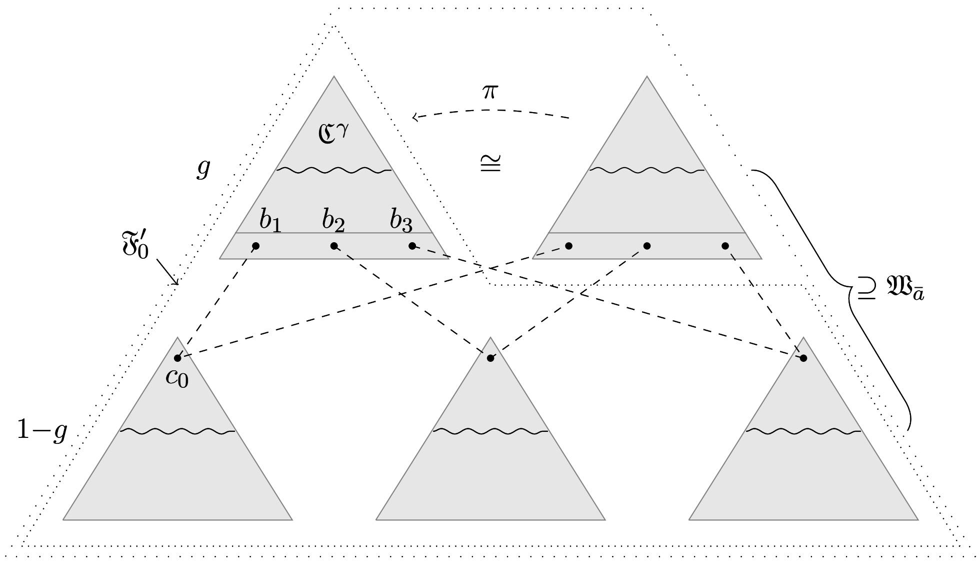

Figure 4.3: Joining the components and Reductions 1 and 2. Elements connected by dashed lines are identified.

put $\mathfrak{h} = \bigcup \mathfrak{h}_i$ as the desired homomorphism, since E^* is total on $\mathfrak{A}_0$ and for $a \in \bar{a}_i$ we also have $\mathfrak{A}_{\mathfrak{h}(a)} = \mathfrak{A}_{\mathfrak{h}_i(a)} \cong \mathfrak{A}_{\mathfrak{p}(a)}$. So *we can restrict attention to tuples* $\bar{a}$ *with* $\mathfrak{W}_{\bar{a}}$ *connected in the above sense.*

Reduction 1. The key fact is that, informally, $\mathfrak{W}_{\bar{a}}$ is contained 'on a boundary of two colors'. That is, there exists $g \in \{0, 1\}$ such that removing all the connections between leaves of color $1-g$ and roots of color g (in other words: any connections between elements of $L_{l(2t+1)}$ and elements of $L_{l(2t+1)+1}$ in components of color $1-g$) does not remove any connection among the elements of $\mathfrak{W}_{\bar{a}}$. This property follows from the fact that each subcomponent 'kills' one of the E_u, therefore, by the arrangement of subcomponents in a component, a connected $\mathfrak{W}_{\bar{a}}$ may be spread over a limited number of layers and the number of layers in a component is chosen high enough so the above property holds.

Reformulating, let $\mathfrak{D}_0^0$ be a structure obtained from $\mathfrak{A}_0^0$ by removing all direct connections between roots of color g and leaves of color $1-g$ and $\mathfrak{D}'_0$ its minimal extension in which equivalences are transitively closed. We have just proved that the inclusion map $\iota : \mathfrak{W}_{\bar{a}} \to \mathfrak{D}'_0$ is a homomorphism, and since for all $a \in \bar{a}$, $\mathfrak{A}_{\mathfrak{p}(\mathfrak{a})} = \mathfrak{A}_{\mathfrak{p}(\iota(\mathfrak{a}))}$, *we can restrict attention to a tuple* $\bar{a}$ *for which* $\mathfrak{W}_{\bar{a}}$ *is connected and search for a homomorphism* $\mathfrak{W}_{\bar{a}} \to \mathfrak{A}_0$ *treating* $\mathfrak{W}_{\bar{a}}$ *as a substructure of* $\mathfrak{D}'_0$.

Reduction 2. Consider the shape of a connected fragment of the graph of components G^{comp} with connections between leaves of color g and roots of color $1-g$ removed. Observe that there is at most one type γ of components of color g, chosen in the previous reduction, containing some element of $\mathfrak{W}_{\bar{a}}$ and all elements of $\mathfrak{W}_{\bar{a}}$ of color $1-g$ are contained in components of the form $\mathfrak{C}^{1-g}_{\cdot;\gamma}$. See Fig. 4.3. Now we can naturally 'project' all the elements of $\mathfrak{W}_{\bar{a}}$ of color g on one chosen component $\mathfrak{C}^\gamma$ of type γ and color g. Call this projection π. Then we remove from $\mathfrak{D}_0^0$ all components of color g other than $\mathfrak{C}^\gamma$ and all components of color $1-g$ of form other than $\mathfrak{C}^{1-g}_{;\gamma}$, obtaining a structure $\mathfrak{F}_0^0$. Let $\mathfrak{F}_0'$ be created by closing transitively all the equivalences in $\mathfrak{F}_0^0$. We claim that π is a homomorphism from $\mathfrak{W}_{\bar{a}}$ to $\mathfrak{F}_0'$. Indeed such projection can be applied to paths in $\mathfrak{D}_0^0$ to get corresponding paths in $\mathfrak{F}_0^0$. Since for all $a \in \bar{a}$ we have $\mathfrak{A}_{\mathfrak{p}(a)} = \mathfrak{A}_{\mathfrak{p}(\pi(a))}$, *we may restrict attention to a tuple $\bar{a}$ for which $\mathfrak{W}_{\bar{a}}$ is connected and search for a homomorphism $\mathfrak{W}_{\bar{a}} \to \mathfrak{A}_0$ treating $\mathfrak{W}_{\bar{a}}$ as a substructure of $\mathfrak{F}_0'$.*

Essential homomorphism construction. By the construction of $\mathfrak{A}_0'$ we can see that $\mathfrak{F}_0^0$ can be considered as a component of height $2l(2t+1)$ and such component can be viewed, as a tree τ whose nodes are subcomponents: we make subcomponent $\mathfrak{B}$ a parent of $\mathfrak{B}'$ iff $\mathfrak{B}'$ contains a witness for an element of $\mathfrak{B}$. We will build a homomorphism $\mathfrak{h} : \mathfrak{W}_{\bar{a}} \to \mathfrak{A}_{a_0}$ inductively using a bottom-up approach on tree τ. For a subcomponent $\mathfrak{B}$ denote by $B^\wedge$ the union of the domains of all the subcomponents belonging to the subtree of τ rooted at $\mathfrak{B}$.

Since we might have cut some connections between an element and some of its witnesses during Reduction 1, we define for each $a \in F_0'$ the surviving part $\mathfrak{V}_a$ of $\mathfrak{W}_a$ by $\mathfrak{V}_a = \mathfrak{F}_0' \restriction V_a$ where $V_a = \{b : \exists i\ \mathfrak{F}_0' \models W^i ab\}$. For a tuple $\bar{b}$ denote $V_{\bar{b}} = \bigcup_{b \in \bar{b}} V_b$ and $\mathfrak{V}_{\bar{b}} = \mathfrak{F}_0' \restriction V_{\bar{b}}$. Note that $V_a \subseteq W_a$, and generally, this inclusion may be strict, but for all $a \in \bar{a}$ we have $\mathfrak{V}_a = \mathfrak{W}_a$, and thus, in particular, the claim below finishes the proof of the currently considered part of (e4), that is the proof of the existence of a homomorphism satisfying the subtree isomorphism property.

Returning to the shape of $\mathfrak{F}_0^0$, it consists of some subcomponents arranged into tree τ glued together by the structure on the surviving parts. Note that all such building blocks (that is both the subcomponents and the surviving parts of the partial witness structures) are transitively closed. Moreover, by the tree structure of τ, if some elements of such a building block are connected by some atom in $\mathfrak{F}_0'$, then they already have been connected by the same atom in $\mathfrak{F}_0^0$, therefore the identity map from $\mathfrak{F}_0^0$ to $\mathfrak{F}_0'$ acts as an isomorphism when restricted to such a building block).

Recall that due to the expansion of the structure defined before the statement of Lemma 4.2, all homomorphisms $\mathfrak{A}_0' \to \mathfrak{A}_0$ respect the numbering of witnesses. This property will be particularly important in the proof of the following claim.

Claim 4.3. *For every subcomponent $\mathfrak{B}_0 \in \tau$ with origin b_0, and for every $\bar{a} \subseteq B_0^\wedge$, $|\bar{a}| \leq t$, there exists a homomorphism $\mathfrak{h} : \mathfrak{V}_{\bar{a}} \to \mathfrak{A}_{\mathfrak{p}(b_0)} \restriction [\mathfrak{p}(b_0)]_{E^*}$ such that for all $a \in \bar{a}$ we have $\mathfrak{A}_{\mathfrak{h}(a)} \cong \mathfrak{A}_{\mathfrak{p}(a)}$, and if $b_0 \in \bar{a}$ then $\mathfrak{h}(b_0) = \mathfrak{p}(b_0)$.*

Proof. Bottom-up induction on subtrees of τ.

Base of induction. In this case $\mathfrak{W}_{\bar{a}} \subseteq \mathfrak{B}_0$ and the claim follows from the inductive assumption of Lemma 4.2.

Inductive step. Let $\mathfrak{B}_1, \ldots, \mathfrak{B}_K$ be the list of those children of $\mathfrak{B}_0$ in τ for which

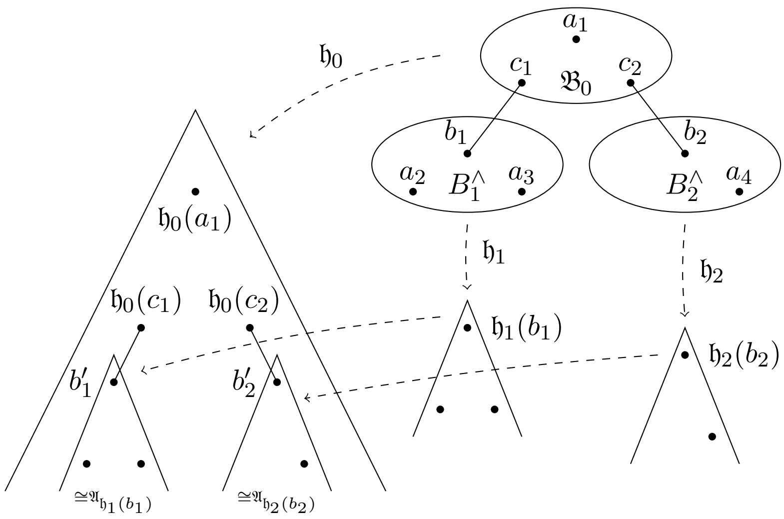

Figure 4.4: Joining homomorphisms

$B_i^\wedge$ contains some elements of $\bar{a}$; denote by b_i the root of $\mathfrak{B}_i$ and let $c_i \in B_0$ be such that b_i is a witness chosen by c_i in the step of providing witnesses/joining the components. If $K = 1$ and $\bar{a} \subseteq B_1^\wedge$ the thesis follows from the inductive assumption of this claim.

Otherwise, by the inductive assumption of this claim applied to $\mathfrak{B}_{(\bar{a}\cap B_i^\wedge)b_i}$ we have homomorphisms $\mathfrak{h}_i : \mathfrak{B}_{(\bar{a}\cap B_i^\wedge)b_i} \to \mathfrak{A}_{\mathfrak{p}(b_i)}$ satisfying $\mathfrak{p}(b_i) = \mathfrak{h}_i(b_i)$ and from the inductive assumption of Lemma 4.2 a homomorphism $\mathfrak{h}_0 : \mathfrak{B}_{(\bar{a}\cap B_0)c_1\ldots c_K}\restriction B_0 \to \mathfrak{A}_{\mathfrak{p}(b_0)}$. We extend the latter in the only possible way to $\mathfrak{h}_0^*$ defined on the whole $\mathfrak{B}_{(\bar{a}\cap B_0)c_1\ldots c_K}$: for each $a \in \bar{a}$ and $c \in V_a \setminus B_0$ (by construction $\mathfrak{B}_a \models W^i ac$ for some i) we set $\mathfrak{h}(c)$ to be the only element satisfying $\mathfrak{A}_0 \models W^i\mathfrak{h}(a)\mathfrak{h}(c)$ (such an element exists since $\mathfrak{A}_{\mathfrak{h}(a)} \cong \mathfrak{A}_{\mathfrak{p}(a)}$—in particular the φ-witness structures of $\mathfrak{h}(a)$ and $\mathfrak{p}(a)$ are isomorphic). Note that the sizes of the tuples used to build the homomorphisms $\mathfrak{h}_i$ are bounded by t, as required.

