

EXTENDED COUNTING
QUANTIFIERS AND
GENERALISATIONS OF THE
FLUTED FRAGMENT OF
FIRST-ORDER LOGIC

A THESIS SUBMITTED TO THE UNIVERSITY OF MANCHESTER
FOR THE DEGREE OF DOCTOR OF PHILOSOPHY
IN THE FACULTY OF SCIENCE AND ENGINEERING

2025

Daumantas Kojelis

Department of Computer Science

Contents

Abstract	6
Declaration	8
Copyright	9
Acknowledgements	10
1 Introduction	13
2 The Two-Variable Fragment	17
2.1 Normal-form and Limits of Expression	18
2.2 Kings, Courts, and Satisfiability	20
2.3 Model Theoretic Characterisations	22
2.4 Bibliographic Notes	25
3 The Fluted Fragment	26
3.1 Normal-form and Fluted Types	27
3.2 The Variable Reduction Technique	28
3.3 Bibliographic Notes	33
4 The Fluted Fragment with Periodic Counting	35
4.1 Preliminaries	36
4.2 The Two-Variable Subfragment	39
4.3 More Than Two Variables	46
4.4 Homogeneity in the Study of Flutedness	50
5 The Adjacent Fragment of First-Order Logic	52
5.1 Preliminaries	54

5.2	Primitive generators of words	56
5.3	Expressive Power	59
5.4	Upper bounds for \mathcal{AF}^ℓ without equality	64
5.4.1	Excursion: the equality-free 3-variable case	74
5.5	Adding Equality	79
5.6	The Guarded Subfragment	93
5.7	Extending the Adjacent Fragment	102
6	The Adjacent Fragment with Counting is Undecidable	105
6.1	Undecidability of Finite Satisfiability	107
6.2	Undecidability of General Satisfiability	110
7	The 3-variable Adjacent Fragment with Counting	120
7.1	Limiting Härtig Quantification	121
7.2	Deciding $\text{Sat}(\mathcal{C}_{\mathcal{UHA}}^2)$	125
7.3	Reducing \mathcal{AFC}^3 to $\mathcal{C}_{\mathcal{UHA}}^2$	133
8	Conclusions	156
	Bibliography	159

List of Tables

5.4.1 Quick reference guide for Sec 5.1, 5.2 and 5.4.	65
5.5.1 Quick reference guide for Sec 5.5.	79
7.3.1 Quick reference guide for Sec 7.3.	144

List of Figures

7.3.1 Possible configuration after Step 2.3 (with element b omitted). Unidirectional lines indicate unidirectional rays. Bidirectional lines are bidirectional rays. Note that elements (e.g. c_1) need not receive or emit rays.	150
7.3.2 Possible configuration after Step 2.4 (with element b omitted). Unidirectional lines indicate unidirectional rays. Bidirectional lines are bidirectional rays. Note that unidirectional rays are only emitted from left to right (with elements of $U_\zeta^{(2)}$ wrapping around).	151
7.3.3 Assigning anti-rays in step 3 when $\zeta \stackrel{\mathcal{L}}{\sim} \eta$ (with element b omitted). Unidirectional lines indicate unidirectional rays. Bidirectional lines are bidirectional rays. When this occurs there is nothing more to set.	154

Abstract

We study the satisfiability problem of the *fluted fragment* – a fragment of first-order logic in which variables are applied to predicates in the order they were quantified in – under two kinds of generalisations. First, we extend the syntax of the language with what are known as *periodic counting quantifiers* and show that the resulting language, \mathcal{FLPC} , has a (finite) satisfiability problem that is in $(\ell-1)$ -NEXPTIME when \mathcal{FLPC} is restricted to ℓ -variable formulas, and TOWER-complete in the general case. To do so we rely on a new observation concerning satisfiable sentences of fluted languages. We show that such sentences admit models which are (in a sense we will make clear) *homogeneous*. We believe that the study of homogeneous models allows us to further extend the expressive capabilities of the fluted fragment whilst also simplifying existing decidability proofs.

On the other hand, we relax the syntactic restrictions of flutedness thus forming the *adjacent fragment* – a fragment in which $\mathbf{p}(x_{i_1} \cdots x_{i_n})$ is allowed as an atom as long as $|i_{j+1} - i_j| \leq 1$ for each $j \in [1, n-1]$, it being understood that variables x_1, x_2, \dots are quantifier in that order. We show that the new language (denoted \mathcal{AF}) generalises the *two-variable fragment* of first-order logic and has a (finite) satisfiability problem that is TOWER-complete. When considering ℓ -variable formulas, the complexity drops to $(\ell-1)$ -NEXPTIME and, in the absence of equality, $(\ell-2)$ -NEXPTIME. In addition, we show that the *guarded adjacent fragment* has a (finite) satisfiability problem that is 2-EXPTIME-complete. We argue that \mathcal{AF} is a maximal decidable (in terms of finite and general satisfiability) fragment obtained by restricting variable sequences.

To supplement our decidability results we show that the adjacent fragment turns to be undecidable (in terms of finite and general satisfiability) in the presence of counting and periodic counting extensions. In particular, we show that the three-variable adjacent fragment with counting quantifiers has a finite satisfiability problem that is Σ_1^0 -complete. If periodic counting quantifiers are allowed,

Σ_1^0 -hardness transfers to the general satisfiability problem. The four-variable adjacent fragment with counting will be shown to be complete of Π_1^0 in terms of the general satisfiability problem. In the presence of periodic counting, the satisfiability problem will be shown to be complete for Σ_1^1 .

Our final result is that concerning the general satisfiability problem of the three-variable adjacent fragment with counting. Intriguingly, we will show that the problem is decidable, albeit, without a complexity-theoretic upper-bound. In the process we will show that two-variable logic can be extended with what we call *uniform Härtig assertions* (i.e. limited forms of *Härtig quantification*) all the while retaining decidability of general satisfiability.

Declaration

No portion of the work referred to in this thesis has been submitted in support of an application for another degree or qualification of this or any other university or other institute of learning.

Copyright

- i. The author of this thesis (including any appendices and/or schedules to this thesis) owns certain copyright or related rights in it (the “Copyright”) and s/he has given The University of Manchester certain rights to use such Copyright, including for administrative purposes.
- ii. Copies of this thesis, either in full or in extracts and whether in hard or electronic copy, may be made **only** in accordance with the Copyright, Designs and Patents Act 1988 (as amended) and regulations issued under it or, where appropriate, in accordance with licensing agreements which the University has from time to time. This page must form part of any such copies made.
- iii. The ownership of certain Copyright, patents, designs, trade marks and other intellectual property (the “Intellectual Property”) and any reproductions of copyright works in the thesis, for example graphs and tables (“Reproductions”), which may be described in this thesis, may not be owned by the author and may be owned by third parties. Such Intellectual Property and Reproductions cannot and must not be made available for use without the prior written permission of the owner(s) of the relevant Intellectual Property and/or Reproductions.
- iv. Further information on the conditions under which disclosure, publication and commercialisation of this thesis, the Copyright and any Intellectual Property and/or Reproductions described in it may take place is available in the University IP Policy (see <http://documents.manchester.ac.uk/DocuInfo.aspx?DocID=24420>), in any relevant Thesis restriction declarations deposited in the University Library, The University Library’s regulations (see <http://www.library.manchester.ac.uk/about/regulations/>) and in The University’s policy on presentation of Theses

Acknowledgements

I would like to thank my fellow co-authors Bartosz Jan Bednarczyk and Ian Pratt-Hartmann for their support and numerous suggestions regarding various publications. In particular, I would like to express gratitude to Ian Pratt-Hartmann for introducing me to the study of decidable fragments, exchanging ideas about formal languages, and reading nearly every line of this thesis.

I extend my gratitude to my colleagues at the University of Manchester who have offered their friendship in daily struggles and joys. I thank my parents, grandparents, great-grandparents and friends back in Lithuania – all of whom showed support during my studies. Most of all, I would like to thank my dear partner Elena who has joined me in the best and worst aspects of years past.

General Notation

Common Notation

- ω_0 – the first infinite ordinal.
- ω_1 – the first uncountable ordinal.
- \aleph_0 – the first infinite cardinal.
- \aleph_1 – the second infinite cardinal.
- \mathbb{N} – the natural numbers $0, 1, 2, \dots$.
- \mathbb{N}^* – the set $\mathbb{N} \cup \{\aleph_0\}$.
- \bar{a} – some word $a_1 \cdots a_n$, it being understood that $n \geq 0$.
- \tilde{a} – the reversal of \bar{a} , i.e. $a_n \cdots a_1$.

Computability Theory

- Σ_1^0 – the subsets of \mathbb{N} definable in Peano arithmetic with only existential first-order quantifiers. Also, by Post's theorem, the RECURSIVELY ENUMERABLE sets.
- Π_1^0 – the class of sets $\{\mathbb{N} \setminus S \mid S \in \Sigma_1^0\}$. Alternatively, the CO-RECURSIVELY ENUMERABLE sets.
- Δ_1^0 – the class of sets $\Sigma_1^0 \cap \Pi_1^0$. Alternatively, the RECURSIVE sets.
- Σ_1^1 – the subsets of \mathbb{N} definable in Peano arithmetic with existential second-order quantification (and any form of first-order quantification).
- Π_1^1 – the class of sets $\{\mathbb{N} \setminus S \mid S \in \Sigma_1^1\}$.

Formal Languages

- \mathcal{FO} – first-order logic.
- fo – first-order logic without the equality symbol.
- $\mathcal{L}_{\omega_1, \omega}$ – infinitary logic; i.e. first-order logic with conjunctions and disjunctions over sets (of first-order formulas) of size less than $|\omega_1|$.
- \mathcal{FO}^2 – two-variable first-order logic.
- \mathcal{C}^2 – two-variable logic with counting quantifiers.
- $\mathcal{FO}_{\text{Pres}}^2$ – two-variable logic with ultimately periodic counting quantifiers.
- $\mathcal{C}_{\text{UHA}}^2$ – the two-variable fragment with uniform Härtig assertions.
- \mathcal{FL} – the fluted fragment of first-order logic.
- fl – the fluted fragment without the equality symbol.
- \mathcal{FLC} – the fluted fragment with counting quantifiers.
- \mathcal{FLPC} – the fluted fragment with periodic counting quantifiers.
- \mathcal{AF} – the adjacent fragment of first-order logic.
- af – the adjacent fragment without the equality symbol.
- \mathcal{AFC} – the adjacent fragment with counting quantifiers.
- \mathcal{AFPC} – the adjacent fragment with periodic counting quantifiers.
- \mathcal{GF} – the guarded fragment of first-order logic.
- \mathcal{GA} – the guarded adjacent fragment of first-order logic.
- \mathcal{L}^ℓ – the language \mathcal{L} restricted to variables x_1, \dots, x_ℓ .

Computational Problems

- $\text{Sat}(\mathcal{L})$ – the satisfiability problem of the language \mathcal{L} .
- $\text{FinSat}(\mathcal{L})$ – the finite satisfiability problem of the language \mathcal{L} .

Chapter 1

Introduction

First-order Logic, \mathcal{FO} , has been an object of intense study in the past few centuries. Numerous results regarding \mathcal{FO} have since seen applications in Mathematics, Philosophy, and, most importantly for us, Computer Science. An essential question regarding the computational aspects of \mathcal{FO} is that posed by D. Hilbert and W. Ackermann – the *Entscheidungsproblem* [Hilbert and Ackermann, 1928]. In short, the D. Hilbert and W. Ackermann ask for an algorithm that decides if a given first-order sentence is logically valid. From a philosophical point of view, the existence of such an algorithm would signal that the process of deriving mathematical proofs (in \mathcal{FO} at least) is nothing more than mechanical. It, perhaps, can be viewed as a fortunate occurrence that no such algorithm can exist as evident by the works of A. Church [Church, 1936] and A. Turing [Turing, 1936].