Using regularity of $\mathfrak{A}$, homomorphisms $\mathfrak{h}_0^*, \mathfrak{h}_1, \ldots, \mathfrak{h}_K$ can be joined together into $\mathfrak{h} : \mathfrak{B}_{\bar{a}b_1\ldots b_K c_1\ldots c_K} \to \mathfrak{A}_{\mathfrak{p}(b_0)}$ (see Fig. 4.4). In order to attach $\mathfrak{h}_i$ to $\mathfrak{h}_0^*$ we define $\mathfrak{h}_i^*$. Let j be such that b_i is the j-th witness for c_i and let b'_i be the j-th witness for $\mathfrak{h}_0(c_i)$ (it exists by $\mathfrak{A}_{\mathfrak{h}_0(c_i)} \cong \mathfrak{A}_{\mathfrak{p}(c_i)}$). Then we have $\mathfrak{A}_{b'_i} \cong \mathfrak{A}_{\mathfrak{p}(b_i)}$ since both b'_i and $\mathfrak{p}(b_i)$ are the j-th witnesses of some elements of $\mathfrak{A}$ being the roots of isomorphic subtrees. Thus, composing $\mathfrak{h}_i$ with such an isomorphism gives a homomorphism $\mathfrak{h}_i^* : \mathfrak{B}_{(\bar{a}\cap B_i^\wedge)b_i} \to \mathfrak{A}_{b'_i}$ with $\mathfrak{h}_i^*(b_i) = b'_i$. Finally we set $\mathfrak{h} = \bigcup \mathfrak{h}_i^*$. Note that $\mathfrak{h}$ is well defined (the value of $\mathfrak{h}$ on each of the b_i has been defined twice).

For each $a \in \mathrm{Dom}\mathfrak{h}_i$ ($= \mathrm{Dom}\mathfrak{h}_i^*$, when $i > 0$) we have $\mathfrak{A}_{\mathfrak{h}(a)} = \mathfrak{A}_{\mathfrak{h}_i^*(a)} \cong \mathfrak{A}_{\mathfrak{h}_i(a)}(\cong \mathfrak{A}_{\mathfrak{p}(a)}$, by the inductive assumptions of this claim and Lemma 4.2). Since $\bar{a} \subseteq \mathrm{Dom}\mathfrak{h}_0 \cup \bigcup_{i>0} \mathrm{Dom}\mathfrak{h}_i^*$, we get that for each $a \in \bar{a}$ we have $\mathfrak{A}_{\mathfrak{p}(a)} \cong \mathfrak{A}_{\mathfrak{h}(a)}$.

Recalling the tree structure on τ we can conclude that $\mathfrak{h}$ is a homomorphism. We give an idea of the proof of this property. Consider an E_u-path in $\mathfrak{F}_0^0$ connecting two elements of $\mathfrak{V}_{\bar{a}b_1\dots b_K c_1\dots c_K}$. We show that the images of these two elements are connected by an E_u-path in $\mathfrak{A}$. Using the tree shape of $\mathfrak{F}_0^0$, we can split it into parts contained in B_0 or some of the $B_i^\wedge$, and parts contained in some of the V_d for $d \in (\bar{a} \cap B_0)c_1 \dots c_K$ (with the splitting points belonging to $V_{(\bar{a}\cap B_0)c_1\dots c_K}$). For the former type of connections, use the fact that $\mathfrak{h}_0, \mathfrak{h}_1^*, \dots, \mathfrak{h}_K^*$ are homomorphisms. For the latter, observe that $\mathfrak{h}_0^*$ sends V_d into the corresponding part of an isomorphic copy of the pattern φ-witness structure from $\mathfrak{A}$ used to define the structure on $\mathfrak{F}_0^0 \restriction V_d$. Similarly a non-transitive relation in $\mathfrak{F}_0^0$ may connect elements contained in B_0 or one of the $B_i^\wedge$, or one of the V_d for $d \in (\bar{a} \cap B_0)c_1 \dots c_K$, and the argument as above shows that it is preserved by $\mathfrak{h}$.

It follows from the construction that $\mathfrak{h}$ has the following property: if $b_0 \in \bar{a}$ then $\mathfrak{h}(b_0) = \mathfrak{h}_0(b_0) = \mathfrak{p}(b_0)$. To finish the proof of the inductive step, we restrict $\mathfrak{h}$ to $V_{\bar{a}}$. □

Now we prove the additional property required for $\mathfrak{h}$ by (e4), namely that $\mathfrak{h}\restriction W_a$ is an isomorphism. Observe that $\mathfrak{h}$ injectively moves $\mathfrak{W}_a$ into the corresponding part of the φ-witness structure for $\mathfrak{h}(a)$ which is isomorphic to the corresponding part of the φ-witness structure for $\mathfrak{p}(a)$ by the subtree isomorphism property. Therefore, since the structure on W_a (prior to taking the transitive closure) was copied from the latter, the inverse of $\mathfrak{h}\restriction W_a$ is a homomorphism and therefore $\mathfrak{h}\restriction W_a$ is an isomorphism.

To prove the 'moreover' part of (e4), it suffices to observe that if $a_0' \in \bar{a}$ then in Reduction 1 we have that $g = 0$ and in Reduction 2 we have that $\gamma = \gamma_{a_0}$. We choose $\mathfrak{C}^\gamma = \mathfrak{C}^{\gamma,0}_{\perp,\perp}$. This way the application of the Reductions does not move a_0'. The claim follows from the fact that $\mathfrak{h}(a_0') = \mathfrak{p}(a_0') = a_0$.

(e5) Apply (e4) to a tuple consisting of just a to obtain an isomorphism $\mathfrak{h} : \mathfrak{W}_a \to \mathfrak{A}_0 \restriction \mathfrak{h}(W_a)$ and then apply an isomorphism between $\mathfrak{A}_{\mathfrak{h}(a)}$ and $\mathfrak{A}_{\mathfrak{p}(a)}$.

This finishes the proof of Lemma 4.2. Let us show how this lemma implies the finite model property for UNFO+EQ. Take $\mathcal{E}_0 = \sigma_{\text{dist}}$, let a_0 be the root of $\mathfrak{A}$. We apply Lemma 4.2 and get a finite structure $\mathfrak{A}_0'$ and a function $\mathfrak{p} : A_0' \to A_0$. Note, that $\mathfrak{A}_0 = \mathfrak{A}$. Let us see that $\mathfrak{A}_0'$ satisfies the conditions of Lemma 2.2. Indeed, (h1) follows from (e5). Condition (h2) follows from (e4). So $\mathfrak{A}_0' \models \varphi$.

4.4. Size of models and complexity

Now we show, that the size of $\mathfrak{A}_0'$ is bounded doubly exponentially in $|\varphi|$. We calculate a recurrence relation on $\mathbf{C}_l$—an upper bound on the size of the structure created in the l-th step of induction. We are interested in an estimate for $\mathbf{C}_{k+1}$ (note that in the base of induction we may have added an additional equivalence relation—this is why we consider $\mathbf{C}_{k+1}$, rather than $\mathbf{C}_k$).

Let $n = |\varphi|$. Consider the l-th induction step. The size of each subcomponent is bounded by $\mathbf{C}_{l-1}$. Consider one component. Layer L_1 consists of at most $\mathbf{C}_{l-1}$ elements, each of them creates at most n elements in layer L_2^{init}, which jointly create at most $\mathbf{C}_{l-1}\cdot n\cdot \mathbf{C}_{l-1}$ elements in layer L_2 and inductively at most $\mathbf{C}_{l-1}^{i} n^{i-1}$ elements in layer L_i. So each component has at most $\mathbf{C}_{l-1}^{l(2t+1)+2} n^{l(2t+1)+2}$ elements. Counting the components, we get an estimate

$$\mathbf{C}_l = \mathbf{C}_{l-1}^{8n^2}\cdot n^{8n^2}\cdot (|\boldsymbol{\gamma}[A]|\cdot 2\cdot \mathbf{C}_{l-1}^{8n^2}\cdot |\boldsymbol{\gamma}[A]| + 1).$$

Solving this recurrence relation we get

$$\mathbf{C}_{k+1} \leq (|\boldsymbol{\gamma}[A]|^2\cdot 4\cdot n^{8n^2})^{(16n^2)^{n+1}},$$

which is doubly exponential in n.

We conclude this chapter with the following theorem.

Theorem 4.4. *The finite satisfiability problem for UNFO+EQ is 2-*ExpTime*-complete.*

Proof. By Thm. 4.1 finite satisfiability and general satisfiability for UNFO+EQ coincide. The lower bound is inherited from pure UNFO [27]. As a side note, the lower bound can be also obtained for the two-variable UNFO with two equivalence relations. In this case the proof is an adaptation of the lower bound proof for GF^2 with equivalence relations in guards [19]. For the upper bound, recall our remarks in the definition of UNFO+TR where we explain that UNFO+EQ may be embedded in UNFO+TR. The latter is in 2-ExpTime by either [1] or [17], as mentioned in the Introduction. We note that a direct algorithm checking satisfiability of an UNFO+EQ formula can be obtained using techniques presented in Chapter 3. For a similar approach one may see the algorithm presented in the proof of Thm. 5.5. □

Chapter 5

Small model theorem for UNFO+TR

In this chapter we show the following small model property.

Theorem 5.1. *Every finitely satisfiable UNFO+TR formula φ has a finite model of size bounded triply exponentially in $|\varphi|$.*

Let us fix a finitely satisfiable normal form UNFO+TR formula φ over a signature $\sigma_{\text{base}} \cup \sigma_{\text{dist}}$ for $\sigma_{\text{dist}} = \{T_1, \ldots, T_k\}$. Recall that we consider structures that, additionally, appropriately interpret the auxiliary symbols T_u^{-1} and E_u from σ_{aux}. Recall also our assumptions on the $\overline{T}_u$. Denote $\mathcal{E} = \{E_1, \ldots, E_k\}$. Fix a regular tree-like model $\mathfrak{A} \models \varphi$, with linearly bounded degree, doubly exponentially bounded transitive paths (in this chapter we denote this bound by $\mathbf{M}_\varphi$) and doubly exponentially many non-isomorphic subtrees, as guaranteed by subsequent application of Lemmas 3.5 and 3.8.

We show how to build a 'small' finite model $\mathfrak{A}' \models \varphi$. In our construction we inductively produce fragments of $\mathfrak{A}'$ in which some of the E_u are total. The induction is over the number of the non-total E_u. Intuitively, if a relation is total then it plays no important role, so we may forget about it during the construction. On the l-th level of induction we produce a substructure for every isomorphism type of a subtree of $\mathfrak{A}$ and any combination of l non-total equivalences, so that its each element has provided its *partial φ-witness structure*. This partial φ-witness structure is an isomorphic copy of the restriction of some φ-witness structure from $\mathfrak{A}$ to the equivalence class of the intersection of the current total equivalences. Every substructure from the l-th level of induction is constructed by an appropriate arrangement of some number of basic building blocks, called *components*. Each of the components is obtained by some number of applications of the inductive assumption to situations in that one of the non-total equivalences is replaced by a total one.

As in the previous chapter we introduce further fresh (non-transitive) binary symbols W^i whose purpose is to relate elements to their witnesses. We number the elements of the φ-witness structures in $\mathfrak{A}$ arbitrarily (recall that each element is a member of its own φ-witness structure) and interpret W^i in $\mathfrak{A}$ so that for each

$a, b \in A$, $\mathfrak{A} \models W^i ab$ iff b is the i-th element of the φ-witness structure for a (from now, for short, we refer to the element b satisfying $W^i ab$ as the i-th witness for a). Do this in such a way that if two subtrees of $\mathfrak{A}$ were isomorphic before interpreting the W^i then they still are after such expansion. Now, if we mark b as the i-th witness for a during the construction (that is set $\mathfrak{A}' \models W^i ab$), then for any homomorphism $\mathfrak{h}$ we have $\mathfrak{A} \models W^i \mathfrak{h}(a)\mathfrak{h}(b)$.

Let us formally state our inductive lemma. In this statement we do not explicitly include any bound on the size of promised finite models, but such a bound will be implicit in the proof and will be presented later.

Lemma 5.2. *Let $a_0 \in A$ and $\mathcal{E}_0 \subseteq \mathcal{E}$. Define $\mathcal{E}_{tot} = \mathcal{E} \setminus \mathcal{E}_0$. Let E^* be a new binary relation symbol, interpreted as the intersection of all relations from $\mathcal{E}_{tot}$ (empty $\mathcal{E}_{tot}$ yields total E^*), and let $\mathfrak{A}_0 = \mathfrak{A}_{a_0}\restriction[a_0]_{E^*}$. Then there exist a finite structure $\mathfrak{A}_0'$, a function $\mathfrak{p} : A_0' \to A_0$ and an element $a_0' \in A_0'$, called the* origin *of $\mathfrak{A}_0'$, such that*

- *(t1) E^* is total on $\mathfrak{A}_0'$.*
- *(t2) $\mathfrak{p}(a_0') = a_0$.*
- *(t3) For each $a' \in A_0'$ and each i, if the i-th witness for $\mathfrak{p}(a')$ lies in A_0 (that is $\mathfrak{A}_0 \models \exists y\, W^i \mathfrak{p}(a')y$) then there exists a unique element $b' \in A_0'$ such that $\mathfrak{A}_0' \models W^i a'b'$. Otherwise there exists no such element. Denote $W_{a'} = \{b' : \exists i\ \mathfrak{A}_0' \models W^i a'b'\}$ and for a tuple $\bar{a}$ let $W_{\bar{a}} = \bigcup_{a \in \bar{a}} W_a$.*
- *(t4) For each $\bar{a} \subseteq A_0'$ satisfying $|\bar{a}| \leq t$ there exists a homomorphism $\mathfrak{h} : \mathfrak{W}_{\bar{a}} \to \mathfrak{A}_0$ such that for each $a \in \bar{a}$ we have $\mathfrak{A}_{\mathfrak{p}(a)} \cong \mathfrak{A}_{\mathfrak{h}(a)}$ and $\mathfrak{h}\restriction W_a$ is an isomorphism (onto its image). Moreover, if $a_0' \in \bar{a}$ then we can choose $\mathfrak{h}$ so that $\mathfrak{h}(a_0') = a_0$.*
- *(t5) For each $a \in A_0'$ we have $\mathfrak{W}_a \cong \mathfrak{W}\restriction A_0$ where $\mathfrak{W}$ is the φ-witness structure for $\mathfrak{p}(a)$. (Note that, by the definition of the W^i, each such isomorphism sends a to $\mathfrak{p}(a)$.)*

Before we prove Lemma 5.2 let us observe that it indeed allows us to build a particular finite model of φ. Apply Lemma 5.2 to $\mathcal{E}_0 = \mathcal{E}$ (which means that $\mathcal{E}_{tot} = \emptyset$ and E^* is total) and a_0 being the root of $\mathfrak{A}$ (which means that $\mathfrak{A}_0 = \mathfrak{A}$). We use Lemma 2.2 to see that the obtained structure $\mathfrak{A}_0'$ is a model of φ. Indeed, Condition (h1) of Lemma 2.2 follows directly from Condition (t5), as the structures $\mathfrak{W}_a$ from (t5) are full φ-witness structures in this case, Condition (h2) is implied by Condition (t4) (since $\bar{a} \subseteq \mathfrak{W}_{\bar{a}}$ and $\mathfrak{h}\restriction W_a$ is an isomorphism, $\mathfrak{h}\restriction\bar{a}$ preserves 1-types).

The proof of Lemma 5.2 goes by induction on $l = |\mathcal{E}_0|$. In the base of induction, $l = 0$, we have $\mathcal{E}_{tot} = \mathcal{E}$. Without loss of generality we may assume that E^*-classes in $\mathfrak{A}$ have cardinality 1. If this is not the case, we simply add an artificial transitive relation T_{k+1} and interpret it as the identity in $\mathfrak{A}$, which means that also E_{k+1} is the identity. We simply take $\mathfrak{A}_0' := \mathfrak{A}_0 = \mathfrak{A}\restriction\{a_0\}$ and set $\mathfrak{p}(a_0) = a_0$. It is readily verified that the conditions (t1)–(t5) are then satisfied.

For the inductive step assume that Lemma 5.2 holds for arbitrary $\mathcal{E}_0$ of size $l - 1$.

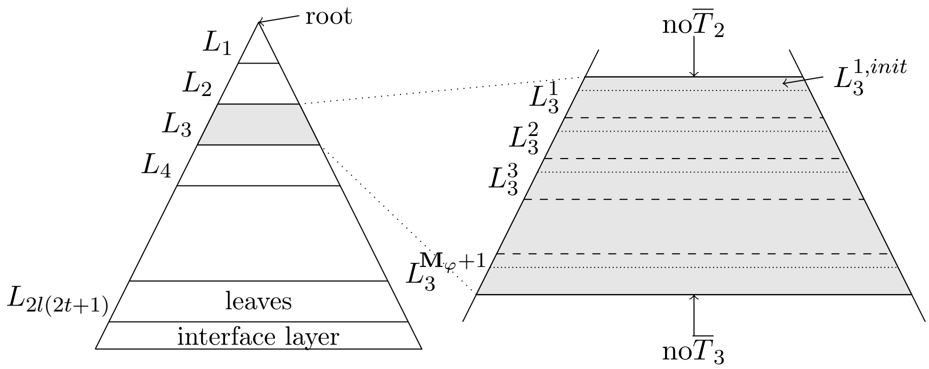

Figure 5.1: A schematic view of a component

We show that then it holds for $\mathcal{E}_0$ of size l. Without loss of generality we assume that $\mathcal{E}_0 = \{E_1, \ldots, E_l\}$. Recall that these are the equivalences induced by $T_1, \ldots, T_l$ and that $\overline{T}_1, \ldots, \overline{T}_{2l}$ is then a list containing $T_1, \ldots, T_l$ and their inverses T_u^{-1}. In the next two sections we present a construction of $\mathfrak{A}'_0$ and then, in the following section, we argue that it is correct. Finally we estimate the size of the produced models and establish the complexity of the finite satisfiability problem.

5.1. Pattern components

We plan to construct $\mathfrak{A}'_0$ out of basic building blocks called *components*. Each component will be an isomorphic copy of some *pattern component*. Let $\boldsymbol{\gamma}[A_0]$ be the set of isomorphism types of subtrees of $\mathfrak{A}$ rooted at A_0. We construct a pattern component $\mathfrak{C}^\gamma$ for every $\gamma \in \boldsymbol{\gamma}[A_0]$. If $\gamma = \gamma_{a_0}$ is the type of $\mathfrak{A}_{a_0}$ then take $a = a_0$; otherwise take any element $a \in A_0$ that is the root of a subtree of type γ.

$\mathfrak{C}^\gamma$ is a finite structure whose universe is divided into $2l(2t+1)$ *inner layers*, denoted $L_1, \ldots, L_{2l(2t+1)}$ and a single *interface layer*, denoted $L_{2l(2t+1)+1}$. Each inner layer L_i is further divided into *sublayers* $L_i^1, L_i^2, \ldots, L_i^{\mathbf{M}_\varphi+1}$. Additionally, in each sublayer L_i^j its initial part $L_i^{j,init}$ is distinguished. In particular, $L_1^{1,init}$ consists of a single element called the *root*. The interface layer $L_{2l(2t+1)+1}$ has no internal division but, for convenience, is sometimes referred to as $L_{2l(2t+1)+1}^{1,init}$. The elements of $L_{2l(2t+1)}$ are called *leaves* and the elements of $L_{2l(2t+1)+1}$ are called *interface elements*. See Fig. 5.1.

$\mathfrak{C}^\gamma$ will have a shape resembling a tree, with structures obtained by the inductive assumption as nodes, though it will not be tree-like in the sense of Chapter 3 (in particular, the internal structure of nodes may be very complicated). All elements of $\mathfrak{C}^\gamma$, except for the interface elements, will have appropriate partial φ-witness structures provided.

We remark that during the process of building a pattern component we do not yet apply the transitive closure to the distinguished and auxiliary relations. Postponing this step is not important from the point of view of the correctness of the construction, but will allow us for a simpler presentation of the proof of its correctness. Given a pattern component $\mathfrak{C}$ we will sometimes denote by $\mathfrak{C}_+$ the structure obtained from $\mathfrak{C}$ by applying all the appropriate transitive closures.

The crucial property we want to enforce is that the root of $\mathfrak{C}^\gamma$ will be *far* from its leaves in the following sense. Denote by $G_l(\mathfrak{S})$, for a σ-structure $\mathfrak{S}$, the Gaifman graph of the structure obtained by removing from $\mathfrak{S}$ the equivalences $E_{l+1}, \ldots, E_k$, the corresponding transitive relations and their inverses. Then there is no connected induced subgraph of $G_l(\mathfrak{C}^\gamma_+)$ of size t containing an element of one of the first $2l$ layers and, simultaneously, an element of one of the last $2l$ inner layers of $\mathfrak{C}^\gamma$.

The role of every inner layer L_i is, speaking informally, to *kill* one of the $\overline{T}_u$, that is to cause that there will be no $\overline{T}_u$-connections from L_i to L_{i+1}. See the right part of Fig. 5.1. The role of sublayers, on the other hand, is to decrease the $\overline{T}_u$-rank of elements. The purpose of the interface layer, $L_{2l(2t+1)+1}$, will be revealed later.

We begin the construction of $\mathfrak{C}^\gamma$ by defining $L_1^{1,init} = \{a'\}$ for a fresh a', setting $\mathrm{atp}^{\mathfrak{C}^\gamma}(a') = \mathrm{atp}^{\mathfrak{A}}(a)$ and $\mathfrak{p}(a') = a$.

Construction of an inner layer: Let $1 \leq i \leq 2l(2t+1)$. Assume we have defined layers $L_1, \ldots, L_{i-1}$, the initial part of sublayer L_i^1, $L_i^{1,init}$, and both the structure of $\mathfrak{C}^\gamma$ and the values of $\mathfrak{p}$ on $L_1 \cup \ldots \cup L_{i-1} \cup L_i^{1,init}$. Let $v = 1 + (i - 1 \bmod 2l)$, and let $w = \lfloor (v+1)/2 \rfloor$. We are going to kill $\overline{T}_v$. Note that E_w is the equivalence corresponding to $\overline{T}_v$. We now expand $L_i^{1,init}$ to full layer L_i.

Step 1: Subcomponents. Assume that we have defined sublayers $L_i^1, \ldots, L_i^{j,init}$, and both the structure of $\mathfrak{C}^\gamma$ and the values of $\mathfrak{p}$ on $L_1 \cup \ldots \cup L_{i-1} \cup L_i^1 \cup \ldots \cup L_i^{j,init}$. For each $b \in L_i^{j,init}$ perform independently the following procedure. Apply the inductive assumption to $\mathfrak{p}(b)$ and the set $\mathcal{E}_0 \setminus E_w$ obtaining a structure $\mathfrak{B}_0$, its origin b_0 and a function $\mathfrak{p}_b : B_0 \to A_{\mathfrak{p}(b)} \cap [\mathfrak{p}(b)]_{E^* \cap E_w} \subseteq A_0$ with $\mathfrak{p}_b(b_0) = \mathfrak{p}(b)$. Identify b_0 with b and add the remaining elements of $\mathfrak{B}_0$ to L_i^j, retaining the structure. Substructures $\mathfrak{B}_0$ of this kind will be called *subcomponents* (note that all appropriate relations are transitively closed in subcomponents). Extend $\mathfrak{p}$ so that $\mathfrak{p}\restriction B_0 = \mathfrak{p}_b$. This finishes the definition of L_i^j.

Step 2: Providing witnesses. For each $b \in L_i^j$ independently perform the following procedure. Let $\mathfrak{B}_0$ be the subcomponent created inductively in *Step 1*, such that $b \in B_0$. Let $\mathfrak{W}$ be the φ-witness structure for $\mathfrak{p}(b)$ in $\mathfrak{A}$. Let $\mathfrak{E} = \mathfrak{W}\restriction[\mathfrak{p}(b)]_{E^*}$ and $\mathfrak{F} = \mathfrak{W}\restriction[\mathfrak{p}(b)]_{E^* \cap E_w}$. Note that $\mathfrak{F}$ is a substructure of $\mathfrak{E}$. By (t5) b has the partial φ-witness structure $\mathfrak{F}'$, isomorphic to $\mathfrak{F}$, provided in $\mathfrak{B}_0$. Extend $\mathfrak{F}'$ in $\mathfrak{C}^\gamma$ to an isomorphic copy $\mathfrak{E}'$ of $\mathfrak{E}$. The structure $\mathfrak{E}'$ will be the structure $\mathfrak{W}_b$ in $\mathfrak{C}^\gamma$ and then in $\mathfrak{A}'_0$. The elements of $E' \setminus F'$ are fresh, and are assigned their sublayers as follows. For $c \in E' \setminus F'$ if $\mathfrak{E}' \models \overline{T}_v bc$ (observe that in this case $\mathfrak{E}' \models \neg \overline{T}_v cb$) then add c to $L_i^{j+1,init}$, otherwise add c to $L_{i+1}^{1,init}$. See Fig. 5.2. Take as the values of $\mathfrak{p}\restriction(E' \setminus F')$ the corresponding elements of $E \setminus F$.