The negative answer to the Entscheidungsproblem prompted the study of *fragments* of \mathcal{FO} , usually derived by (i) limiting quantifier prefixes, (ii) allowing only variables x and y to appear, or (iii) relativising quantification (as in the guarded fragment). In this thesis we continue to study the Entscheidungsproblem for fragments \mathcal{FO} (in fact, fragments of infinitary logic $\mathcal{L}_{\omega_1, \omega}$). We will do so in terms of (*finite*) *satisfiability* – the problem of determining if a sentence has a (finite) model satisfying it. Given a language \mathcal{L} , we will write $\text{Sat}(\mathcal{L})$ and $\text{FinSat}(\mathcal{L})$ to reference the general and finite satisfiability problems for \mathcal{L} respectively.

The fragments of \mathcal{FO} we will be most concerned with in this thesis do not conform with the limitations (i)–(iii) above. We will instead follow the ideas of V. W. Quine and look at *argument-sequence logics* – a lesser-known family of formalisms obtained by limiting variable sequences which are allowed to appear as

arguments in predicates. One argument-sequence logic which will figure throughout the thesis is the *fluted fragment*, \mathcal{FL} . The history of the fragment is long and complicated (an account is given in Chapter 3). For our current purposes, it suffices to roughly define the fluted fragment as a fragment of first-order logic in which variables appear as arguments (in predicates) in the order which they were quantified in. As it turns out, imposing such limitations restores decidability of the (finite) satisfiability problem [Pratt-Hartmann et al., 2016, Pratt-Hartmann et al., 2019]. The fluted fragment is of particular interest in the field of computation as it is a multi-variable generalisation of various *description logics*. Indeed, after a routine translation, formulas of the description logic \mathcal{ALCH} (i.e. \mathcal{ALC} with role hierarchies) are contained in the two-variable sub-fragment of \mathcal{FL} . In the past few years the fluted fragment has been studied under various syntactic and semantic extensions such as counting quantifiers (i.e. quantifiers requesting a particular amount of existential witnesses) [Pratt-Hartmann, 2021], transitive relations [Pratt-Hartmann and Tendera, 2019, Pratt-Hartmann and Tendera, 2022], and even the combination of the two [Pratt-Hartmann and Tendera, 2023].

We show decidability of (finite) satisfiability for \mathcal{FL} in Chapter 3 by giving a variable reduction procedure similar to that in [Pratt-Hartmann et al., 2019]. One of our main contribution (in Chapter 4) concerning the fluted paradigm is a new observation about its satisfiable sentences (possibly with counting quantifiers). In short, such sentences always have a model that behaves (in a sense that we will make clear) *homogeneously*. Section 4.1 gives background on the fluted fragment with counting extensions. Section 4.2 shows that 2-variable fluted sentences admit what we call *globally homogeneous* models. With that, we establish that the fluted fragment with periodic counting quantifiers (i.e. existential quantifiers requesting that the number of witnesses) has a (finite) satisfiability problem that is NEXPTIME-complete. Section 4.3 established that satisfiable ℓ -variable fluted sentences always have a *locally $(\ell-1)$ -homogeneous* model. We adapt the variable reduction technique presented in Chapter 3 to fit the setting of periodic counting thus showing that (finite) satisfiability for the ℓ -variable subfragment can be checked in $(\ell-1)$ -NEXPTIME. With that we conclude that the fluted fragment with periodic counting quantifiers has a (finite) satisfiability problem that is TOWER-complete. We finish the chapter by advocating, in Section 4.4, for the usefulness of homogenous models in establishing new results about fluted languages and simplifying old ones.

The remainder of the thesis is concerned with a new formalism we call the *adjacent fragment* of first-order logic, \mathcal{AF} , and its extensions. The fragment is probably best explained from the comfort of flutedness. Roughly put, if $\psi(x_1 \cdots x_\ell)$ is a formula in the fluted fragment, then $\psi(x_{i_1} \cdots x_{i_\ell})$ is in the adjacent fragment as long as $|i_{j+1} - i_j| \leq 1$ for each $j \in [1, \ell-1]$. As it will become apparent, \mathcal{AF} consumes not only \mathcal{FL} , but also the *two-variable fragment* of first-order logic, \mathcal{FO}^2 . In fact, to show decidability of satisfiability for \mathcal{AF} , we will rely on the fact that the two-variable fragment has a decidable (finite) satisfiability problem. Necessary background on \mathcal{FO}^2 is provided in Chapter 2.

We formally introduce the adjacent fragment in Chapter 5 and prove the following results. Sections 5.1–5.2 give the necessary tools for studying the adjacent fragment. We study the expressive power of \mathcal{AF} in Section 5.3 by providing appropriate notions of bisimulations in \mathcal{AF} . It is in that section that \mathcal{FO}^2 -formulas are shown to be logically equivalent to those of \mathcal{AF} . Section 5.4 studies the upper bound for the (finite) satisfiability problem of ℓ -variable equality-free adjacent fragment. In particular, we show that the finite and general satisfiability problems coincide and reside in $(\ell-2)$ -NEXPTIME by means of variable reductions similar to those in Chapter 3, and combinatoric tricks reminiscent of those in Chapter 2. In Section 5.5 we show that the finite and general satisfiability problems also coincide if the equality symbol is allowed. Again, by the variable reduction technique, we will have that the ℓ -variable adjacent fragment (with equality) has a (finite) satisfiability problem that is in $(\ell-1)$ -NEXPTIME. With the above we will conclude that \mathcal{AF} (with equality and unlimited variables) has a satisfiability problem that is TOWER-complete (Theorem 5.5.4). Section 5.5 studies the *guarded* adjacent fragment, \mathcal{GAF} . It is shown that the (finite) satisfiability problem for \mathcal{GAF} is 2-EXPTIME-complete thus making it no easier than that of the full guarded fragment. Lastly, Section 5.7 argues that the adjacent fragment is a maximal decidable (in terms of satisfiability) fragment amongst those obtained by limiting variable sequences. In that sense, we answer W. V. Quine’s question on where one *limit of decision* lies.

Having established the above, we turn to syntactic extensions for the adjacent fragment. In line with our thesis, we will consider counting and periodic counting extensions. Since the (finite) satisfiability problem is decidable for \mathcal{FL} with both counting [Pratt-Hartmann, 2021] and periodic counting (see Chapter 4), it is natural to conjecture that the same is true for \mathcal{AF} with the appropriate

extensions. The conjecture is further supported by the fact that the base cases: the two-variable fragment with counting, \mathcal{C}^2 , and the two-variable fragment with periodic counting, $\mathcal{FO}_{\text{Pres}}^2$, both have decidable (finite) satisfiability problems [Pratt-Hartmann, 2010, Benedikt et al., 2024]. The significance of a decidability result comes from the fact that the three-variable guarded fragment with counting quantifiers is undecidable [Grädel, 1999]. Thus, \mathcal{AF} with counting extensions is a possible candidate for generalising description logics such as \mathcal{ALCHIQ} (i.e. \mathcal{ALC} with role hierarchies, role inverses and cardinality restrictions) in a multivariable fashion without losing decidability of (finite) satisfiability.

In Chapter 6 we show that the above conjecture is, for the most part, false. Section 6.1 shows that the finite satisfiability problem for the three-variable adjacent fragment with counting quantifiers is Σ_1^0 -complete. Σ_1^0 -hardness is retained for the satisfiability problem when periodic counting quantifiers are allowed. Section 6.2 houses results concerning Π_1^0 -completeness of the satisfiability problem for the 4-variable adjacent fragment with counting quantifiers. It will be shown that the problem is Σ_1^1 -complete if periodic counting quantifiers are allowed.

Notice that nothing is said about the satisfiability problem for the three-variable adjacent fragment with counting quantifiers. We will show, in Chapter 7, that the language in question is decidable for general satisfiability, albeit without a complexity theoretic upper-bound. Doing so will require analysing the two-variable fragment with what we call *uniform Härtig assertions*. We define the assertions and language in Section 7.1 and show that it is decidable for satisfiability in Section 7.2. Results obtained in this chapter will conclude our study of the adjacent paradigm under counting extensions.

Chapter 2

The Two-Variable Fragment

The *two-variable fragment*, \mathcal{FO}^2 , is a fragment of first-order logic formed by restricting the set of (free or bound) variables to x and y . It is usually left implicit that \mathcal{FO}^2 -formulas are built over function-free signatures. We do not deviate from this restriction, though, additionally, we forbid the use of constant symbols.

To get a feel for the expressive power of \mathcal{FO}^2 , consider the sentence “Every musician practices with another musician”. The sentence can be easily formalised in the two-variable language as follows:

$$\forall x \left(\text{mus}(x) \rightarrow \exists y (\text{mus}(y) \wedge \text{practices}(xy) \wedge \text{practices}(yx)) \right). \quad (\mathcal{FO}^2\text{-ex1})$$

Additionally, one can, in \mathcal{FO}^2 , assert that an element is unique (at least in terms of the 1-variable, quantifier-free formulas it satisfies). Thus, sentences such as “There is only a single person named Ravel that is a composer” can be written:

$$\forall xy \left((\text{ravel}(x) \wedge \text{comp}(x) \wedge \text{ravel}(y) \wedge \text{comp}(y)) \rightarrow x = y \right). \quad (\mathcal{FO}^2\text{-ex2})$$

In this chapter we will be concerned with the finite and general satisfiability problems for \mathcal{FO}^2 . More specifically, we give an adaptation of a classical result by E. Grädel, P. Kolaitis and M. Vardi [Grädel et al., 1997a] which classifies $\text{Sat}(\mathcal{FO}^2)$ and $\text{FinSat}(\mathcal{FO}^2)$ as NEXPTIME problems. We will then characterise the expressive power of \mathcal{FO}^2 by providing an appropriate notion of *bisimulations*. The tools developed in this chapter will serve as a foundation for results to come.

2.1 Normal-form and Limits of Expression

Consider the quantifier prefixes in (\mathcal{FO}^2 -ex1), (\mathcal{FO}^2 -ex2). In the context of the satisfiability problem, it was shown by D. Scott in [Scott, 1962] that one only needs to consider finite sets of prenex normal-form formulas $\forall x \exists y \theta$ and $\forall xy \theta$. Alternatively, such sets can be viewed as *normal-form* sentences:

$$\forall xy \alpha(xy) \wedge \bigwedge_{t \in T} \forall x \exists y \beta_t(xy), \quad (\mathcal{FO}^2\text{-nmf})$$

where α, β_t are quantifier-free \mathcal{FO}^2 -formulas indexed by a finite set T .