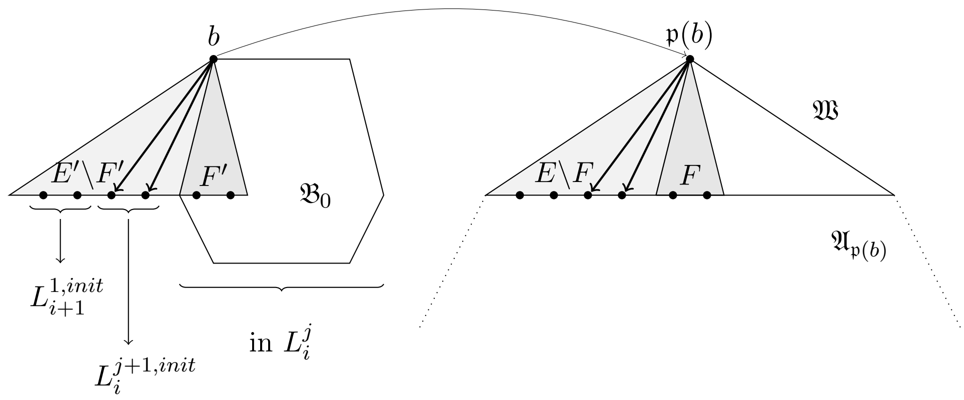

Figure 5.2: Providing witnesses. Thick arrows denote $\overline{T}_u$-connections.

An attentive reader may be afraid that when adding witnesses for elements of the last sublayer $L_i^{\mathbf{M}_\varphi+1}$ of L_i we may want to add one of them to the non-existing layer $L_i^{\mathbf{M}_\varphi+2}$. There is however no such danger, which follows from the following claim.

Claim 5.3. *(i) Let* $b \in L_i^{j,init}$ *and let* $\mathfrak{B}_0$ *be the subcomponent created for* b *in* Step 1. *Then for all* $b' \in \mathfrak{B}_0$ *we have* $\mathfrak{r}_v(\mathfrak{p}(b)) \geq \mathfrak{r}_v(\mathfrak{p}(b'))$. *(ii) Let* $b \in L_i^j$ *and let* $\mathfrak{E}', \mathfrak{F}'$ *be the partial* φ*-witness structures for* b *considered in* Step 2. *Then for any* $c \in E' \setminus F'$ *such that* $\mathfrak{E}' \models \overline{T}_v bc$ *(so* $c \in L_i^{j+1}$*) the inequality* $\mathfrak{r}_v(\mathfrak{p}(b)) > \mathfrak{r}_v(\mathfrak{p}(c))$ *holds.*

Proof. (i) By the inductive assumption (t1) for $\mathfrak{B}_0$, we have $\mathfrak{B}_0 \models E_w bb'$. By (t4) there exists a homomorphism $\mathfrak{h} : \{b, b'\} \to \mathfrak{A}_{\mathfrak{p}(b)}$ and thus $\mathfrak{A}_{\mathfrak{p}(b)} \models E_w \mathfrak{h}(b)\mathfrak{h}(b')$. Since b is the origin of $\mathfrak{B}_0$, we may assume that $\mathfrak{h}(b) = \mathfrak{p}(b)$. This means that $\mathfrak{r}_v(\mathfrak{p}(b)) \geq \mathfrak{r}_v(\mathfrak{h}(b')) = \mathfrak{r}_v(\mathfrak{p}(b'))$ (the last equality follows from the fact that $\mathfrak{A}_{\mathfrak{h}(b')} \cong \mathfrak{A}_{\mathfrak{p}(b')}$ and that $\overline{T}_v$-rank is clearly preserved by an isomorphism).

(ii) By the choice of $\mathfrak{E}'$ and $\mathfrak{F}'$ we have $\mathfrak{E}' \models \overline{T}_v bc \wedge \neg \overline{T}_v cb$, thus by the choice of $\mathfrak{p}$ $\mathfrak{A} \models \overline{T}_v \mathfrak{p}(b)\mathfrak{p}(c) \wedge \neg \overline{T}_v \mathfrak{p}(c)\mathfrak{p}(b)$ and finally $\mathfrak{r}_v(\mathfrak{p}(b)) > \mathfrak{r}_v(\mathfrak{p}(c))$. □

Hence, when moving from L_i^j to L_i^{j+1} the $\overline{T}_v$-ranks of pattern elements for the elements of these sublayers strictly decrease. Since these ranks are bounded by $\mathbf{M}_\varphi$, then, even if the $\overline{T}_v$-ranks of the patterns of some elements of L_i^1 are equal to $\mathbf{M}_\varphi$, then, if $L_i^{\mathbf{M}_\varphi+1}$ is non-empty, the $\overline{T}_v$-ranks of the patterns of its elements must be 0, which means that they cannot have witnesses connected to them one-directionally by $\overline{T}_v$.

The construction of $\mathfrak{C}^\gamma$ is finished when the interface layer, $L_{2l(2t+1)+1}$ is defined (recall that it has only its 'initial part').

5.2. Joining the components

In this chapter we take some number of copies of pattern components and arrange them into the desired structure $\mathfrak{A}'_0$, identifying interface elements of some components with the roots of some other. Some care is needed in this process in order to avoid any modifications of the internal structure of closures $\mathfrak{C}_+$ of components $\mathfrak{C}$, which could potentially result from transitivity of relations. In particular we need to ensure that if for some u a pair of elements of a component $\mathfrak{C}$ is not connected by $\overline{T}_u$ inside $\mathfrak{C}$, then it will not become connected by a chain of $\overline{T}_u$-edges external to $\mathfrak{C}$.

We create a pattern component $\mathfrak{C}^\gamma$ for every $\gamma \in \boldsymbol{\gamma}[A_0]$. Let max be the maximal number of interface elements across all the $\mathfrak{C}^\gamma$. For each $\mathfrak{C}^\gamma$ we number its interface elements.

For each $\gamma \in \boldsymbol{\gamma}[A_0]$ we take copies $\mathfrak{C}^{\gamma;g}_{i,\gamma'}$ of $\mathfrak{C}^\gamma$ for $g \in \{0,1\}$ (g is often called a *color*), $1 \le i \le max$ and $\gamma' \in \boldsymbol{\gamma}[A_0]$. We also take an additional copy $\mathfrak{C}^{\gamma_{a_0};0}_{\perp,\perp}$ of $\mathfrak{C}^{\gamma_{a_0}}$. Its root will become the origin of the whole $\mathfrak{A}'_0$.

For each γ, g consider components of the form $\mathfrak{C}^{\gamma;g}_{\cdot,\cdot}$. Perform the following procedure for each i—the number of an interface element. Let b be the i-th interface element of any such component, let γ' be the type of $\mathfrak{A}_{\mathfrak{p}(b)}$. Identify the i-th interface elements of all $\mathfrak{C}^{\gamma;g}_{\cdot,\cdot}$ with the root c_0 of $\mathfrak{C}^{\gamma',1-g}_{i,\gamma}$. See Fig. 5.3

Note that the values of $\mathfrak{p}(c_0)$ and $\mathfrak{p}(b)$ (the latter equals to the value of $\mathfrak{p}$ on the i-th interface element in all the $\mathfrak{C}^{\gamma;g}_{\cdot,\cdot}$) may differ. However, by construction, $\mathfrak{A}_{\mathfrak{p}(b)} \cong \mathfrak{A}_{\mathfrak{p}(c_0)}$ (in particular, the 1-types of b and c_0 match). For the element c^* obtained in this identification step we define $\mathfrak{p}(c^*) = \mathfrak{p}(c_0)$.

Define the graph of components used in the above construction, G^{comp}, by joining two components by an edge iff we identified an interface element of one with the root of the other. Take $\mathfrak{A}^0_0$ as the structure restricted to the components accessible from $\mathfrak{C}^{\gamma_{a_0};0}_{\perp,\perp}$ in G^{comp}. Note that in $\mathfrak{A}^0_0$ we still do not take the transitive closures of relations. We define $\mathfrak{A}'_0$ by transitively closing all appropriate relations in $\mathfrak{A}^0_0$. Later we will keep using the convention of marking some auxiliary structures in which the transitive closures are not yet applied with the superscript 0. Finally, we choose as the origin a'_0 of $\mathfrak{A}'_0$ the root of the pattern component $\mathfrak{C}^{\gamma_0;0}_{\perp,\perp}$.

5.3. Correctness of the construction

(t1) By the construction, after taking the transitive closures, E^* is total on each pattern component. Next, by the definition of the graph of components G^{comp}, E^* is total on $\mathfrak{A}'_0$.

(t2) As a'_0 we take the root of $\mathfrak{C}^{\gamma_{a_0};0}_{\perp,\perp}$. Recall that we explicitly map the root of the pattern component $\mathfrak{C}^{\gamma_{a_0}}$ by $\mathfrak{p}$ to a_0.

(t3) The interpretations of the W^i are defined in the step of providing witnesses where, implicitly, we take care of this condition for every element a' of the inner layers by extending the fragment of the partial φ-witness structure for a' created on the previous level of induction by a copy of a further fragment of the same pattern φ-witness structure. Note also that during the step of joining the components all the interface elements become identified with some roots, which are elements of inner layers, and that the identifications do not spoil the required property.

(t4) For simplicity, let us ignore the 'moreover' part of this condition for some time. We will explain how to take care of it near the end of this proof. Now we find a homomorphism $\mathfrak{h}$ such that $\mathfrak{A}_{\mathfrak{p}(a)} \cong \mathfrak{A}_{\mathfrak{h}(a)}$ for all $a \in \bar{a}$ (we say that such a homomorphism has *the subtree isomorphism property.* Later we will show that its restrictions to the substructures $\mathfrak{W}_a$ are indeed isomorphisms. The proof starts with several homomorphic reductions which show that instead of $\mathfrak{A}'_0$ we can consider a structure looking like a pattern component but twice as high.

Reduction 0. Consider a tuple $\bar{a} \subseteq A'_0$ such that $|\bar{a}| \leq t$. Observe that for each $a \in \bar{a}$ the structure $\mathfrak{W}_a$ is connected in $G_l(\mathfrak{A}'_0 \restriction W_{\bar{a}})$ (recall the definition of Gaifman graph $G_l(\mathfrak{S})$ and the interpretation of the symbols W^i). Let $\mathfrak{W}_{\bar{a}_1}, \ldots, \mathfrak{W}_{\bar{a}_K}$ be the connected components of $\mathfrak{W}_{\bar{a}}$ in $G_l(\mathfrak{A}'_0 \restriction W_{\bar{a}})$. If we have homomorphisms $\mathfrak{h}_i : \mathfrak{W}_{\bar{a}_i} \to \mathfrak{A}_0$ satisfying the subtree isomorphism property then we can take $\mathfrak{h} = \bigcup \mathfrak{h}_i : \mathfrak{W}_{\bar{a}} \to \mathfrak{A}_0$ which is a homomorphism, since E^* is total on $\mathfrak{A}_0$, that still has the subtree isomorphism property. Owing to this reduction *we can restrict attention to tuples $\bar{a}$ with $\mathfrak{W}_{\bar{a}}$ connected (in the above sense)* .

Reduction 1. By the construction, for all $1 \leq i \leq 2l(2t+1)$ and $v = 1+(i-1 \bmod 2l)$, there is no $\overline{T}_v$-path in any component from an element of L_i to an element of L_{i+1}. Thus, if we divide (inner) layers of components into groups of size $2l$, a transitive path may join at most elements of two neighboring groups. Obviously, non-transitive relations join only tuples consisting of elements of at most two consecutive layers, and, in particular, each of the $\mathfrak{W}_a$ lies in at most two consecutive layers. It follows, that given a connected $\mathfrak{W}_{\bar{a}}$, $|\bar{a}| \leq t$, by our choice of the number of layers in a component, there exists $g \in \{0, 1\}$ such that removing all the connections between leaves of color $1 - g$ and roots of color g (in other words: any connections between elements of $L_{2l(2t+1)}$ and elements of $L_{2l(2t+1)+1}$ in components of color $1 - g$) does not remove any connections among the elements of $\mathfrak{W}_{\bar{a}}$.

More formally, let $\mathfrak{D}^0_0$ be the structure obtained from $\mathfrak{A}^0_0$ by removing all the connections as described above, and let $\mathfrak{D}'_0$ be the transitive closure of $\mathfrak{D}^0_0$. Then the inclusion map $\iota : \mathfrak{W}_{\bar{a}} \to \mathfrak{D}'_0$ is a homomorphism. Clearly $\mathfrak{A}_{\mathfrak{p}(\iota(a))} = \mathfrak{A}_{\mathfrak{p}(a)}$ since $a = \iota(a)$. Thus *we can restrict attention to a tuple $\bar{a}$ for which $\mathfrak{W}_{\bar{a}}$ is connected and search for a homomorphism $\mathfrak{W}_{\bar{a}} \to \mathfrak{A}_0$ treating $\mathfrak{W}_{\bar{a}}$ as a substructure of $\mathfrak{D}'_0$.*

Reduction 2. Observe that by our scheme of arranging the copies of pattern components there is at most one type γ of components of color g (where g is the color from

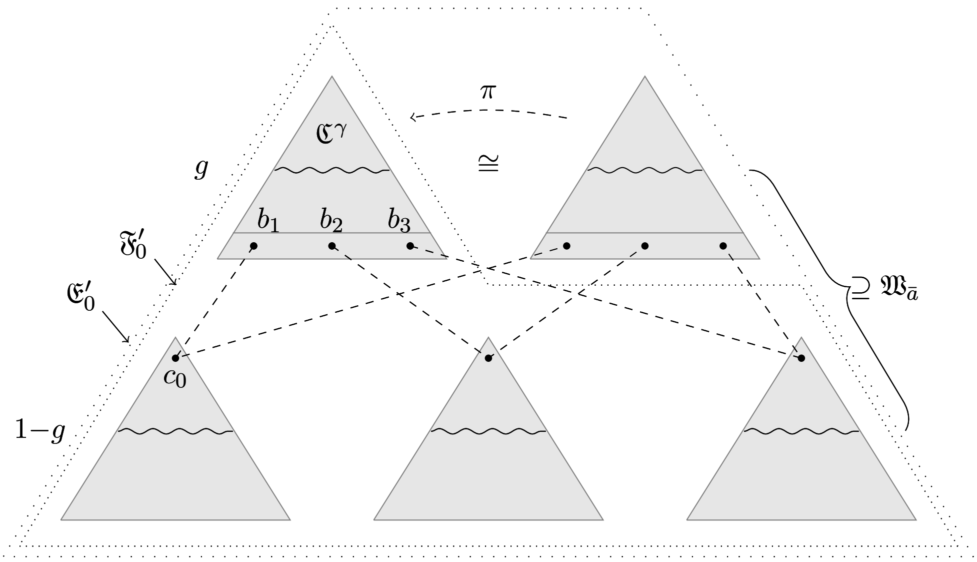

Figure 5.3: Joining the components and Reductions 1 and 2. Elements connected by dashed lines are identified.

the previous reduction) that contains some element of a connected $\mathfrak{W}_{\bar{a}}$ (consider the shape of a connected fragment of the graph of components G^{comp} with connections between leaves of color g and roots of color $1-g$ removed). Furthermore, all elements of $\mathfrak{W}_{\bar{a}}$ of color $1-g$ are contained in components of the form $\mathfrak{C}^{\cdot;1-g}_{\cdot;\gamma}$. Choose one component of type γ of color g and call it $\mathfrak{C}^\gamma$. Consider the structure $\mathfrak{E}_0^0$ (resp. $\mathfrak{F}_0^0$) obtained as the restriction of the structure $\mathfrak{D}_0^0$ from the previous reduction to the union of the domains of the components of the form $\mathfrak{C}^{\gamma;g}_{\cdot;\cdot}$ (resp. the domain of $\mathfrak{C}^\gamma$) and the domains of all the components of the form $\mathfrak{C}^{\cdot;1-g}_{\cdot;\gamma}$. Let $\mathfrak{E}_0'$ (resp. $\mathfrak{F}_0'$) be their transitive closures. Consider a projection π that projects all elements of $\mathfrak{E}_0^0$ of color g onto $\mathfrak{C}^\gamma$ and is the identity on the others. We claim that $\pi{\restriction}W_{\bar{a}} : \mathfrak{W}_{\bar{a}} \to \mathfrak{F}_0'$ is a homomorphism. To see this, observe that the paths connecting elements of $\mathfrak{W}_{\bar{a}}$ in $\mathfrak{D}_0^0$ are contained in E_0^0 and $\pi : \mathfrak{E}_0^0 \to \mathfrak{F}_0^0$ is a homomorphism. See Fig. 5.3. Clearly for each $a \in \bar{a}$ we have $\mathfrak{A}_{\mathfrak{p}(a)} = \mathfrak{A}_{\mathfrak{p}(\pi(a))}$ since $\mathfrak{p}(a) = \mathfrak{p}(\pi(a))$. Thus, finally, *we can restrict attention to a tuple $\bar{a}$ for which $\mathfrak{W}_{\bar{a}}$ is connected and search for a homomorphism $\mathfrak{W}_{\bar{a}} \to \mathfrak{A}_0$ treating $\mathfrak{W}_{\bar{a}}$ as a substructure of $\mathfrak{F}_0'$.*

Essential homomorphism construction. Note that $\mathfrak{F}_0^0$ looks like a single component but is twice as high. Consider the *tree of subcomponents* of $\mathfrak{F}_0^0$, τ, defined as follows: make a subcomponent $\mathfrak{B}$ the parent of $\mathfrak{B}'$ if $\mathfrak{B}'$ contains a witness for an element of $\mathfrak{B}$. Observe that so obtained τ is indeed a tree. For a subcomponent $\mathfrak{B} \in \tau$ denote by $B^\wedge$ the union of the domains of all the subcomponents belonging to the subtree of τ rooted at $\mathfrak{B}$.

Since we might have cut some connections between an element and some of its witnesses during Reduction 1, we define for each $a \in F_0'$ the surviving part $\mathfrak{V}_a$ of $\mathfrak{W}_a$ by $\mathfrak{V}_a = \mathfrak{F}_0'{\restriction}V_a$ where $V_a = \{b : \exists i\ \mathfrak{F}_0' \models W^i ab\}$. For a tuple $\bar{b}$ denote $V_{\bar{b}} = \bigcup_{b \in \bar{b}} V_b$ and $\mathfrak{V}_{\bar{b}} = \mathfrak{F}_0'{\restriction}V_{\bar{b}}$. Note that $V_a \subseteq W_a$, and generally, this inclusion may be strict, but for all $a \in \bar{a}$ we have $\mathfrak{V}_a = \mathfrak{W}_a$, and thus, in particular, the claim below finishes the proof of the currently considered part of (t4), that is the proof of the existence of a homomorphism satisfying the subtree isomorphism property.

Returning to the shape of $\mathfrak{F}_0^0$, it consists of some subcomponents arranged into tree τ glued together by the structure on the surviving parts. Note that all such building blocks (that is both the subcomponents and the surviving parts of the partial witness structures) are transitively closed. Moreover, by the tree structure of τ, if some elements of such a building block are connected by some atom in $\mathfrak{F}_0'$, then they already have been connected by the same atom in $\mathfrak{F}_0^0$, therefore the identity map from $\mathfrak{F}_0^0$ to $\mathfrak{F}_0'$ acts as an isomorphism when restricted to such a building block.

Claim 5.4. *For every subcomponent $\mathfrak{B}_0 \in \tau$ with origin b_0, and $\bar{a} \subseteq B_0^\wedge$, $|\bar{a}| \leq t$, there exists a homomorphism $\mathfrak{h} : \mathfrak{B}_{\bar{a}} \to \mathfrak{A}_{\mathfrak{p}(b_0)} \restriction [\mathfrak{p}(b_0)]_{E^*}$ such that for all $a \in \bar{a}$ we have $\mathfrak{A}_{\mathfrak{h}(a)} \cong \mathfrak{A}_{\mathfrak{p}(a)}$, and if $b_0 \in \bar{a}$ then $\mathfrak{h}(b_0) = \mathfrak{p}(b_0)$.*

Proof. Bottom-up induction over tree.

Induction base: If $\mathfrak{B}_0$ is a leaf of τ then $\mathfrak{B}_{\bar{a}} \subseteq \mathfrak{B}_0$ and the claim follows by the inductive assumption of Lemma 5.2 (note that here we implicitly use the fact that the identity map is an isomorphism between $\mathfrak{F}_0^0 \restriction B_0$ and $\mathfrak{F}_0' \restriction B_0$).

Induction step: Let $\mathfrak{B}_1, \ldots, \mathfrak{B}_K$ be the list of all the children of $\mathfrak{B}_0$ in τ such that $B_i^\wedge$ contains some element of $\bar{a}$. If $K = 1$ and $\bar{a} \subseteq B_1^\wedge$ the thesis follows from the inductive assumption of this claim.

Otherwise, for $1 \leq i \leq K$ (note that it is possible that $K = 0$), denote by b_i the origin of $\mathfrak{B}_i$ and let $c_i \in \mathfrak{B}_0$ be such that b_i is a witness chosen by c_i in the step of providing witnesses or during the step of joining the components. By the inductive assumption of this claim there exist homomorphisms $\mathfrak{h}_i : \mathfrak{B}_{(\bar{a} \cap B_i^\wedge) b_i} \to \mathfrak{A}_{\mathfrak{p}(b_i)}$ such that $\mathfrak{h}_i(b_i) = \mathfrak{p}(b_i)$. From the inductive assumption of Lemma 5.2 we have a homomorphism $\mathfrak{h}_0 : \mathfrak{B}_{(\bar{a} \cap B_0) c_1 \ldots c_K} \restriction B_0 \to \mathfrak{A}_{\mathfrak{p}(b_0)} \restriction [\mathfrak{p}(b_0)]_{E^*}$. We extend it in the only possible way to $\mathfrak{h}_0^*$ defined on the whole $\mathfrak{B}_{(\bar{a} \cap B_0) c_1 \ldots c_K}$: for each $a \in \bar{a}$ and $c \in V_a \setminus B_0$ (by construction $\mathfrak{B}_a \models W^i ac$ for some i) we set $\mathfrak{h}(c)$ to be the only element satisfying $\mathfrak{A}_0 \models W^i \mathfrak{h}(a) \mathfrak{h}(c)$ (such an element exists since $\mathfrak{A}_{\mathfrak{h}(a)} \cong \mathfrak{A}_{\mathfrak{p}(a)}$—in particular the φ-witness structures of $\mathfrak{h}(a)$ and $\mathfrak{p}(a)$ are isomorphic). Note that the sizes of the tuples used to build the homomorphisms $\mathfrak{h}_i$ are bounded by t, as required.

We construct from the above maps a homomorphism $\mathfrak{h} : \mathfrak{B}_{\bar{a} b_1 \ldots b_K c_1 \ldots c_K} \to \mathfrak{A}_{\mathfrak{p}(b_0)}$. See Fig. 4.4. Here, the crucial property is that $\mathfrak{A}$ has a regular shape. Indeed, for each i, $1 \leq i \leq K$, for the witness b_i' for $\mathfrak{h}_0(c_i)$ (that is an element b_i' satisfying $\mathfrak{A}_0 \models W^j \mathfrak{h}_0(c_i) b_i'$ for the appropriate j), corresponding to the witness b_i for c_i (such that $\mathfrak{F}_0' \models W^j c_i b_i$), we have $\mathfrak{A}_{b_i'} \cong \mathfrak{A}_{\mathfrak{p}(b_i)}$. This is the case since, by the inductive assumption of Lemma 5.2 we have $\mathfrak{A}_{\mathfrak{p}(c_i)} \cong \mathfrak{A}_{\mathfrak{h}_0(c_i)}$, by construction we have $\mathfrak{A}_{b_i^{\mathfrak{p}}} \cong \mathfrak{A}_{\mathfrak{p}(b_i)}$ (here $b_i^{\mathfrak{p}}$ is the j-th witness of $\mathfrak{p}(c_i)$; note that during the step of providing witnesses we set $\mathfrak{p}(b_i) = b_i^{\mathfrak{p}}$; we need to consider $\cong$ since the value of $\mathfrak{p}(b_i)$ may change due to a possible identification applied in the step of joining the components) and the numbering of witnesses is preserved by subtree isomorphisms. Thus there is a homomorphism $\mathfrak{h}_i^* : \mathfrak{B}_{(\bar{a}\cap B_i^\wedge)b_i} \to \mathfrak{A}_{b_i'}$ with $\mathfrak{h}_i^*(b_i) = b_i'$.

We naturally join $\mathfrak{h}_0^*, \mathfrak{h}_1^*, \ldots, \mathfrak{h}_K^*$ into $\mathfrak{h}$: $\mathfrak{h} = \bigcup \mathfrak{h}_i^*$. Note that such $\mathfrak{h}$ is well defined, even though the value $\mathfrak{h}$ on each of the b_i is defined twice, since b_i belongs to both $\mathrm{Dom}\mathfrak{h}_0^*$ and $\mathrm{Dom}\mathfrak{h}_i^*$ ($\mathfrak{h}$ has been defined on the other elements exactly once). For each $a \in \mathrm{Dom}\mathfrak{h}_i$ ($= \mathrm{Dom}\mathfrak{h}_i^*$, when $i > 0$) we have $\mathfrak{A}_{\mathfrak{h}(a)} = \mathfrak{A}_{\mathfrak{h}_i^*(a)} \cong \mathfrak{A}_{\mathfrak{h}_i(a)} (\cong \mathfrak{A}_{\mathfrak{p}(a)}$, by the inductive assumptions of this claim and Lemma 5.2). Since $\bar{a} \subseteq \mathrm{Dom}\mathfrak{h}_0 \cup \bigcup_{i>0} \mathrm{Dom}\mathfrak{h}_i^*$, we can conclude that for each $a \in \bar{a}$ we have $\mathfrak{A}_{\mathfrak{p}(a)} \cong \mathfrak{A}_{\mathfrak{h}(a)}$.