We say that two formulas φ and ψ are *equisatisfiable* when φ is (finitely) satisfiable if and only if ψ is as well. Restricting attention to sentences of the form (\mathcal{FO}^2 -nmf) comes at no loss of generality (at least in the realm of finite and general satisfiability), we show the following using standard rewriting techniques:

Lemma 2.1.1. (Scott's reduction [Scott, 1962]). *Suppose φ is an \mathcal{FO}^2 -sentence. Then, we may compute, in polynomial time, an equisatisfiable sentence ψ in normal-form.*

Proof. Write $\varphi_0 := \varphi$ and suppose that φ_0 has a subformula $\theta(\bar{u}) := Qv \chi(\bar{u}v)$, where $Q \in \{\exists, \forall\}$, χ is a quantifier-free \mathcal{FO}^2 -formula, $v \in \{x, y\}$, and where \bar{u} is either the variable u , where $u \neq v$, or the empty string. Now, let \mathbf{p} be a fresh predicate of arity $|\bar{u}|$. Writing $\bar{\exists} := \exists$, $\bar{\forall} := \forall$, and \bar{u}^* for the string \bar{u} but with each occurrence of u replaced by x we define ψ_1 as

$$\forall x Qy (\mathbf{p}(\bar{u}^*) \rightarrow \chi(\bar{u}^*y)) \wedge \forall x \bar{Q}y (\chi(\bar{u}^*y) \rightarrow \mathbf{p}(\bar{u}^*)),$$

and set φ_1 to be φ_0 but with $\theta(\bar{u})$ replaced by the atom $\mathbf{p}(\bar{u})$. It is easy to verify that $\varphi_1 \wedge \psi_1 \models \varphi_0$. To see that (finite) satisfiability of φ_0 entails that of $\varphi_1 \wedge \psi_1$ fix $\mathfrak{A} \models \varphi_0$ and let \mathfrak{A}' be the expansion of \mathfrak{A} obtained by setting $\mathfrak{A}' \models \mathbf{p}[\bar{a}]$ if and only if $\mathfrak{A} \models \theta[\bar{a}]$ for each $\bar{a} \in A^{|\bar{u}|}$. It is immediate that $\mathfrak{A}' \models \varphi_1 \wedge \psi_1$.

Repeating the rewriting procedure on φ_1 and subsequent sentences we will obtain a series of sentences ψ_1, \dots, ψ_m where, for each $i \in [1, m]$,

$$\psi_i := \forall xy \alpha_i(xy) \wedge \forall x \exists y \beta_i(xy),$$

and φ_m that is composed of proposition letters. Writing ψ for

$$\forall xy \left(\bigwedge_{i \in [1, m]} \alpha_i(xy) \wedge \varphi_m \right) \wedge \bigwedge_{i \in [1, m]} \forall x \exists y \beta_i(xy)$$

we will have a sentence of the form (\mathcal{FO}^2 -nmf) that is equisatisfiable to φ_0 . \square

Returning to (\mathcal{FO}^2 -ex2) it is natural to ask how far uniqueness requirements can be taken in \mathcal{FO}^2 . For instance, can we ensure, via an \mathcal{FO}^2 -sentence, the existence of exactly $n \geq 2$ elements satisfying the same 1-variable quantifier-free formulas? The answer turns out to be no; at least for sentences of the form (\mathcal{FO}^2 -nmf). We proceed with some preliminaries. Take any normal-form sentence φ and denote its signature by $\text{sig}(\varphi)$. Writing $\Sigma := \text{sig}(\varphi)$ and taking $i \in \{1, 2\}$ define an *atomic i -type* ζ over Σ to be a maximal consistent set of literals from $\Sigma \cup \{=\}$ in the variables $\{x\}$ when $i = 1$ and $\{x, y\}$ when $i = 2$. Whenever convenient, we identify ζ as the conjunction of literals $\bigwedge \zeta$. Fixing a model $\mathfrak{A} \models \varphi$ it is clear that each i -tuple of elements satisfies a unique atomic i -type over Σ . We denote the atomic i -type realised by $\bar{a} \in A^i$ in \mathfrak{A} by $\text{tp}^{\mathfrak{A}}(\bar{a})$. The following is immediate:

Fact 2.1.2. *Suppose θ_1 and θ_2 are quantifier-free \mathcal{FO}^2 -formulas having free variables in $\{x\}$ and $\{x, y\}$ respectively. Fixing some structure \mathfrak{A} we have, for all $i \in \{1, 2\}$ and $\bar{a} \in A^i$, that $\mathfrak{A}, \bar{a} \models \theta_i$ if and only if $\text{tp}^{\mathfrak{A}}(\bar{a}) \models \theta_i$.*

Now, define a *star* σ over Σ to be a collection of atomic 2-types over Σ satisfying the following:

- there is a 1-type ζ such that $\xi \models \zeta$ for each $\xi \in \sigma$.

In \mathfrak{A} we again have that each $a \in A$ realises a unique star $\text{str}^{\mathfrak{A}}(a)$ defined as $\{\text{tp}^{\mathfrak{A}}(ab) \mid b \in A\}$. Given a quantifier-free \mathcal{FO}^2 -formula θ , we will write $\sigma \models \forall y \theta$ if $\xi \models \theta$ for all $\xi \in \sigma$, and, similarly, $\sigma \models \exists y \theta$ if $\xi \models \theta$ for some $\xi \in \sigma$. The following will then be useful in the forthcoming proofs:

Fact 2.1.3. *Take any structure \mathfrak{A} and let θ be any quantifier-free \mathcal{FO}^2 -formula. Fixing $a \in A$ and $\sigma := \text{str}^{\mathfrak{A}}(a)$ we have, for $Q \in \{\exists, \forall\}$, that $\mathfrak{A}, a \models Qy \theta$ if and only if $\sigma \models Qy \theta$.*

Thus, the satisfaction of normal-form formulas by \mathfrak{A} can be checked by taking the stars realised in \mathfrak{A} and checking if they entail $\forall y \alpha$ and $\exists y \beta_t$ for each $t \in T$.

Let us fix $\mathfrak{A} \models \varphi$, where φ is a normal-form \mathcal{FO}^2 -sentence. Taking any atomic 1-type ζ suppose it is realised in \mathfrak{A} by at least 2 elements, say c_1 and c_2 . Taking some fresh element $c' \notin A$ we claim that there is a model $\mathfrak{A}' \models \varphi$ over the domain $A' := A \cup \{c'\}$ such that $\mathfrak{A}'|_A = \mathfrak{A}$ and $\text{tp}^{\mathfrak{A}'}(c') = \zeta$. To achieve this we set $\mathfrak{A}'|_A := \mathfrak{A}$ whilst also copying over 2-types that c_1 is involved in over to c' as follows. First, set $\text{tp}^{\mathfrak{A}'}(c'c_1) := \text{tp}^{\mathfrak{A}}(c_2c_1)$. Secondly, set $\text{tp}^{\mathfrak{A}'}(c'a) := \text{tp}^{\mathfrak{A}}(c_1a)$ for each $a \in A \setminus \{c_1\}$. Notice that, by our assignment, $\text{str}^{\mathfrak{A}'}(a) = \text{str}^{\mathfrak{A}}(a)$ for each $a \in A$. Recalling that φ is of the form (\mathcal{FO}^2 -nmf) it is immediate, by Fact 2.1.3, that $\mathfrak{A}', a \models \forall y \alpha \wedge \bigwedge_{t \in T} \exists y \beta_t$. We thus need only show that $\mathfrak{A}', c' \models \exists y \beta_t$ for each $t \in T$ and that $\mathfrak{A}', c' \models \forall y \alpha$. Write $\sigma := \text{str}^{\mathfrak{A}'}(c')$. By our construction, $\sigma = \text{str}^{\mathfrak{A}}(c_1) \cup \{\text{tp}^{\mathfrak{A}}(c_2c_1)\}$. Since $\text{str}^{\mathfrak{A}}(c_1) \subseteq \sigma$, we have, by Fact 2.1.3, that $\mathfrak{A}, c_1 \models \exists y \beta_t$ implies $\mathfrak{A}', c' \models \exists y \beta_t$ for each $t \in T$. On the other hand, recall that $\mathfrak{A}, c \models \forall y \alpha$ for $c \in \{c_1, c_2\}$. Since $\sigma \subseteq \text{str}^{\mathfrak{A}}(c_1) \cup \text{str}^{\mathfrak{A}}(c_2)$, we have, again by Fact 2.1.3, that $\mathfrak{A}', c' \models \forall y \alpha$. We have thus established the following:

Lemma 2.1.4. (Pumping) *Suppose φ is a normal-form \mathcal{FO}^2 -formula and take $\mathfrak{A} \models \varphi$. Then, elements having a 1-type realised multiple times in \mathfrak{A} can be duplicated as given above thus forming a model $\mathfrak{A}' \models \varphi$.*

2.2 Kings, Courts, and Satisfiability

Putting (\mathcal{FO}^2 -ex2) and Lemma 2.1.4 together it is obvious that, when considering the satisfiability problem, atomic 1-types realised in a structure a single time need to be handled with care. Fix \mathfrak{A} to be a model of any normal-form sentence φ and let $K \subseteq A$ be the set of elements having atomic 1-types which are realised only once in \mathfrak{A} . We go by the terminology used in [Grädel et al., 1997a] and call members of K *kings* and those of $A \setminus K$ *common*. Writing $\|\varphi\|$ for the number of symbols used in φ it should be clear that $|K| \leq 2^{\|\varphi\|}$ as there are at most $2^{\|\varphi\|}$ different atomic 1-types over $\text{sig}(\varphi)$. Define a *court* \mathfrak{C} of \mathfrak{A} to be any minimal substructure of \mathfrak{A} that houses the kings along with their witnesses for $\exists y \beta_t$ for each $t \in T$. More formally, a court \mathfrak{C} of \mathfrak{A} is a structure over $C := K \cup S$, where S is any minimal subset of A satisfying the following: for each $k \in K$ and $t \in T$ there is some $a \in S$ such that $\mathfrak{C} \models \beta_t[ka]$. It should be clear that $|C|$ is no larger than $|K| + |K||T|$. We invite the reader to regard \mathfrak{C} as the substructure of \mathfrak{A} that acts “irregularly”. Building on top of Lemma 2.1.4 we show that elements $A \setminus C$ can be substituted by a finite surrogate set E in which existential requirements

are fulfilled, as it were, in a circular pattern:

Lemma 2.2.1. (Circular Witnessing, [Grädel et al., 1997a]) *Suppose φ is of the form $(\mathcal{FO}^2\text{-nmf})$ and $\mathfrak{A} \models \varphi$. Then, any court \mathfrak{C} of \mathfrak{A} can be expanded into a structure $\mathfrak{D} \models \varphi$ over $C \cup E$, where $|E| \leq 2^{|\varphi|} \cdot 3|T|$.*

Proof. Let S be the set of atomic 1-type in \mathfrak{A} that are realised by common elements. That is to say, elements realising the 1-types of S in \mathfrak{A} can be duplicated as given in Lemma 2.1.4. Without loss of generality let $T = [1, |T|]$. We define $E := E_0 \cup E_1 \cup E_2$, where, for each $i \in [0, 2]$, E_i is a set of fresh elements $\{e_{\zeta,t} \mid t \in T, \zeta \in S\}$, and set the domain of \mathfrak{D} to be $D := C \cup E$. Taking $e_{\zeta,*} \in E$ we first specify that $\text{tp}^{\mathfrak{D}}(e_{\zeta,*}) := \zeta$. Thus, each 1-type $\zeta \in S$ is realised in \mathfrak{D} at least $3|T|$ times. Setting $\mathfrak{D}|_C := \mathfrak{C}$ notice that, by the definition of the court, each king element already has a witness for each of the $t \in T$ existential requirements $\exists y \beta_t$. We thus turn to providing existential witnesses for elements in $D \setminus K$.