The fact that $\mathfrak{h}$ is a homomorphism follows from the tree structure of τ. In particular, there cannot be any connections (before taking the transitive closures) between (non-origin) elements of two different $B_i^\wedge$ (for $1 \le i \le K$). The full proof that $\mathfrak{h}$ is a homomorphism is tedious, therefore we show two representative cases that use all the major ideas required. First, consider $a, a' \in V_{\bar{a}b_1\ldots b_K c_1 \ldots c_K}$ such that $a \in B_i^\wedge$, $a' \in B_j^\wedge$ for some i, j such that $c_i \neq c_j$. Assume that $\mathfrak{F}_0' \models T_u aa'$ for some u. We will prove that $\mathfrak{A}_0 \models\ T_u\mathfrak{h}(a)\mathfrak{h}(a')$. By a standard argument, owing to the tree structure of τ (some more care is needed since there may be some connections in the structures $\mathfrak{B}_b$), there exist $d_i \in V_{c_i} \cap B_0$ and $d_j \in V_{c_j} \cap B_0$ such that $\mathfrak{F}_0' \models T_u ab_i$, $\mathfrak{F}_0^0 {\restriction} V_{c_i} \models T_u b_i d_i$, $\mathfrak{B}_0 \models T_u d_i d_j$, $\mathfrak{F}_0^0 {\restriction} V_{c_j} \models T_u d_j b_j$ and $\mathfrak{F}_0' \models T_u b_j a'$ (we assumed that $a, b_i, d_i, d_j, b_j, a'$ are pairwise different; otherwise some parts of such path become trivial). Since $\mathfrak{h}_i^*$, $\mathfrak{h}_0$ and $\mathfrak{h}_j^*$ are homomorphisms, $\mathfrak{A}_0 \models T_u\mathfrak{h}(a)\mathfrak{h}(b_i) \wedge T_u\mathfrak{h}(d_i)\mathfrak{h}(d_j) \wedge T_u\mathfrak{h}(b_j)\mathfrak{h}(a')$. Now we show that $\mathfrak{A}_0 \models T_u\mathfrak{h}(b_i)\mathfrak{h}(d_i)$. By construction of $\mathfrak{B}_a$, there exist indices i_b and i_d such that $\mathfrak{F}_0' \models W^{i_b}c_i b_i \wedge W^{i_d}c_i d_i$ and therefore by the choice of $\mathfrak{h}_0$, $\mathfrak{A}_0 \models\ W^{i_d}\mathfrak{h}(c_i)\mathfrak{h}(d_i)$ and by the choice of the extension of $\mathfrak{h}_0$ to $\mathfrak{h}_0^*$, $\mathfrak{A}_0 \models W^{i_b}\mathfrak{h}(c_i)\mathfrak{h}(b_i)$. Let $b_i^{\mathfrak{p}}$ be the i_b-th witness of $\mathfrak{p}(c_i)$ and $d_i^{\mathfrak{p}}$ be the i_d-th witness of $\mathfrak{p}(c_i)$. By construction, $\mathfrak{A}_0 \models T_u b_i^{\mathfrak{p}} d_i^{\mathfrak{p}}$. But $\mathfrak{A}_{\mathfrak{p}(c_i)} \cong \mathfrak{A}_{\mathfrak{h}(c_i)}$ and by the uniqueness of the numbers of the witnesses, any isomorphism between these subtrees sends $b_i^{\mathfrak{p}}$ to $\mathfrak{h}(b_i)$ and $d_i^{\mathfrak{p}}$ to $\mathfrak{h}(d_i)$, therefore $\mathfrak{A}_0 \models T_u\mathfrak{h}(b_i)\mathfrak{h}(d_i)$. Similarly $\mathfrak{A}_0 \models T_u\mathfrak{h}(d_j)\mathfrak{h}(b_j)$. Joining the pieces together, by transitivity of T_u, $\mathfrak{A}_0 \models T_u\mathfrak{h}(a)\mathfrak{h}(a')$. Secondly, we consider the case when $\mathfrak{F}_0' \models R(\bar{a}')$ for some non-transitive symbol R and $\bar{a}' \subseteq V_{\bar{a}b_1\ldots b_K c_1\ldots c_K}$. By construction, $R(\bar{a}')$ was set either in the process of building some subcomponent or during the step of providing witnesses. Thus $\bar{a}'$ is either contained in B_0 or some of the $B_i^\wedge$ or some of the V_a for some $a \in (\bar{a} \cap B_0)c_1 \ldots c_K$. Now we can prove, using arguments similar to ones used for appropriate parts of the path in the previous case, that $\mathfrak{A}_0 \models R(\mathfrak{h}(\bar{a}'))$

Since by construction $\mathfrak{h}_0 \subseteq \mathfrak{h}$, if $b_0 \in \bar{a}$ then $\mathfrak{h}(b_0) = \mathfrak{h}_0(b_0) (= \mathfrak{p}(b_0)$ by the inductive assumption of Lemma 5.2). To finish the inductive step, we restrict $\mathfrak{h}$ to $V_{\bar{a}}$. □

Now we prove the additional property required for $\mathfrak{h}$ by (t4), that is, that for each $a \in \bar{a}$, $\mathfrak{h}{\restriction}W_a$ is an isomorphism. By the numbering of witnesses, as explained before the statement of this lemma, $\mathfrak{h}$ moves W_a into the part of the witness structure of $\mathfrak{h}(a)$

contained in A_0 and is one-to-one by the uniqueness of the numbers of witnesses in a witness structure. The other way around, we can use a similar argument as in the first case presented in the proof that the map built in Claim 5.4 is a homomorphism. That is, if for some $\bar{a}' \subseteq W_a$ and some (arbitrary) relation R, $\mathfrak{A}_0 \models R(\mathfrak{h}(\bar{a}'))$, then, since $\mathfrak{A}_{\mathfrak{h}(a)} \cong \mathfrak{A}_{\mathfrak{p}(a)}$ and any isomorphism preserves the numbering of witnesses and the structure on W_a was copied from a part of the witness structure for $\mathfrak{p}(a)$ (together with such numbering), $\mathfrak{A}'_0 \models R(\bar{a}')$ and therefore the inverse of $\mathfrak{h}\restriction\mathfrak{W}_a$ is also a homomorphism, so $\mathfrak{h}\restriction\mathfrak{W}_a$ is an isomorphism.

Now we return to the 'moreover' part of (t4). Let us assume that $a'_0 \in \bar{a}$. We will slightly modify the above proof. Reductions 0 and 1 do not move a'_0 and we keep them unchanged. Notice that in Reduction 1 we have that $g = 0$. Now, in Reduction 2 we have that $\gamma = \gamma_{a_0}$ and we choose $\mathfrak{C}^\gamma = \mathfrak{C}^{\gamma,0}_{\perp,\perp}$. This way application of π does not move a'_0. To finish the proof, it is sufficient to see that by Claim 5.4 $\mathfrak{h}(a'_0) = \mathfrak{p}(a'_0) = a_0$.

(t5) Apply (t4) to a tuple consisting of just a to obtain an isomorphism $\mathfrak{h} : \mathfrak{W}_a \to \mathfrak{A}_0\restriction\mathfrak{h}(W_a)$ and then apply an isomorphism between $\mathfrak{A}_{\mathfrak{h}(a)}$ and $\mathfrak{A}_{\mathfrak{p}(a)}$.

5.4. Size of models and complexity

To complete the proof of Thm. 5.1 we need to show an appropriate upper bound on the size of finite models produced by our construction. The following routine estimation shows that $|A'_0|$ is triply exponential in $n = |\varphi|$, regardless of the choice of the initial tree-like model $\mathfrak{A}$. We calculate a bound $\mathbf{C}_l$ on the size of the structure obtained in the proof of Lemma 5.2 for $|\mathcal{E}_0| = l$. We are interested in $\mathbf{C}_{k+1}$, which is the desired bound on the size of $\mathfrak{A}'_0$ (we use $\mathbf{C}_{k+1}$ here, rather than $\mathbf{C}_k$, because we may potentially introduce the auxiliary identity relation in the base step of induction). By the construction any pattern component is a tree of subcomponents consisting of at most $2l(2t+1)(\mathbf{M}_\varphi + 1)$ sublayers. In the first sublayer we have at most $\mathbf{C}_{l-1}$ elements, in the second one—at most $\mathbf{C}_{l-1}n$ subcomponents; this jointly gives $\mathbf{C}^2_{l-1}n$ elements. Iterating, we have at most $\mathbf{C}^i_{l-1}n^{i-1}$ elements in i-th sublayer, which jointly gives an estimate $(\mathbf{C}_{l-1}n)^{2l(2t+1)(\mathbf{M}_\varphi+1)+1}$ on both the number of inner elements and the number of interface elements in a pattern component. Multiplying it by the number of components used in the joining phase, and then estimating t and l in the exponent by n and $n+1$ respectively, we get a bound $\mathbf{C}_l = 2|\boldsymbol{\gamma}[A]|^2(\mathbf{C}_{l-1}n)^{4(n+1)(2n+1)(\mathbf{M}_\varphi+1)+2}$. Solving this recurrence relation, and recalling that $\mathbf{M}_\varphi$ and $|\boldsymbol{\gamma}[A]|$ are doubly exponential in $|\varphi|$ we obtain a triply exponential bound on $\mathbf{C}_{k+1}$.

This finishes the proof of Thm. 5.1. We do not know if our construction is optimal with respect to the size of models. The best we can do for the lower bound is to enforce models of at most doubly exponential size (actually, it can be done in UNFO even without transitive relations).

Thm. 5.1 immediately gives the decidability of the finite satisfiability problem for UNFO+TR and suggests a simple 3-NExpTime-procedure: convert a given formula φ into normal form φ', guess a finite structure of size bounded triply exponentially and verify that it is a model of φ'. We can however do better and show a doubly exponential upper bound matching the known complexity of the general satisfiability problem.

Theorem 5.5. *The finite satisfiability problem for UNFO+TR is 2-ExpTime-complete.*

Proof. The lower bound is inherited from pure UNFO [27]. As a side note, the lower bound can be also obtained for the two-variable UNFO with one transitive relation. The proof is an adaptation of the lower bound proof for GF^2 with transitive relations in guards [20].

For the upper bound, we describe an algorithm in AExpSpace. Fix φ in normal form. We have proved that φ has a finite model iff it has a tree-like model with bounded transitive paths (as in Lemma 3.5). We will look for the latter. We advise the reader to recall the proof of Lemma 3.8, as we presently use a similar apparatus. In our procedure we produce, in an alternating fashion, a finite tree $\mathfrak{A}^*$, corresponding to some number of the upper levels of a model. Simultaneously, we define a function $\mathfrak{g}^*$ returning for an element of A^* its 1-type together with some φ-declaration and one stopwatch for each of the $\overline{T}_u$ (cf. the proof of Lemma 3.8).

More precisely, let $\mathbf{M}_\varphi$ be the bound on transitive paths obtained in Lemma 3.5 and $\mathbf{C}$ be a bound on $|\mathrm{Rng}\mathfrak{g}^*|$ (we use $(\overline{T}_u, \mathbf{M}_\varphi)$-stopwatches in $\mathfrak{g}^*$). The alternating algorithm works as follows. Calculate $\mathbf{M}_\varphi$ and $\mathbf{C}$. Note that both are doubly exponential in $|\varphi|$. Construct the root of $\mathfrak{A}$ and guess its 1-type α, a φ-declaration $\mathfrak{d}$ containing all the formulas of the form $\varphi_0^j(\bar{x}) \wedge \bigwedge_{i \in Q} x_i = y \wedge \bigwedge_{i \in \boldsymbol{Q} \setminus Q} x_i \neq y$ for any $Q \subseteq \boldsymbol{Q}$ and $1 \leq j \leq z$ (recall that φ_0 is equivalent to $\varphi_0^1 \vee \ldots \vee \varphi_0^z$ with the φ_0^j being conjunctions of some $\mathcal{R}$ and $\mathcal{T}$ formulas). Set $\mathfrak{g}^*(a) = (\alpha, \mathfrak{d}, (0)_{u=1}^{2k})$. Now construct the downward family of a, $F = \{a, a_1, \ldots, a_s\}$, for some $s < |\varphi|$, guess its (transitively closed) structure, and guess the values $\mathfrak{g}^*(a_1), \ldots, \mathfrak{g}^*(a_s)$. Check whether F is a φ-witness structure for a, the 1-types assigned by $\mathfrak{g}^*$ agree with the structure, the declarations assigned by $\mathfrak{g}^*$ satisfy the LCCs and the stopwatches assigned by $\mathfrak{g}^*$ satisfy the local condition described in the definition of $(\overline{T}_u, \mathbf{M}_\varphi)$-stopwatch labeling. If not, reject. Next universally choose one of the a_i. Then proceed as for a—guess the downward family of a_i and values of $\mathfrak{g}^*$, and check their consistency as above, universally choose one of the children of a_i and so on. We additionally keep a counter containing the number of the current level in $\mathfrak{A}^*$. If it reaches $\mathbf{C} + 1$, we accept.