We begin by first picking some $e_{\zeta,*} \in E$. Find some $a \in A$ such that $\text{tp}^{\mathfrak{A}}(a) = \zeta$. We will call a our *reference element in \mathfrak{A} for $e_{\zeta,*}$* . Since $\mathfrak{A}, a \models \exists y \beta_t$ for each $t \in T$, pick $\sigma^- \subseteq \text{str}^{\mathfrak{A}}(a)$ to be any smallest subset of atomic 2-types such that $\sigma^- \models \exists y \beta_t$ for each $t \in T$. At the end of our construction we will have that $e_{\zeta,*}$ realises some star $\sigma \supseteq \sigma^-$ thus guaranteeing, by Fact 2.1.3, that $\mathfrak{D}, e_{\zeta,*} \models \exists y \beta_t$ for each $t \in T$. Writing $n = |\sigma^-|$ we enumerate the atomic 2-types of σ^- as $(\xi_k)_{k \in [1, n]}$ and define $(b_k)_{k \in [1, n]}$ to be some sequence of elements from A such that $\mathfrak{A} \models \xi_k[ab_k]$. Since there are no repeated entries in σ^- , each element in $(b_k)_{k \in [1, n]}$ is distinct. Fixing some $k \in [1, n]$ we proceed as follows. In case $a = b_k$ we already have that $\text{tp}^{\mathfrak{D}}(e_{\zeta,*}e_{\zeta,*}) = \text{tp}^{\mathfrak{A}}(aa) = \xi_k$ as, by our assignment of atomic 1-types, $\text{tp}^{\mathfrak{D}}(e_{\zeta,*}) = \zeta = \xi_k(xx)$. If $b_k \in K$ we set $\text{tp}^{\mathfrak{D}}(e_{\zeta,*}b_k) := \xi_k$. Lastly, if b_k is common, we proceed as follows. Write $\eta := \text{tp}^{\mathfrak{A}}(b_k)$. Suppose that $e_{\zeta,*} \in E_i$ (where $i \in [0, 2]$) and let $+_3$ be addition modulo 3. Recalling that $n \leq |T|$, we have that the set $E_{i+_3 1}$ contains the element $e_{\eta,k}$. We can thus set $\text{tp}^{\mathfrak{D}}(e_{\zeta,*}e_{\eta,k}) := \xi_k$. That is to say, we let $e_{\eta,k}$ mimic the witness b_k . Since there are $|T| \geq n$ copies of $e_{\eta,k}$, and each element in $(b_k)_{k \in [1, n]}$ is distinct, we can perform the assignments above in parallel for each $k \in [1, n]$. Moreover, since we are looking for witnesses in a circular pattern around the three disjoint sets E_0 , E_1 and E_2 , we have that repeating the procedure above for each $e_{\zeta,*} \in E$ does not override any previous assignments.

Now, pick some $c \in C \setminus K$. Since $c \in A$, we can simply pick c itself to be the reference element in \mathfrak{A} for c and (since $c \notin E_0 \cup E_1 \cup E_2$) find witnesses in $K \cup E_0$

as above. (Note that if some $b_k \in K$, the reassignment $\text{tp}^{\mathfrak{D}}(cb_k) := \text{tp}^{\mathfrak{A}}(cb_k)$ is inconsequential as $\text{tp}^{\mathfrak{D}}(cb_k) = \text{tp}^{\mathfrak{C}}(cb_k)$).

In conclusion, for each $d \in D$ and $t \in T$ we have $\mathfrak{D}, d \models \exists y \beta_t$. Additionally, none of the atomic 2-types assigned violate the universal requirement α as they were copied over from the model $\mathfrak{A} \models \forall xy \alpha$. It might be, of course, that some two elements $d_1, d_2 \in D$ do not have an atomic 2-type set between them. If that is indeed the case, notice that, by our construction, $d_1 \notin K$ or $d_2 \notin K$. We may then pick any two distinct elements $a_1, a_2 \in A$ having $\text{tp}^{\mathfrak{D}}(d_1) = \text{tp}^{\mathfrak{A}}(a_1)$ and $\text{tp}^{\mathfrak{D}}(d_2) = \text{tp}^{\mathfrak{A}}(a_2)$ and set $\text{tp}^{\mathfrak{D}}(d_1 d_2) := \text{tp}^{\mathfrak{A}}(a_1 a_2)$. Clearly, $\mathfrak{D} \models \alpha[d_1 d_2]$ as $\mathfrak{A} \models \alpha[a_1 a_2]$. By repeating this assignment for each $d_1, d_2 \in D$ that has not been considered by our construction, we will have that $\mathfrak{D} \models \varphi$ as required. \square

Given an \mathcal{FO}^2 -sentence φ it is then easy to devise a decision procedure. Compute an equisatisfiable sentence ψ of the form (\mathcal{FO}^2 -nmf). This can be done in polynomial time as per Lemma 2.1.1. Then, guess a model \mathfrak{A} of size at most $2^{|\psi|} + 2^{|\psi|}|T| + 2^{|\psi|} \cdot 3|T|$. By Lemma 2.2.1 such a model will exist if ψ is satisfiable. Checking that $\mathfrak{A} \models \psi$ can be checked in time $O(|A|^2)$ as, by Fact 2.1.3, we need only extract the stars realised in \mathfrak{A} and check them against α and β_t for each $t \in T$. We have thus shown the following:

Theorem 2.2.2. [Grädel et al., 1997a] *The (finite) satisfiability problem for \mathcal{FO}^2 is in NEXPTIME. In fact, $\text{Sat}(\mathcal{FO}^2) = \text{FinSat}(\mathcal{FO}^2)$, or, in other words, \mathcal{FO}^2 has the finite model property.*

The NEXPTIME upper-bound attributed to $\text{Sat}(\mathcal{FO}^2)$ above is tight. Indeed, in [Fürer, 1983] it was shown $\text{Sat}(\mathcal{FO}^2)$ is NEXPTIME-hard. In conclusion:

Theorem 2.2.3. *The (finite) satisfiability problem for \mathcal{FO}^2 is NEXPTIME-complete.*

2.3 Model Theoretic Characterisations

We finish the chapter by introducing well-known tools for studying the expressive power of \mathcal{FO}^2 . Let \mathfrak{A} be any structure and take $\bar{a} \in A^k$ for $k \leq 2$. We call the pair (\mathfrak{A}, \bar{a}) a *pointed structure*. When considering formulas with free variables x, y , we treat a_1 as a substitute for x , and a_2 as the substitute for y . For convenience let us fix two pointed structures (\mathfrak{A}, \bar{a}) and (\mathfrak{B}, \bar{b}) for the rest of the section. We write $(\mathfrak{A}, \bar{a}) \simeq^{\mathcal{FO}^2} (\mathfrak{B}, \bar{b})$ if:

1. **Harmony:** $\text{tp}^{\mathfrak{A}}[\bar{a}] = \text{tp}^{\mathfrak{B}}[\bar{b}]$,
2. **Forth:** for all $i \in [1, |\bar{a}|]$ and $c \in A$ exists $d \in B$ s.t. $(\mathfrak{A}, a_i c) \simeq^{\mathcal{FO}^2} (\mathfrak{B}, b_i d)$,
3. **Back:** for all $i \in [1, |\bar{b}|]$ and $d \in B$ exists $c \in A$ s.t. $(\mathfrak{A}, a_i c) \simeq^{\mathcal{FO}^2} (\mathfrak{B}, b_i d)$.

We say that (\mathfrak{A}, \bar{a}) and (\mathfrak{B}, \bar{b}) are \mathcal{FO}^2 -bisimilar just in case $(\mathfrak{A}, \bar{a}) \simeq^{\mathcal{FO}^2} (\mathfrak{B}, \bar{b})$. Now, let us write $(\mathfrak{A}, \bar{a}) \equiv^{\mathcal{FO}^2} (\mathfrak{B}, \bar{b})$ just in case $\mathfrak{A} \models \psi[\bar{a}] \Leftrightarrow \mathfrak{B} \models \psi[\bar{b}]$ for every \mathcal{FO}^2 -formula ψ . That is to say, $(\mathfrak{A}, \bar{a}) \equiv^{\mathcal{FO}^2} (\mathfrak{B}, \bar{b})$ just in case the two pointed structures satisfy the same 2-variable formulas. Using standard techniques we link the notion of \mathcal{FO}^2 -bisimulations and satisfaction of 2-variable formulas:

Lemma 2.3.1. $(\mathfrak{A}, \bar{a}) \simeq^{\mathcal{FO}^2} (\mathfrak{B}, \bar{b})$ implies $(\mathfrak{A}, \bar{a}) \equiv^{\mathcal{FO}^2} (\mathfrak{B}, \bar{b})$. The converse is true over ω -saturated structures.

Proof. We show the first statement by structural induction on formulas $\psi \in \mathcal{FO}^2$. More specifically, we will show that $(\mathfrak{A}, \bar{a}) \simeq^{\mathcal{FO}^2} (\mathfrak{B}, \bar{b})$ implies $\mathfrak{A} \models \psi[\bar{a}] \Leftrightarrow \mathfrak{B} \models \psi[\bar{b}]$. Firstly, if ψ is atomic, then $\mathfrak{A} \models \psi[\bar{a}] \Leftrightarrow \text{tp}^{\mathfrak{A}}[\bar{a}] \models \psi$ and $\mathfrak{B} \models \psi[\bar{b}] \Leftrightarrow \text{tp}^{\mathfrak{B}}[\bar{b}] \models \psi$. By harmony, $\text{tp}^{\mathfrak{A}}[\bar{a}] = \text{tp}^{\mathfrak{B}}[\bar{b}]$ thus securing the required property. The cases where $\psi := \neg\theta$ and $\psi := \chi \wedge \theta$ are trivial. Thus, we are only left to deal with¹ $\psi := \exists y \theta$. Keeping $(\mathfrak{A}, \bar{a}) \simeq^{\mathcal{FO}^2} (\mathfrak{B}, \bar{b})$ suppose that $\mathfrak{A} \models \psi[\bar{a}]$. Thus, there is some $c \in A$ such that $\mathfrak{A} \models \theta[a_1 c]$. By forth we may find $d \in B$ such that $(\mathfrak{A}, a_1 c) \simeq^{\mathcal{FO}^2} (\mathfrak{B}, b_1 d)$. By our inductive hypothesis, $\mathfrak{B} \models \theta[b_1 d]$ as required. The backwards direction is handled similarly.

For the second statement assume that \mathfrak{A} and \mathfrak{B} are ω -saturated and that $(\mathfrak{A}, \bar{a}) \equiv^{\mathcal{FO}^2} (\mathfrak{B}, \bar{b})$. We claim that the set

$$\mathbf{S} := \bigcup_{n \leq 2} \{(\bar{c}, \bar{d}) \in A^n \times B^n \mid (\mathfrak{A}, \bar{c}) \equiv^{\mathcal{FO}^2} (\mathfrak{B}, \bar{d})\}$$

is an \mathcal{FO}^2 -bisimulation between the two pointed structures. Clearly, the harmony condition is satisfied by pairs in \mathbf{S} . To show the back and forth properties take some $(\bar{c}, \bar{d}) \in \mathbf{S}$. Picking $i \in [1, |\bar{c}|]$ and $c' \in A$ we show that there is some $d' \in B$ such that $(\mathfrak{A}, c_i c') \equiv^{\mathcal{FO}^2} (\mathfrak{B}, d_i d')$. To do so we claim that the set of formulas

$$\Gamma := \{\psi \in \mathcal{FO}^2 \mid \mathfrak{A} \models \psi[c_i c']\}$$

is realised by $d_i d'$ in \mathfrak{B} for some $d' \in B$. We first show that the claim holds for every finite subset $\gamma \subseteq \Gamma$. Let us identify γ as the conjunction of formulas it

¹The case $\psi := \exists x \theta$ is dealt with by swapping instances of x and y .

contains. Clearly, $\mathfrak{A}, c_i \models \exists y \gamma$. Since $(\mathfrak{A}, \bar{c}) \equiv^{\mathcal{FO}^2} (\mathfrak{B}, \bar{d})$ we have that $\mathfrak{B}, d_i \models \exists y \gamma$. By compactness, Γ is then consistent with the theory of (\mathfrak{B}, d_i) . Since \mathfrak{B} is ω -saturated, we may find an element $d' \in B$ such that $\mathfrak{B} \models \Gamma[d_i d']$. Clearly, $(\mathfrak{A}, c_i c') \equiv^{\mathcal{FO}^2} (\mathfrak{B}, d_i d')$ thus securing $(c_i c', d_i d') \in \mathbf{S}$. The back condition is established analogously. \square

Let us take any \mathcal{FO} -formula ψ . We say that ψ is *invariant under \mathcal{FO}^2 -bisimulations* just in case $(\mathfrak{C}, \bar{c}) \simeq^{\mathcal{FO}^2} (\mathfrak{D}, \bar{d})$ implies $\mathfrak{C} \models \psi[\bar{c}] \Leftrightarrow \mathfrak{D} \models \psi[\bar{d}]$ for all pointed pairs of structures (\mathfrak{C}, \bar{c}) and (\mathfrak{D}, \bar{d}) . The following lemma attests that our notion of \mathcal{FO}^2 -bisimulations captures which \mathcal{FO} -definable properties are in the 2-variable fragment.

Lemma 2.3.2. *Suppose that $\psi \in \mathcal{FO}$ is invariant under \mathcal{FO}^2 -bisimulations. Then ψ is logically equivalent to an \mathcal{FO}^2 -formulas*

Proof. In case ψ is inconsistent, then the required formula is \perp . Thus let us suppose the opposite. Defining Ψ to be the \mathcal{FO}^2 consequences of ψ ; i.e.:

$$\Psi := \{\theta \in \mathcal{FO}^2 \mid \psi \models \theta\}.$$

Our end-goal is to show $\Psi \models \psi$. By our assumption that ψ is consistent, we have that so is Ψ . Thus, there is a pointed structure (\mathfrak{C}, \bar{c}) satisfying Ψ . Let us write:

$$\Gamma := \{\theta \in \mathcal{FO}^2 \mid \mathfrak{C} \models \theta[\bar{c}]\}.$$

Clearly, $\Gamma \cup \{\psi\}$ is consistent as otherwise, by compactness, there is a finite subset $\gamma \subseteq \Gamma$ such that $\gamma \models \neg\psi$. But then, identifying γ as the conjunction of its (finitely many) components we have that $\psi \models \neg\gamma$. Then, by definition, $\neg\gamma \in \Psi$ thus contradicting $\mathfrak{C} \models \gamma[\bar{c}]$. Having argued that $\Gamma \cup \{\psi\}$ is consistent, we may take (\mathfrak{D}, \bar{d}) to be a pointed structure such that $\mathfrak{D} \models \Gamma[\bar{d}]$ and $\mathfrak{D} \models \psi[\bar{d}]$. Clearly, $(\mathfrak{C}, \bar{c}) \equiv^{\mathcal{FO}^2} (\mathfrak{D}, \bar{d})$. Thus, writing \mathfrak{C}' and \mathfrak{D}' for the ω -saturated extensions of \mathfrak{C} and \mathfrak{D} respectively, we have, by Lemma 2.3.1, that $(\mathfrak{C}', \bar{c}) \simeq^{\mathcal{AF}} (\mathfrak{D}', \bar{d})$. Thus, by our initial assumption of invariance, $\mathfrak{C}' \models \psi[\bar{c}]$ and thus $\mathfrak{C} \models \psi[\bar{c}]$. Since (\mathfrak{C}, \bar{c}) was chosen arbitrarily, we conclude that $\Psi \models \psi$.

To see that ψ is equivalent to a formula in \mathcal{FO}^2 we need only apply compactness. Indeed, we have that there is a finite subset $\psi' \subseteq \Psi$ such that $\psi' \models \psi$. \square

2.4 Bibliographic Notes

The origins of \mathcal{FO}^2 date back to 1962, when D. Scott showed that the satisfiability problem for \mathcal{FO}^2 was reducible to that of an already decidable fragment – the *Gödel fragment* [Gödel, 1933, Scott, 1962]. It is now known, however, that the reduction does not imply decidability of satisfiability for the *full* two-variable fragment as the Gödel fragment becomes undecidable for satisfiability in the presence of the equality predicate [Goldfarb, 1984]. The first proof of decidability for the satisfiability problem of the *full* fragment of \mathcal{FO}^2 was not until 1975 by M. Mortimer [Mortimer, 1975]. The work confirmed that \mathcal{FO}^2 with equality has a decidable satisfiability problem; in fact, M. Mortimer showed that, if an \mathcal{FO}^2 -sentence is satisfiable, then it is satisfiable in a model of finite size. Whilst no complexity bounds were specifically outline in the article, an upper-bound of 2-NEXPTIME can be read-off. It wasn't until 1997 that NEXPTIME-completeness for the problem was established by E. Grädel, P. Kolaitis and M. Vardi [Grädel et al., 1997a]. The decidability procedure in [Grädel et al., 1997a] is more or less what we have in Lemma 2.2.1.

Ever since, the two-variable fragment has been scrutinised under various syntactic and semantic extensions. Of particular interest to us are counting extensions. It has been known since 1997 that the two-variable fragment augmented with counting quantifiers has a decidable satisfiability problem, albeit, without the finite model property [Grädel et al., 1997b]. The satisfiability problem was classified as being in 2-NEXPTIME in 2000 [Pacholski et al., 2000], with a result showing NEXPTIME-completeness following in 2005 [Pratt-Hartmann, 2005]. Turning to recent affairs, in 2020 it has been shown that the satisfiability problem remains decidable in the presence of non-first-order counting extensions such as periodic counting quantifiers [Benedikt et al., 2020, Benedikt et al., 2024]. But this result appears to be approaching the limits: it has been known since 1999 that the 2-variable language with equicardinality expressions (such as Härtig quantifiers) has a highly undecidable satisfiability problem [Grädel et al., 1999].

Chapter 3

The Fluted Fragment

The *fluted fragment (without equality)*, fl , is a fragment of \mathcal{FO} in which, roughly put, variables appear in predicates following the order in which they were quantified. For illustrative purposes, we translate the sentence “Every conductor nominates their favorite soloist to play at every concert” into this language as follows:

$$\forall x_1 \left(\text{cnd}(x_1) \rightarrow \exists x_2 (\text{sl}(x_2) \wedge \text{fv}(x_1 x_2) \wedge \forall x_3 (\text{cnc}(x_3) \rightarrow \text{nm}(x_1 x_2 x_3))) \right). \quad (fl\text{-ex1})$$

As a non-example, the sentences axiomatising transitivity, symmetry and reflexivity of a relation are not in the fluted fragment.

The fluted fragment is a member of *argument-sequence logics* – a family of decidable (in terms of satisfiability) languages which also includes the *ordered* [Herzig, 1990, Jaakkola, 2021] and *forward* [Bednarczyk, 2021] fragments. Formally, the fluted fragment is the union of sets of formulas $fl^{[\ell]}$ defined by simultaneous induction as follows:

- (i) any atom $\mathbf{r}(x_k, \dots, x_\ell)$, where x_k, \dots, x_ℓ is a contiguous subsequence of x_1, x_2, \dots and \mathbf{r} is a predicate of arity $\ell - k + 1$, is in $fl^{[\ell]}$;
- (ii) $fl^{[\ell]}$ is closed under Boolean combinations;
- (iii) if $\varphi \in fl^{[\ell+1]}$, then $\exists x_{\ell+1} \varphi$ is in $fl^{[\ell]}$.

Whenever convenient, we identify the formula $\neg \exists x \neg \varphi$ as $\forall x \varphi$. We write $fl := \bigcup_{\ell \geq 0} fl^{[\ell]}$ for the set of all fluted formulas and define the ℓ -variable fluted fragment to be the set $fl^\ell := fl \cap \mathcal{FO}^\ell$. We will assume that all fluted formulas are built over function- and constant-free signatures. Historically, the equality symbol is disallowed in fl . We follow this standard in the current chapter.

Since fluted atom features a suffix of the variable quantification order, variables convey no meaningful information and can thus be omitted in fluted formulas. As an example, the sentence (*fl-ex1*) may be written as

$$\forall \left(\text{cnd} \rightarrow \exists (\text{s1} \wedge \text{fv} \wedge \forall (\text{cnc} \rightarrow \text{nm})) \right),$$

without ambiguity (up to a shift of variable indices).

3.1 Normal-form and Fluted Types

We proceed similarly as in Section 2.1. Fix some $\ell \geq 2$ and write \forall^ℓ for $\forall x_1 \cdots \forall x_\ell$. We say that an $\text{fl}^{\ell+1}$ -sentence is in *normal-form* if it takes the following shape:

$$\bigwedge_{r \in R} \forall^\ell (\alpha_R \rightarrow \forall \gamma_r) \wedge \bigwedge_{t \in T} \forall^\ell (\beta_t \rightarrow \exists \delta_t), \quad (\text{fl-nmf})$$

where α_r, β_t are quantifier-free fl^ℓ -formulas whilst γ_r, δ_t are quantifier-free $\text{fl}^{\ell+1}$ -formulas indexed by finite sets R and T . When considering satisfiability, we can, without loss of generality, confine ourselves to normal-form formulas as:

Lemma 3.1.1. *Take any $\text{fl}^{\ell+1}$ -sentence φ . Then we may compute, in polynomial time, an equisatisfiable $\text{fl}^{\ell+1}$ -sentence ψ in normal form (*fl-nmf*).*

Proof. Writing $\varphi_0 := \varphi$, take any subformula $\theta := Q\chi$ of φ_0 , where $Q \in \{\exists, \forall\}$ and χ is quantifier-free. Supposing there are k free variables in θ , let \mathfrak{q} be a fresh predicate of arity k . Writing $\bar{\exists} := \forall$ and $\bar{\forall} := \exists$ we define ψ_1 to be

$$\forall^\ell (\mathfrak{q} \rightarrow Q\chi) \wedge \forall^\ell (\neg \mathfrak{q} \rightarrow \bar{Q} \neg \chi)$$

and set φ_1 to be φ_0 but with θ replaced by \mathfrak{q} . Clearly, $\varphi_1 \wedge \psi_1 \models \varphi_0$. Conversely, if $\mathfrak{A} \models \varphi_0$, we may expand \mathfrak{A} to \mathfrak{A}' by setting $\mathfrak{q}^{\mathfrak{A}'} := \{\bar{a} \in A^k \mid \mathfrak{A} \models \theta[\bar{a}]\}$. Then, $\mathfrak{A}' \models \varphi_1 \wedge \psi_1$ as required. Processing φ_1 and subsequent formulas in the same way, we are left with a sentence φ_m composed solely of proposition letters along with sentences ψ_1, \dots, ψ_m of the form above. Then, $\psi := \psi_1 \wedge \cdots \wedge \psi_m \wedge \forall^\ell (\top \rightarrow \forall \varphi_m)$ is then of the required form. \square

Fix some relational signature Σ and take any predicate symbol $\mathfrak{p} \in \Sigma$ of arity $k \leq \ell$. Then, for $i \geq 0$, the formulas $p(x_{i-k+1} \cdots x_i)$ and $\neg p(x_{i-k+1} \cdots x_i)$ are *i-literals*. A *fluted i-type* over Σ is then a maximal consistent set of fluted *i-literals*

over Σ . Given a Σ -structure \mathfrak{A} we have that each i -tuple $\bar{a} \in A^i$ realises a unique fluted i -type in \mathfrak{A} denoted by $\text{ftp}^{\mathfrak{A}}[\bar{a}]$. The following is immediate:

Fact 3.1.2. *Suppose θ is a quantifier-free fl^i -formula. Taking any structure \mathfrak{A} and $\bar{a} \in A^i$ we have that $\mathfrak{A} \models \theta[\bar{a}]$ if and only if $\text{ftp}^{\mathfrak{A}}[\bar{a}] \models \theta$.*

Recall that $\tilde{a} = a_i \cdots a_1$. It is important to note that, when $\bar{a} \neq \tilde{a}$, the fluted i -types of \bar{a} contains no information about the fluted i -type of \tilde{a} . We will write FTP_i^Σ for the set of all fluted i -type over Σ . Given $\zeta \in \text{FTP}_i^\Sigma$ we denote by $\zeta|_{[j,i]}$, the fluted $(i-j+1)$ -type obtained by deleting entries of ζ of arity greater than $i-j+1$. It is implicitly assumed that $\zeta|_{[j,i]}$ is in the variables x_1, \dots, x_{i-j+1} .