It is clear that the described algorithm can be implemented in AExpSpace: we only need to store the structure and the values of $\mathfrak{g}^*$ on a single family, plus a counter. All of these can be written using exponentially many bits.

Correctness proof. To see that if φ has a model $\mathfrak{A}$ with bounded transitive paths then the algorithm accepts, it is sufficient to make the guesses in accordance with $\mathfrak{A}^*$—the structure induced on the first $\mathbf{C}+1$ levels of $\mathfrak{A}$ with $\mathfrak{g}^*$ defined as follows $A^* \ni a \mapsto (\mathrm{atp}^{\mathfrak{A}}(a), \mathrm{dec}^{\mathfrak{A}}_{\varphi}(a), (\mathcal{S}_u)_{u=1}^{2k})$ where $\mathcal{S}_u$ is the $(\overline{T}_u, \mathbf{M}_\varphi)$-stopwatch labeling of $\mathfrak{A}$. The fact that such a strategy leads to an accepting run of the algorithm is almost straightforward. In particular, the local consistency of declarations follows from Lemma 3.3(ii). The opposite implication uses ideas similar to the ones from the proof of Lemma 3.8. Assume that the algorithm has an accepting run. From this run we can naturally infer a tree-like structure $\mathfrak{A}^*$ consisting of $\mathbf{C}$+1 levels, and a function $\mathfrak{g}^*$. Note that on each path from the root to a leaf in $\mathfrak{A}^*$ some value of $\mathfrak{g}^*$ appears at least twice. Cut each branch at the first position on which the value of $\mathfrak{g}^*$ reappears and make a link from this point to the first occurrence of this value on the considered branch. Naturally unravel so obtained structure into an infinite tree-like structure $\mathfrak{A}$. Define on $\mathfrak{A}$ function $\mathfrak{g}$ just copying the values of $\mathfrak{g}^*$. We show that $\mathfrak{A} \models \varphi$ and has transitive paths bounded by $\mathbf{M}_\varphi$. Note that the downward families in $\mathfrak{A}$ and the values of $\mathfrak{g}$ on them are copies of some downward families in $\mathfrak{A}^*$ and their values of $\mathfrak{g}^*$, so each $a \in A$ has a φ-witness structure ($\mathfrak{A}$ satisfies all the $\forall\exists$-conjuncts of φ) and also $\mathfrak{g}$ gives a locally consistent set of declarations and $(\overline{T}_u, \mathbf{M}_\varphi)$-stopwatch labelings. The latter guarantee that $\mathfrak{A}$ has bounded transitive paths; the former, together with the choice of the declaration $\mathfrak{d}$ for the root of $\mathfrak{A}^*$, allows us to conclude that $\mathfrak{A}$ satisfies the $\forall$-conjunct of φ. □

As remarked in the Introduction, we can state our results in a slightly stronger way, for a setting in which we may not only require some binary symbols to be interpreted as arbitrary transitive relations, but we can, more specifically, require some of them to be equivalences and some other—partial orders. Indeed, assuming that T is transitive we can enforce it in UNFO to be a (strict) partial order, writing $\neg\exists xy(Txy \wedge Tyx)$. Non-strict partial orders can be then simulated by disjunctions $Txy \vee x = y$. To have equivalences we can just use the symbols from σ_{aux} (cf. Chapter 2).

Corollary 5.6. *The finite satisfiability problem for UNFO with transitive relations, equivalences and partial orders is 2-*ExpTime*-complete.*

We note that our approach does not allow us to deal with *linear* orders. Actually, the presence of a strict linear order $<$ makes the satisfiability problem for UNFO undecidable, as it allows for a reduction from UNFO with inequalities, which is known to be undecidable [27]: $x \neq y$ can be then expressed as $x < y \vee y < x$. See also [1].

Chapter 6

Extensions

In this chapter we present two extensions of UNFO+EQ and UNFO+TR which can be shown decidable by slight adaptations of our constructions. We then show that combining these two extensions gives undecidability.

6.1. 1-dimensional guarded negation fragment

We observe now that our small model constructions can be adapted for a slightly bigger logic.

The *guarded negation fragment* of first-order logic, GNFO, is defined in [4] by the following grammar:

$$\varphi = R(\bar{x}) \mid x = y \mid \varphi \wedge \varphi \mid \varphi \vee \varphi \mid \exists x \varphi \mid \gamma(\bar{x}, \bar{y}) \wedge \neg\varphi(\bar{y}),$$

where γ is an atomic formula. Since equality statements of the form $x = x$ can be used as guards, GNFO may be viewed as an extension of UNFO. However, the satisfiability problem for GNFO with equivalences or with transitive relations is undecidable. It follows from the fact that even the two-variable guarded fragment, which is contained in GNFO, becomes undecidable when extended by equivalences or transitive relations [19, 18, 22].

To regain decidability we consider the *base-guarded negation fragment with equivalences (transitive relations)*, BGNFO+EQ (BGNFO+TR). BGNFO+TR is investigated in [1]. In these variants all guards must belong to σ_{base}, and all symbols from σ_{dist} must be interpreted as equivalences/transitive relations. Recall that the general satisfiability problem for BGNFO+TR was shown decidable in [1], and as explained in Chapter 2 this implies decidability of the general satisfiability problem for BGNFO+EQ. In this thesis we do not solve the finite satisfiability problem for neither full BGNFO+EQ nor full BGNFO+TR. We, however, do solve this problem for their one-dimensional restrictions.

We say that a first-order formula is *one-dimensional* if its every maximal block of quantifiers leaves at most one variable free. E.g., $\neg\exists yzR(x,y,z))$ is one-dimensional, and $\neg\exists zR(x,y,z))$ is not. By *one-dimensional guarded negation fragment*, GNFO_1 we mean the subset of GNFO containing its all one-dimensional formulas. Not all UNFO formulas are one-dimensional, but they can be easily converted to the already mentioned UN-normal form [27], which contains only one-dimensional formulas. The cost of this conversion is linear. This allows us to view UNFO as a fragment of GNFO_1. We can define the one-dimensional restrictions BGNFO_1+EQ (BGNFO_1+TR) of BGNFO+EQ (BGNFO+TR) in a natural way.

Our proofs can be adapted to cover the case of BGNFO_1+EQ and BGNFO_1+TR. The adaptations are not difficult. What is crucial is that in the current constructions, during the step of providing witnesses, we build isomorphic copies of whole witness structures, which means that we preserve not only positive atoms but also their negations. Thus, we preserve witness structures.

Theorem 6.1. *(i) $BGNFO_1$+EQ has a doubly exponential finite model property, and its satisfiability (= finite satisfiability) problem is 2*-ExpTime-*complete.*
(ii) The finite satisfiability problem for $BGNFO_1$+TR is 2-ExpTime-*complete.*

Since the adaptations of the proofs covering (i) and (ii) are almost the same, we will show only (ii).

Using a natural adaptation of the standard Scott translation [25] we can transform any BGNFO_1+TR sentence into a normal form sentence φ of the shape as in (2.1), where the φ_i are quantifier-free GNFO formulas. Assume that some finite structure $\mathfrak{A}$ is a model of φ. First, we need a slightly stronger version of condition (h1) in Lemma 2.2—each of the considered homomorphisms should additionally be an isomorphism when restricted to a guarded substructure. We need to extend the notion of a declaration so that it treats subformulas of the form $\gamma(\bar{x},\bar{y}) \wedge \neg\varphi'(\bar{y})$ like non-transitive atomic formulas. This allows us to perform surgery making the transitive paths bounded and then to construct a regular tree-like model $\mathfrak{A}' \models \varphi$ as it is done in the proofs of Lemma 3.5 and Lemma 3.8, respectively. The key facts are that Lemma 3.3 holds (with the new declarations) and that φ_0 is equivalent to a disjunction of some formulas generated by declarations. Finally we apply, without any changes, the construction from the proof of Lemma 5.2 to $\mathfrak{A}'$ and φ obtaining eventually a finite structure $\mathfrak{A}''$. Note that during the step of providing witnesses we build isomorphic copies of partial witness structures, which means that we preserve not only positive atoms but also their negations. Thus the elements of A'' have all witness structures required by φ. Consider now the conjunct $\forall x_1,\ldots,x_t\neg\varphi_0(\bar{x})$, and take arbitrary elements $a_1,\ldots,a_t \in A''$. From Lemma 5.2 we know that there is a homomorphism $\mathfrak{h} : \mathfrak{A}''\restriction\{a_1,\ldots,a_t\} \to \mathfrak{A}'$ preserving 1-types. If $\gamma(\bar{z},\bar{y}) \wedge \neg\varphi'(\bar{y})$ is a subformula of φ_0 with γ a σ_{base}-guard and $\mathfrak{A}'' \models \gamma(\bar{b},\bar{c}) \wedge \neg\varphi'(\bar{c})$ for some $\bar{b},\bar{c} \subseteq \bar{a}$ then, by our construction, all elements of $\bar{b} \cup \bar{c}$ are members of the φ-witness structure for some element. As mentioned above such witness structures are isomorphic copies of substructures from $\mathfrak{A}$ and $\mathfrak{h}$ works on them as an isomorphism, and thus $\mathfrak{h}$ preserves on $\bar{c}$ not only 1-types and positive atoms but also negations of atoms in witnesses structures. Since $\mathfrak{A}' \models \neg\varphi_0(\mathfrak{h}(a_1), \ldots, \mathfrak{h}(a_t))$ this means that $\mathfrak{A}'' \models \neg\varphi_0(a_1, \ldots, a_t)$.

The algorithm for checking finite satisfiability presented in the proof of Thm. 5.5 also works without any changes and, moreover, its correctness proof does not need any modifications. It is the case since a key role is played here by Lemma 3.3 that still holds with the new version of declarations.

6.2. Inclusions of binary relations

For a given signature σ denote $\sigma^{-1} = \{B^{-1} : B \text{ binary and } B \in \sigma_{\text{base}} \cup \sigma_{\text{dist}}\}$, $\sigma^{-1}_{\text{base}} = \{B^{-1} : B \text{ binary and } B \in \sigma_{\text{base}}\}$ and $\sigma^{-1}_{\text{dist}} = \{T^{-1} : T \text{ binary and } T \in \sigma_{\text{dist}}\})$. The *(finite) satisfiability problem for* UNFO+EQ (resp. UNFO+TR) *with binary inclusions* is defined as follows. Given an UNFO+EQ (resp. UNFO+TR) formula φ over σ and a set $\mathcal{H}$ of inclusions of the form $B \subseteq B'$, for binary symbols $B, B' \in \sigma_{\text{base}} \cup \sigma_{\text{dist}} \cup \sigma^{-1}$, check if there exists a (finite) model of φ in which for every $B \subseteq B' \in \mathcal{H}$ the interpretation of B is contained in the interpretation of B', where the interpretations of the symbols from σ^{-1} are the inverses of the interpretations of the corresponding symbols from σ.

We first remark that all our constructions from Chapters 3, 4 and 5, without *literally any changes*, respect inclusions of the form $T \subseteq T'$, $B \subseteq B'$ and $B \subseteq T$ for any $T, T' \in \sigma_{\text{dist}} \cup \sigma^{-1}_{\text{dist}}$ and $B, B' \in \sigma_{\text{base}} \cup \sigma^{-1}_{\text{base}}$. The only problematic inclusions are those of the form $T \subseteq B$ for $B \in \sigma_{\text{base}} \cup \sigma^{-1}_{\text{base}}$ and $T \in \sigma_{\text{dist}} \cup \sigma^{-1}_{\text{dist}}$. However, treating the relations that appear on the right side of such inclusions in a special way, we can apply our constructions so that they respect also this kind of constraints. The corresponding (general) satisfiability problems are also decidable (by much simpler arguments).