3.2 The Variable Reduction Technique

Let us for the rest of the section fix a normal-form $fl^{\ell+1}$ -sentence φ over some constant- and function-free signature Σ . We describe a variable reduction procedure and run it on φ thus producing an fl^ℓ -sentence ψ of the form

$$\psi_1 \wedge \psi_2 \wedge \psi_3. \quad (\psi)$$

In addition, we will have that ψ is satisfiable if and only if φ is. By running the procedure again on ψ (as well as subsequent formulas) we will eventually reach a sentence in fl^2 . Having $fl^2 \subsetneq \mathcal{FO}^2$ along with the fact that $\text{Sat}(\mathcal{FO}^2)$ is decidable (Theorem 2.2.2), we will conclude that $\text{Sat}(fl^{\ell+1})$ is decidable as well. But we are getting ahead of ourselves; let us return to the original sentence φ . The advertised ℓ -variable sentence ψ will simply be a collection of facts about models of φ .

Take any model \mathfrak{A} of φ as an example. (Note that ψ is constructed solely from the syntactic properties of φ ; \mathfrak{A} is only here for motivational purposes). Whilst defining the advertised sentence we will also expand \mathfrak{A} to \mathfrak{A}^+ which will be a model of ψ . We start by introducing a series of $(\ell-1)$ -ary predicates $(\mathbf{q}_\zeta)_{\zeta \in \text{FTP}_\ell^\Sigma}$ and by setting $\mathbf{q}_\zeta^{\mathfrak{A}^+} := \{\bar{a} \in A^{\ell-1} \mid \mathfrak{A} \models \zeta[a\bar{a}]\}$ for each $\zeta \in \text{FTP}_\ell^\Sigma$. To explain the motivation behind \mathbf{q}_ζ let us fix $\bar{a} \in A^{\ell-1}$. The predicate \mathbf{q}_ζ simply remembers if \bar{a} can be extended (by appending an element to the left) to realise the fluted ℓ -type ζ . It should be clear that \mathfrak{A}^+ models the following fl^ℓ -sentence:

$$\bigwedge_{\zeta \in \text{FTP}_\ell^\Sigma} \forall^\ell (\zeta \rightarrow \mathbf{q}_\zeta). \quad (\psi_1)$$

Let us now consider the universal requirements of φ . To this end take $R' \subseteq R$ and write $\alpha_{R'}$ for $\bigwedge_{r \in R'} \alpha_r$ and γ_r for $\bigwedge_{r \in R'} \gamma_r$. Supposing some fluted ℓ -type ζ satisfies $\zeta \models \alpha_{R'}$, we must have in \mathfrak{A} that each $(\ell+1)$ -tuple $a\bar{a}b$ with $\text{ftp}^{\mathfrak{A}}[a\bar{a}] = \zeta$ realises a fluted $(\ell+1)$ -type ξ such that $\xi \models \gamma_{R'}$. Of course, in the ℓ -variable setting, we cannot enforce that the appropriate $(\ell+1)$ -types are attributed. We thus shift our attention to \mathfrak{A}^+ . Notice that the satisfaction of \mathbf{p}_ζ by $\bar{a} \in A^{\ell-1}$ implies that, for every element $b \in A$, the fluted ℓ -type $\eta := \text{ftp}^{\mathfrak{A}}[\bar{a}b]$ must be included in some fluted $(\ell+1)$ -type ξ that is compatible with the requirements $\gamma_{R'}$. More precisely, $\xi|_{[2, \ell+1]} = \eta$, where $\xi \models \gamma_{R'}$. In the sequel we will argue that this property gives enough information in order to assign $(\ell+1)$ -types without violating $\gamma_{R'}$. But for now we have shown that \mathfrak{A}^+ is a model of the following:

$$\bigwedge_{R' \subseteq R} \bigwedge_{\zeta \in \text{FTP}_\ell^\Sigma}^{\zeta \models \alpha_{R'}} \forall^{\ell-1}(\mathbf{q}_\zeta \rightarrow \forall \bigvee_{\xi \in \text{FTP}_{\ell+1}^\Sigma}^{\xi \models \gamma_{R'}} \xi|_{[2, \ell+1]}). \quad (\psi_2)$$

We translate the existential requirements of φ similarly as in ψ_2 . Let us again fix some $R' \subseteq R$ and, additionally, take some existential requirement $t \in T$. In case ζ is a fluted ℓ -type satisfying $\zeta \models \alpha_{R'}$ we must have that, for each ℓ -tuple $a\bar{a}$ satisfying $\text{ftp}^{\mathfrak{A}}[a\bar{a}] = \zeta$, there is an element $b \in A$ with which we have $\text{ftp}^{\mathfrak{A}}[a\bar{a}b] \models \gamma_{R'} \wedge \delta_t$. Again, in \mathfrak{A}^+ , we can make no reference to the possible fluted $(\ell+1)$ -types that $a\bar{a}b$ may realise. We see, however, that $\bar{a}b$ realises a fluted ℓ -type that is compatible with $\text{ftp}^{\mathfrak{A}}[a\bar{a}b]$. Again, we will later argue that this property is sufficient for assigning fluted $(\ell+1)$ -types. Right now, we conclude the definition by specifying the following whilst noting that \mathfrak{A}^+ is a model of:

$$\bigwedge_{t \in T} \bigwedge_{R' \subseteq R} \bigwedge_{\zeta \in \text{FTP}_\ell^\Sigma}^{\zeta \models \alpha_{R'} \wedge \beta_t} \forall^{\ell-1}(\mathbf{q}_\zeta \rightarrow \exists \bigvee_{\xi \in \text{FTP}_{\ell+1}^\Sigma}^{\xi \models \gamma_{R'} \wedge \delta_t} \xi|_{[2, \ell+1]}). \quad (\psi_3)$$

Recalling that $\psi = \psi_1 \wedge \psi_2 \wedge \psi_3$ we have shown the following:

Lemma 3.2.1. (Variable Reduction \Leftarrow , [Pratt-Hartmann et al., 2019]) *Suppose that $\mathfrak{A} \models \varphi$. Then, \mathfrak{A} can be expanded to a model $\mathfrak{A}^+ \models \psi$.*

When dealing with equality-free formulas, we may freely duplicate elements in their models. Let \mathfrak{B} be a Σ -structure, and H a non-empty set of indices. We define the structure $\mathfrak{B} \times H$ over the Cartesian product $B \times H$ as follows: for any $p \in \Sigma$ of arity k , and any k -tuples $b_1 \cdots b_k$ over B and $h_1 \cdots h_k$ over H , set

$\mathfrak{B} \times H \models \mathbf{r}[\langle b_1, h_1 \rangle \cdots \langle b_k, h_k \rangle]$ if and only if $\mathfrak{B} \models \mathbf{r}[b_1 \cdots b_k]$. The following is then almost immediate.

Lemma 3.2.2. *Take any equality-free \mathcal{FO} -formula χ and a non-empty set H . Fixing any structure \mathfrak{B} we have that, for each $\bar{b} \in B^k$ and $\bar{h} \in H^k$, $\mathfrak{B} \times H \models \chi[\langle b_1, h_1 \rangle \cdots \langle b_k, h_k \rangle]$ if and only if $\mathfrak{B} \models \chi[\bar{b}]$.*

Proof. We proceed by structural induction on χ . Suppose first that χ is an atomic formula. Then, by definition, $\mathfrak{B} \models \chi[\bar{b}]$ if and only if $\mathfrak{B} \times H \models \chi[\langle b_1, h_1 \rangle \cdots \langle b_k, h_k \rangle]$ for each $\bar{b} \in B^k$ and $\bar{h} \in H^k$.

The cases $\chi := \neg\chi_1$ and $\chi := \chi_1 \wedge \chi_2$ are then routine. We are thus left with $\chi := \exists x_{k+1} \theta(x_1 \cdots x_{k+1})$. Take any $\bar{b} \in B^k$ and suppose $\mathfrak{B} \models \chi[\bar{b}]$. Then, there is some element $b_{k+1} \in B$ satisfying $\mathfrak{B} \models \theta[\bar{b}b_{k+1}]$. Applying the inductive hypothesis on θ we have that $\mathfrak{B} \models \theta[\bar{b}b_{k+1}]$ if and only if $\mathfrak{B} \times H \models \theta[\langle b_1, h_1 \rangle \cdots \langle b_{k+1}, h_{k+1} \rangle]$ for each $h_1, \dots, h_{k+1} \in H$. Thus, $\mathfrak{B} \times H \models \chi[\langle b_1, h_1 \rangle \cdots \langle b_k, h_k \rangle]$ for all $h_1, \dots, h_k \in H$ as required. \square

With Lemma 3.2.2 at hand we are now ready to prove the converse direction.

Lemma 3.2.3. (Variable Reduction \Rightarrow , [Pratt-Hartmann et al., 2019]) *Suppose that $\mathfrak{B} \models \psi$. Then, there is a model $\mathfrak{C} \models \varphi$ such that $|C|/|B| \leq |T|$.*

Proof. Let us write \mathfrak{B}^- for the Σ -reduct of \mathfrak{B} . Notice that $\text{sig}(\psi)$ does not contain any symbols of arity greater than ℓ . Thus, we may assume that extensions to $(\ell+1)$ -ary symbols in \mathfrak{B} and \mathfrak{B}^- are left undefined. We proceed by specifying them as suggested by ψ . We proceed by first identifying pseudo-witnesses (in regards to the existential requirements indexed by T) for tuples $\bar{a} \in A^\ell$. Taking any $\bar{a} \in A^\ell$ let $\zeta := \text{ftp}^{\mathfrak{B}^-}[\bar{a}]$. By ψ_1 , this implies that $\mathfrak{B} \models \mathbf{q}_\zeta[a_2 \cdots a_\ell]$. Now, let $R' := \{r \in R \mid \zeta \models \alpha_r\}$ and $T' := \{t \in T \mid \zeta \models \beta_t\}$. By ψ_3 , for each $t \in T'$ there is an element $b_{\bar{a},t} \in A$ and a fluted $(\ell+1)$ -type $\xi_{\bar{a},t}$ such that

- $\xi_{\bar{a},t} \models \delta_t$,
- $\xi_{\bar{a},t} \models \bigwedge_{r \in R'} \gamma_r$, and
- $\mathfrak{B} \models \xi_{\bar{a},t}|_{[2,\ell+1]}[a_2 \cdots a_\ell b_t]$

Let us write $\eta := \xi_{\bar{a},t}|_{[2,\ell+1]}$ and denote by η^+ the result of incrementing variable indices in η by 1. Since $\xi_{\bar{a},t} \setminus \eta^+$ only contains $(\ell+1)$ -ary symbols, we are able to assign the fluted $(\ell+1)$ -type $\xi_{\bar{a},t}$ to $a\bar{a}b$ without fear of contradictions. Notice

that this would provide \bar{a} with a witness for the t -th existential requirement of φ . Of course, we cannot assign a fluted $(\ell+1)$ -type just yet as $b_{\bar{a},t}$ might also be picked to be a witness for some other existential requirement, say $t' \in T$, thus making $b_{\bar{a},t} = b_{\bar{a},t'}$. In that case, there is no guarantee that $\xi_{\bar{a},t} = \xi_{\bar{a},t'}$. But we are getting ahead of ourselves. It suffices, at this point, that we have, for each ℓ -tuple, identified elements and fluted $(\ell+1)$ -type that can serve as witnesses for the existential requirements of φ .

In order to prevent clashes when defining fluted $(\ell+1)$ -types we will shift focus to the structure \mathfrak{C} which is a “blown-up” version of \mathfrak{B}^- . Formally, we define $\mathfrak{C} := \mathfrak{B}^- \times T$ and proceed with the assignments of fluted $(\ell+1)$ -types as follows. Take any $\bar{c} \in C^\ell$ and write $c_i = \langle a_i, h_i \rangle$ for $i \in [1, \ell]$. Writing $\bar{a} := a_1 \cdots a_\ell$ and $\zeta := \text{ftp}^{\mathfrak{B}^-}[\bar{a}]$, we see, by Lemma 3.2.2, that $\text{ftp}^{\mathfrak{C}}[\bar{c}] = \zeta$. Now, let $R' := \{r \in R \mid \zeta \models \alpha_r\}$ and $T' := \{t \in T \mid \zeta \models \beta_t\}$. We provide witnesses to $(\xi_{\bar{a},t})_{t \in T}$ for \bar{c} in \mathfrak{C} thus satisfying the existential requirements of φ . Taking any $t \in T'$ we have previously identified $b_{\bar{a},t} \in B$ as a possible witness for \bar{a} for the fluted $(\ell+1)$ -type $\xi_{\bar{a},t}$. We claim that $\bar{c}\langle b_{\bar{a},t}, t \rangle$ can be assigned $\xi_{\bar{a},t}$ in \mathfrak{C} . Indeed, since $\bar{a}b_{\bar{a},t}$ can be assigned $\xi_{\bar{a},t}$, and $\text{ftp}^{\mathfrak{B}^-}[a_2 \cdots a_\ell b_{\bar{a},t}] = \text{ftp}^{\mathfrak{C}}[c_2 \cdots c_\ell \langle b_{\bar{a},t}, t \rangle]$ by Lemma 3.2.2, we can consistently set $\text{ftp}^{\mathfrak{C}}[\bar{c}\langle b_{\bar{a},t}, t \rangle] := \xi_{\bar{a},t}$. Notice that the second coordinate of witnesses picked for \bar{c} will differ depending on $t \in T'$. Thus, we can perform this assignment for each $t \in T'$ without fear of clashes. It is then evident that $\mathfrak{C}, \bar{c} \models \exists \delta_t$ for each $t \in T'$. As for $t \in T \setminus T'$ we have that $\mathfrak{C} \not\models \beta_t[\bar{c}]$. Moreover, the universal requirements of φ are not violated by any fluted $(\ell+1)$ -type assigned so far. Indeed, we have already argued that $\xi_{\bar{a},t} \models \bigwedge_{r \in R'} \gamma_r$, whilst for $r \in R \setminus R'$ we have $\mathfrak{C} \not\models \alpha_r[\bar{c}]$. Now, let us take $\bar{c}' \in C^\ell$ such that $\bar{c}' \neq \bar{c}$. Clearly, the fluted $(\ell+1)$ -type given to $\bar{c}d$ has no bearing on that of $\bar{c}'d$ for any $d \in C$. We may thus repeat the procedure for each $\bar{c} \in C^\ell$; and, in doing so, ensure that the existential requirements of φ are met in \mathfrak{C} .

Let us now suppose that some $\bar{c} \in C^{\ell+1}$ does not yet hold a fluted $(\ell+1)$ -type. Take $\zeta := \text{ftp}^{\mathfrak{B}^-}[c_1 \cdots c_\ell]$ and write $R' := \{r \in R \mid \zeta \models \alpha_r\}$. Evidently, the fluted $(\ell+1)$ -type we assign to \bar{c} must be consistent with γ_r for all $r \in R'$. Let us write $c_i = \langle a_i, h_i \rangle$ for $i \in [1, \ell+1]$, $\bar{a} := a_1 \cdots a_{\ell+1}$, and turn back to the structures \mathfrak{B} and \mathfrak{B}^- . By Lemma 3.2.2, $\text{ftp}^{\mathfrak{B}^-}[a_1 \cdots a_\ell] = \zeta$ and thus, by ψ_1 , $\mathfrak{B} \models \mathfrak{q}_\zeta[a_2 \cdots a_\ell]$. Then, by ψ_2 , $\text{ftp}^{\mathfrak{B}^-}[a_2 \cdots a_{\ell+1}] = \xi|_{[2, \ell+1]}$, where ξ is some fluted $(\ell+1)$ -type satisfying $\xi \models \bigwedge_{r \in R'} \alpha_r$. Again by Lemma 3.2.2 we have $\text{ftp}^{\mathfrak{C}}[c_2 \cdots c_{\ell+1}] = \xi|_{[2, \ell+1]}$. Using similar arguments as before, we can consistently assign $\text{ftp}^{\mathfrak{C}}[\bar{c}] := \xi$ thus

ensuring that $\mathfrak{C} \models \gamma_r[\bar{c}]$ or $\mathfrak{C} \not\models \alpha_r[c_1 \cdots c_\ell]$ for each $r \in R$. Repeating the assignment for each $(\ell+1)$ -tuple as described above will result in $\mathfrak{C} \models \varphi$. \square

Let us take stock of the previous few lemmas. Take φ_ℓ to be an fl^ℓ -sentence over the signature Σ_ℓ . We may, by Lemma 3.1.1, assume that it is in normal-form. By computing the formula $\varphi_{\ell-1} := \psi$ as described above, we have, by Lemmas 3.2.1 and 3.2.3, that $\varphi_{\ell-1}$ is satisfiable if and only if φ_ℓ is. To determine the computational resources needed for deciding satisfiability of φ_ℓ we first compute the size of $\varphi_{\ell-1}$. In the construction of $\varphi_{\ell-1}$ we consider elements of $\text{FTP}_{\ell-1}^{\Sigma_\ell}$, $\text{FTP}_\ell^{\Sigma_\ell}$, and all possible subsets of R and T . Since $|R| \leq \|\varphi_\ell\|$ it should be clear that there are no more than $2^{\|\varphi_\ell\|}$ subsets of R (with the same bound holding for $|T|$). As for the number of fluted ℓ - and $(\ell-1)$ -types, notice that a single one is of size at most $|\Sigma_\ell|$. Clearly, $|\text{FTP}_{\ell-1}^{\Sigma_\ell}|$ and $|\text{FTP}_\ell^{\Sigma_\ell}|$ are at most $2^{|\Sigma_\ell|} \leq 2^{\|\varphi_\ell\|}$. Thus, it is easy to verify that $\|\varphi_{\ell-1}\|$ is bounded by $2^{O(\|\varphi_\ell\|)}$. When talking about the size of the structures, we see, by Lemma 3.2.3, that if $\varphi_{\ell-1}$ is satisfiable in, say \mathfrak{B} , then φ_ℓ is satisfiable in \mathfrak{C} which satisfies $|C| \leq |B| \cdot |T| \leq |B| \cdot \|\varphi_\ell\|$.

We can now run the variable reduction procedure on $\varphi_{\ell-1}$ and subsequent formulas. What we obtain is a sequence of equisatisfiable sentences $\varphi_\ell, \dots, \varphi_2$ each exponentially larger than the last. Since φ_2 is in \mathcal{FO}^2 , we can, by Theorem 2.2.2, determine the satisfiability status of φ_2 , and thus also φ_ℓ , in non-deterministic time $2^{O(\|\varphi_2\|)}$. To better grasp this bound let us define $\mathfrak{t} : \mathbb{N} \times \mathbb{N} \rightarrow \mathbb{N}$ as:

$$\begin{aligned} \mathfrak{t}(0, n) &:= n, \\ \mathfrak{t}(m+1, n) &:= 2^{\mathfrak{t}(m, n)}. \end{aligned}$$

That is, $\mathfrak{t}(m, n)$ is a tower of m exponentiations ending with the final term n . A straightforward inductive argument reveals that $\|\varphi_i\|$ is bounded by $\mathfrak{t}(\ell-i, O(\|\varphi_\ell\|))$ for $i \in [2, \ell]$. Thus, by theorem 2.2.2, we may determine satisfiability of φ_ℓ in non-deterministic time $\mathfrak{t}(\ell-1, O(\|\varphi_\ell\|))$. In addition, by Lemma 2.2.1, we have that if φ_2 is satisfiable, it is satisfiable in a structure of size at most $2^{\|\varphi_2\|}$. We see, again by induction, that φ_ℓ is then satisfiable in a structure of size at most $2^{\|\varphi_2\|} \cdot \|\varphi_3\| \cdots \|\varphi_\ell\|$. But then, using the result of our previous induction, we have that this structure is no larger than $\mathfrak{t}(\ell-1, O(\|\varphi_\ell\|))$.

We can, however, do better complexity wise. In [Pratt-Hartmann et al., 2019, Lemma 7] it was shown that satisfiable fl^3 -sentences such as φ_3 have a model of size $2^{O(\|\varphi_3\|)}$. Thus, $\text{Sat}(fl^3)$ resides in non-deterministic time $2^{O(\|\varphi_3\|)}$. By stopping

our variable reduction procedure a step early we may then check the satisfiability of status φ_ℓ in non-deterministic time $t(\ell-2, O(\|\varphi_\ell\|))$ thus proving the following:

Theorem 3.2.4. [Pratt-Hartmann et al., 2019, Theorem 12] *Given $\ell \geq 3$, the satisfiability problem (= finite satisfiability problem) for fl^ℓ is in $(\ell-2)$ -NEXPTIME.*

Noting that $\text{Sat}(fl^\ell)$ is $\lfloor \ell/2 \rfloor$ -NEXPTIME-hard [Pratt-Hartmann et al., 2019, Theorem 2] we may conclude:

Theorem 3.2.5. [Pratt-Hartmann et al., 2019, Theorem 13] *The satisfiability problem (= finite satisfiability problem) for fl is TOWER-complete.*

It is understood that the variable reduction described above only works because, in Lemma 3.2.3, we can build a model of the original sentence. It is worthwhile to then ask how “well-behaved” can this model be? The answer, as we will find out in the following section, is “quite well-behaved”. Taking $\mathfrak{C} \models \varphi$ as in Lemma 3.2.3 and any $\bar{c} \in C^\ell$ we leave the reader to ponder the following: when can two distinct ℓ -tuples $c\bar{c}$ and $c'\bar{c}$ over C have $d \in C$ as a witness for the same existential requirement of φ ? The answer is somewhat obfuscated by the choice of pseudo-witnesses and the “blowing-up” of the model.