Theorem 6.2. *(i) UNFO+EQ with binary inclusions has the finite model property. Its (finite) satisfiability problem is 2-*ExpTime*-complete. (ii) The (finite) satisfiability problem for UNFO+TR with inclusions of binary relations is 2-*ExpTime*-complete.*

We show how to modify our constructions to prove this theorem.

(ii) For a binary $B \in \sigma$ we assume that $(B^{-1})^{-1} = B$. For a given set of inclusions $\mathcal{H}$ let $\mathcal{H}^+$ denote the smallest set such that (a) $\mathcal{H} \subseteq \mathcal{H}^+$, (b) if $B_1 \subseteq B_2 \in \mathcal{H}^+$ then $B_1^{-1} \subseteq B_2^{-1} \in \mathcal{H}^+$, (c) if $B_1 \subseteq B_2 \in \mathcal{H}^+$ and $B_2 \subseteq B_3 \in \mathcal{H}^+$ then $B_1 \subseteq B_3 \in \mathcal{H}^+$. For any structure $\mathfrak{A}$ we have that $\mathfrak{A} \models \mathcal{H}$ iff $\mathfrak{A} \models \mathcal{H}^+$. If $T \subseteq B \in \mathcal{H}$, for $T \in \sigma_{\text{dist}} \cup \sigma^{-1}_{\text{dist}}$ and $B \in \sigma_{\text{base}} \cup \sigma^{-1}_{\text{base}}$ then B is called *pseudo-transitive*. Pseudo-transitive relations must be treated in a special way. Let us see some details.

In our constructions we often proceed as follows. We first build a structure $\mathfrak{A}^0$ by joining together some number of copies of fragments of a pattern model $\mathfrak{A}^{\mathfrak{p}}$ (there are no further connections among such copies; they just share some elements). Relations from σ_{dist} are not transitively closed in $\mathfrak{A}^0$. In particular if $\mathfrak{A}^{\mathfrak{p}} \models \mathcal{H}^+$ then $\mathfrak{A}^0 \models \mathcal{H}^+$.

Then, we take the transitive closure of the appropriate relations in $\mathfrak{A}^0$ obtaining $\mathfrak{A}^\tau$. However, the structure $\mathfrak{A}^\tau$ does not necessarily respect the constraints of the form $T \subseteq B \in \mathcal{H}^+$ for $T \in \sigma_{\mathrm{dist}} \cup \sigma_{\mathrm{dist}}^{-1}$ and $B \in \sigma_{\mathrm{base}} \cup \sigma_{\mathrm{base}}^{-1}$. To solve this problem, we apply the *pseudo-transitive closure* to $\mathfrak{A}^\tau$, that is, for any $T \subseteq B \in \mathcal{H}^+$, such that $T \in \sigma_{\mathrm{dist}} \cup \sigma_{\mathrm{dist}}^{-1}$ and $B \in \sigma_{\mathrm{base}} \cup \sigma_{\mathrm{base}}^{-1}$, and any $a, b \in A^\tau$ if $\mathfrak{A}^\tau \models Tab$ then we join a and b by B. Denote the resulting structure as $\mathfrak{A}$.

It is easy to verify that $\mathfrak{A} \models \mathcal{H}^+$. Now, observe that for each $B \in \sigma_{\mathrm{base}} \cup \sigma_{\mathrm{base}}^{-1}$ and $a, b \in A^\tau (= A)$ we have $\mathfrak{A} \models Bab$ iff $\mathfrak{A}^\tau \models Bab \vee \bigvee_{T : T \subseteq B \in \mathcal{H}^+} Tab$. Let us denote by φ^τ the formula obtained from φ by replacing each σ_{base}-atom with an appropriate disjunction as above. Clearly $\mathfrak{A}^\tau \models \varphi^\tau$ iff $\mathfrak{A} \models \varphi$ and, similarly, since $\mathfrak{A}^{\mathfrak{p}} \models \mathcal{H}^+$, $\mathfrak{A}^{\mathfrak{p}} \models \varphi$ iff $\mathfrak{A}^{\mathfrak{p}} \models \varphi^\tau$. Thus, in all stages of our finite model constructions we may proceed as it was described for UNFO+TR without inclusions, but for the formula φ^τ, and then apply the pseudo-transitive closure to obtain the desired structure.

A similar approach may be applied in the algorithm searching for models respecting inclusions: we search for a model of φ^τ, requiring additionally that the downward families respect inclusions from $\mathcal{H}^+$.

As a side note, we remark that if the transformation of φ into φ^τ is applied to a formula φ from BGNFO$_1$+TR then φ^τ does not necessarily belong to BGNFO$_1$+TR (as a base guard may be replaced by a disjunction containing transitive atoms). Thus, the above arguments do not apply to the guarded case.

(i) Observe that the inclusion $E \subseteq E^{-1}$ is equivalent to E being symmetric. Thus we can embed UNFO+EQ with binary inclusions in UNFO+TR with binary inclusions by adding for each equivalence symbol E a conjunct $\forall x\, Exx$, an inclusion constraint $E \subseteq E^{-1}$ and treating E as a symbol of the distinguished signature. Note that a formula has the same models before and after such transformation. Thus, the (finite) satisfiability problem for UNFO+EQ with binary inclusions is 2-ExpTime-complete. For the finite model property, observe for a satisfiable formula φ (call φ' its image by the above reduction), that the structure obtained by unraveling a model for φ' has no one-directional transitive connections. Therefore, by an application of the further steps of the construction, φ' (and therefore also φ) has a finite model (in fact, since the structures obtained in the constructions have transitive paths bounded by 0, one may use the estimation presented in Chapter 5 to observe that it is doubly exponential).

6.3. Combining the two decidable extensions

Note that BGNFO$_1$+EQ and BGNFO$_1$+TR can express inclusions $B \subseteq T$ and $B \subseteq B'$, and our constructions respect the inclusions of the form $T \subseteq T'$ for $B, B' \in$

$\sigma_{\text{base}} \cup \sigma_{\text{base}}^{-1}$ and $T, T' \in \sigma_{\text{dist}} \cup \sigma_{\text{dist}}^{-1}$. However, it turns out that if we extend them with inclusions of the form $T \subseteq B$ then the (finite) satisfiability problems become undecidable. Since, similarly to the previous section, one can embed BGNFO$_1$+EQ with binary inclusions in BGNFO$_1$+TR with binary inclusions, it suffices to consider the former. The undecidability can be easily shown by a reduction from the already mentioned (finite) satisfiability problem for GF^2+EQ. We build the proof on a conversion of GF formulas to equivalent GNFO formulas presented in [4]. All the negations in a GF^2+EQ formula φ can be guarded by the guards of quantifiers. If φ uses an equivalence guard E then we can add an inclusion $E \subseteq B_E$, for a fresh $B_E \in \sigma_{\text{base}}$ and then use B_E to guard negations. More precisely, subformulas of φ of the form $\exists y(Exy \wedge \psi(x,y))$ are replaced by $\exists y(Exy \wedge \psi^*(x,y))$, where ψ^* is obtained by replacing negated subformulas $\neg\varsigma(x,y)$ (not in the scope of a deeper quantifier) of ψ by, properly base-guarded, $B_E xy \wedge \neg\varsigma(x,y)$. (In subformulas of the original formula of the form $\exists y(Bxy \wedge \psi(x,y))$, for non-equivalence B, we just add the guard Bxy to all binary negations in ψ, not in the scope of a deeper quantifier.) Note that since GF^2+EQ uses only two variables all its formulas are one-dimensional, which is not changed by the described reduction.

Theorem 6.3. *The (finite) satisfiability problems for BGNFO$_1$+EQ with binary inclusions and BGNFO$_1$+TR with binary inclusions are undecidable.*

In fact, the above proof, together with [22, 19], may be used to show that (finite) satisfiability for BGNFO$_1$+EQ (resp. BGNFO$_1$+TR) with three (resp. two) $T \subseteq B$ inclusions, where $T \in \sigma_{\text{dist}}$ and $B \in \sigma_{\text{base}}$, is undecidable.

Bibliography

[1] A. Amarilli, M. Benedikt, P. Bourhis, and M. Vanden Boom, *Query answering with transitive and linear-ordered data*, in Proceedings of the Twenty-Fifth International Joint Conference on Artificial Intelligence, IJCAI 2016, 2016, pp. 893–899.

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

[3] J. Baget, M. Leclère, M. Mugnier, and E. Salvat, *Extending decidable cases for rules with existential variables*, in IJCAI 2009, Proceedings of the 21st International Joint Conference on Artificial Intelligence, 2009, pp. 677–682.

[4] V. Bárány, B. ten Cate, and L. Segoufin, *Guarded negation*, J. ACM, 62 (2015), p. 22.

[5] M. Bojańczyk, C. David, A. Muscholl, T. Schwentick, and L. Segoufin, *Two-variable logic on data words*, ACM Trans. Comput. Log., 12 (2011), p. 27.

[6] M. Bojańczyk, A. Muscholl, T. Schwentick, and L. Segoufin, *Two-variable logic on data trees and xml reasoning*, J. ACM, 56 (2009).

[7] D. Danielski and E. Kieroński, *Unary negation fragment with equivalence relations has the finite model property*, in Thirty-Third Annual ACM/IEEE Symposium on Logic in Computer Science , LICS 2018, 2018.

[8] M. Dzięciołowski, *Satisfability issues for unary negation logic*, bachelor's thesis, University of Wrocław, 2017.

[9] B. Glimm and Y. Kazakov, *Role conjunctions in expressive description logics*, in Logic for Programming, Artificial Intelligence, and Reasoning, 15th International Conference, LPAR 2008, 2008, pp. 391–405.

[10] B. Glimm, C. Lutz, I. Horrocks, and U. Sattler, *Conjunctive query answering for the description logic SHIQ*, J. Artif. Intell. Res., 31 (2008), pp. 157–204.

[11] T. Gogacz, Y. Ibañéz-García, and F. Murlak, *Finite query answering in expressive description logics with transitive roles*, (2018). Accepted for KR'18.

[12] E. Grädel, *On the restraining power of guards*, J. Symb. Log., 64 (1999), pp. 1719–1742.

[13] E. Grädel, P. Kolaitis, and M. Y. Vardi, *On the decision problem for two-variable first-order logic*, Bulletin of Symbolic Logic, 3 (1997), pp. 53–69.

[14] L. Hella and A. Kuusisto, *One-dimensional fragment of first-order logic*, in Proceedings of Advances in Modal Logic, 2014, 2014, pp. 274–293.

[15] I. Horrocks and U. Sattler, *A description logic with transitive and inverse roles and role hierarchies*, J. Log. Comput., 9 (1999), pp. 385–410.

[16] I. Horrocks and U. Sattler, *A tableaux decision procedure for $\mathcal{SHOIQ}$*, Journal of Automated Reasoning, 39 (2007), pp. 249–276.

[17] J. C. Jung, C. Lutz, M. Martel, and T. Schneider, *Querying the unary negation fragment with regular path expressions*, in International Conference on Database Theory, ICDT 2018, 2018, pp. 15:1–15:18.

[18] Y. Kazakov, *Saturation-based decision procedures for extensions of the guarded fragment*, PhD thesis, Universität des Saarlandes, Saarbrücken, Germany, 2006.

[19] E. Kieroński, *Results on the guarded fragment with equivalence or transitive relations*, in Computer Science Logic, vol. 3634 of LNCS, Springer, 2005, pp. 309–324.

[20] E. Kieroński, *On the complexity of the two-variable guarded fragment with transitive guards*, Inf. Comput., 204 (2006), pp. 1663–1703.

[21] E. Kieroński, J. Michaliszyn, I. Pratt-Hartmann, and L. Tendera, *Two-variable first-order logic with equivalence closure*, SIAM Journal of Computing, 43 (2014), pp. 1012–1063.

[22] E. Kieroński and M. Otto, *Small substructures and decidability issues for first-order logic with two variables.*, Journal of Symbolic Logic, 77 (2012), pp. 729–765.

[23] E. Kieroński and L. Tendera, *Finite satisfiability of the two-variable guarded fragment with transitive guards and related variants*, ACM Trans. Comput. Logic, 19 (2018), pp. 8:1–8:34.

[24] I. Pratt-Hartmann, *The finite satisfiability problem for two-variable, first-order logic with one transitive relation is decidable*, Mathematical Logic Quarterly, (2018).

[25] D. Scott, *A decision method for validity of sentences in two variables*, Journal Symbolic Logic, 27 (1962), p. 477.

[26] W. Szwast and L. Tendera, *The guarded fragment with transitive guards*, Annals of Pure and Applied Logic, 128 (2004), pp. 227–276.

[27] B. ten Cate and L. Segoufin, *Unary negation*, Logical Methods in Comp. Sc., 9 (2013).

[28] S. Tobies, *Complexity results and practical algorithms for logics in knowledge representation*, PhD thesis, RWTH Aachen University, Germany, 2001.