3.3 Bibliographic Notes

The origins of *argument-sequence logics* reach back to W. V. Quine’s work on *homogenous m-adic formulas*; that is formulas that feature m -ary predicates, atoms with the variable sequence $x_i \cdots x_{i+m-1}$, and quantification following the sequence x_1, x_2, \dots . In [Quine, 1969] W. V. Quine explained how the decision procedure for the satisfiability problem of the *monadic fragment* can be reworked to accommodate his language. It is clear that at the time W. V. Quine was looking for the *limits of decision* of first-order logic with restricted access to *predicate functors* (read: variable permutations) when forming formulas.

Moving away from the fragment of homogenous m -adic formulas to more expressive languages we turn to what is known as the *fluted fragment* which has a somewhat confusing history. The first use of the term *fluted* (in the form of *fluted schemata*) seems to originate from the work of W. V. Quine in [Quine, 1976]. However, in terms of the fluted fragment, it is the more expressive definition given W. C. Purdy in [Purdy, 1996] that appears to have stuck. It is now known that

what W. V. Quine meant by fluted schemata is the *ordered fragment* – a PSPACE-complete (in terms of satisfiability) argument-sequence logic rediscovered by A. Herzig [Herzig, 1990]. The definition of the fluted fragment as defined by W. C. Purdy is the one considered in the thesis.

It is somewhat unfortunate that the mismatch in definitions of W. V. Quine and W. C. Purdy is only the start of the fragments confusing story. In recent years W. C. Purdy’s results concerning “nice” features of the fluted fragment have been refuted. For instance, the result in [Purdy, 2002] claiming that the fluted fragment holds the *craig interpolation property* has been refuted by B. Bednarczyk and R. Jaakkola in [Bednarczyk and Jaakkola, 2022]. Similarly, the claim that the satisfiability problem for the fluted fragment rests in NEXPTIME does not hold as evident by I. Pratt-Hartmann’s, W. Szostak’s and L. Tendera’s result first in [Pratt-Hartmann et al., 2016] and later in the journal article [Pratt-Hartmann et al., 2019]. Indeed, as mentioned in the section before, the satisfiability problem for fl is TOWER-hard.

Note that there are still gaps in our knowledge of the fluted fragment. In [Pratt-Hartmann et al., 2019] it was shown that, for $\ell \geq 3$, $\text{Sat}(fl^\ell)$ is in $(\ell-2)$ -NEXPTIME and $\lfloor \ell/2 \rfloor$ -NEXPTIME-hard. This leaves a gap in complexity for ℓ as small as 5. The complexity gap persists at the moment of writing and it is not at all clear which one of the two bounds is lax.

Chapter 4

The Fluted Fragment with Periodic Counting

The following chapter is an expanded version of Sections 1–4, and 6 of the conference paper [Kojelis, 2025]. All results presented are due to the Ph. d. candidate.

In this chapter we establish that the class of models of satisfiable sentences in the fluted fragment with periodic counting, i.e. \mathcal{FLPC} , always contain a “nice” structure in which elements behave (in a sense that we will make clear) *homogeneously*. Utilising this new-found behaviour we will show that the fluted fragment extended with periodic counting quantifiers has a decidable (finite) satisfiability problem. Intriguingly, even though periodic counting quantifiers generalise standard counting quantifiers, our methodology allows us to dispense with Presburger quantifiers, which were required to establish decidability of satisfiability for \mathcal{FL} with standard counting [Pratt-Hartmann, 2021].

As will become apparent in further chapters of this thesis, the satisfiability problem for the fluted fragment with counting extensions becomes undecidable when minimal syntactic relaxations are allowed. We highlight that the work in this chapter is closely related to [Benedikt et al., 2024] in which decidability of satisfiability is established for \mathcal{FO}_{Pres}^2 – the two-variable fragment with periodic counting – but without a sharp complexity-theoretic bound (NEXPTIME-hard, but in 2-NEXPTIME). Our homogeneity conditions, which stem from the unidirectional nature of fluted logics, allow us to establish NEXPTIME-completeness for the finite and general satisfiability problems of \mathcal{FLPC}^2 .

The current chapter is structured as follows. We formally define the fluted fragment with periodic counting in Section 4.1. Section 4.2 gives homogeneity

results, which we use to show that the (finite) satisfiability problem for the two-variable subfragment of \mathcal{FLPC} is decidable. Section 4.3 presents a generalised homogeneity condition which we use to show that the (finite) satisfiability problem for multi-variable subfragments of \mathcal{FLPC} is decidable. In fact, the satisfiability problems turn to be TOWER-complete no variable bound is imposed. Section 4.4 argues for the usefulness of homogenous models when analysing fluted languages.

4.1 Preliminaries

We use \mathbb{N} to denote the set of non-negative integers $\{0, 1, 2, \dots\}$, and \mathbb{N}^* to denote \mathbb{N} along with the first infinite cardinal; i.e. $\mathbb{N}^* = \mathbb{N} \cup \{\aleph_0\}$. When $n, p \in \mathbb{N}$ we write n^{+p} for the *linear set* $\{n + ip \mid i \in \mathbb{N}\}$. In the extended integers \mathbb{N}^* , the cardinal \aleph_0 is the maximum element under the canonical ordering “ $<$ ” and

- $0 \cdot \aleph_0 = \aleph_0 \cdot 0 = 0$;
- $n + \aleph_0 = \aleph_0 + n = \aleph_0$ for all $n \in \mathbb{N}^*$; and
- $n \cdot \aleph_0 = \aleph_0 \cdot n = \aleph_0$ for all $n \in \mathbb{N}^* \setminus \{0\}$.

A *linear Diophantine inequation* is an expression of the form

$$a_1 v_1 + \dots + a_n v_n + b \bowtie c_1 v_1 + \dots + c_n v_n + d,$$

where $(a_i)_{i=1}^n, b, (c_i)_{i=1}^n, d$ are constant values taken from \mathbb{N}^* , $\bar{v} = v_1, \dots, v_n$ is a vector of variables, and \bowtie is any of the relations $=, \neq, \leq, <, \geq, >$ (each interpreted as one would assume). It is known that when the cardinal \aleph_0 is disallowed, a solution for a set of such inequations may be found in NP TIME [Papadimitriou, 1981]. The picture does not change when \aleph_0 is permitted as a solution and/or constant. Indeed, we may reduce the problem of finding a solution over \mathbb{N}^* to that of finding it over \mathbb{N} as follows. First guess which variables should be mapped to \aleph_0 and which should have a finite value. Then, check that each inequation featuring a variable assigned \aleph_0 holds and discard them. What will be left is a system of inequations with constants in \mathbb{N} and in variables assumed to be finite. See [Pratt-Hartmann, 2023, Chapter 7.4] for greater detail. We allow systems of inequations to contain disjunctions.

Recall that a word \bar{a} over A is a tuple $\bar{a} = a_1 \cdots a_n$, where $a_i \in A$ for each $i \in [1, n]$. In case $n = 1$, we often write a instead of \bar{a} . By \tilde{a} we mean $a_n \cdots a_1$. If \bar{a} and \bar{b} are words we write $\bar{a}\bar{b}$ for the concatenation of the two.

Now, take some structure \mathfrak{A} and an ℓ -tuple \bar{a} of elements from A . Suppose $B = \{b \in A \mid \mathfrak{A} \models \varphi[\bar{a}b]\}$ for some first-order formula $\varphi(x_1 \cdots x_{\ell+1})$. Fixing $n, p \in \mathbb{N}$ we extend the syntax of first-order logic with *periodic counting quantifier* $\exists_{[n+p]}$. Semantically, $\mathfrak{A}, \bar{a} \models \exists_{[n+p]} x_{\ell+1} \varphi$ if and only if $|B| \in n+p$. We refrain from further generalisation to *ultimately periodic counting quantifiers* $\exists_{[n_1+p_1 \cup \dots \cup n_k+p_k]}$ (as in [Benedikt et al., 2024]) as they can be expressed as a disjunction of formulas using periodic counting quantifiers. Thus, a sentence such as “Every orchestra hires an even number of people to play violin” may be written in a language with periodic counting as follows:

$$\forall x_1 \left(\text{orc}(x_1) \rightarrow \exists_{[0+2]} x_2 \left(\text{per}(x_2) \wedge \exists x_3 \left(\text{vio}(x_3) \wedge \text{hire}(x_1 x_2 x_3) \right) \right) \right). \quad (\mathcal{FLPC}\text{-ex1})$$

Fix $\ell \geq 0$. We say that an expression $\mathbf{p}(\bar{x})$ is a fluted ℓ -atom if \mathbf{p} is a predicate letter of arity at most ℓ or the designated equality symbol “=”, and \bar{x} is a suffix of the variable sequence $x_1 \cdots x_\ell$. Notice that, in case \bar{x} is the empty word, \mathbf{p} is a proposition letter. Formally, the fluted fragment with periodic counting is the union of sets of formulas $\mathcal{FLPC}^{[\ell]}$ defined by simultaneous induction as follows:

- (i) each fluted ℓ -atom is in $\mathcal{FLPC}^{[\ell]}$;
- (ii) $\mathcal{FLPC}^{[\ell]}$ is closed under Boolean combinations;
- (iii) if $\varphi \in \mathcal{FLPC}^{[\ell+1]}$, then $\exists_{[n+p]} x_{\ell+1} \varphi$ is in $\mathcal{FLPC}^{[\ell]}$ for every $n, p \in \mathbb{N}$.

We write $\mathcal{FLPC} = \bigcup_{\ell \geq 0} \mathcal{FLPC}^{[\ell]}$ for the set of all fluted formulas with periodic counting and define the ℓ -variable fluted fragment with periodic counting to be the set $\mathcal{FLPC}^\ell := \mathcal{FLPC} \cap \mathcal{FO}^\ell$. We will implicitly restrict attention to signatures which feature no function and/or constant symbols. Lastly, we use $\forall x \varphi$ interchangeably with $\exists_{[0+0]} x \neg \varphi$ whenever convenient.

As mentioned in Chapter 3, variables in fluted logics convey no meaningful information. We will thus again employ variable-free notation. As an example, the formula ($\mathcal{FLPC}\text{-ex1}$) can be written as:

$$\forall \left(\text{orc} \rightarrow \exists_{[0+2]} \left(\text{per} \wedge \exists \left(\text{vio} \wedge \text{hire} \right) \right) \right),$$

without ambiguity (up to a shift of variable indices).

Fix a first-order formula φ with periodic counting quantifiers. We assume that numeric values are encoded in binary and write $\|\varphi\|$ for the number of symbols used in φ . Moreover, we will assume that linear sets $\{n+ip \mid i \in \mathbb{N}\}$ are compactly represented by the pair of integers n and p (again, in binary). We point out that the signature of φ (denoted $\text{sig}(\varphi)$) is no larger than $\|\varphi\|$.

Now, let φ be a formula in $\mathcal{FLPC}^{\ell+1}$. We say that φ is in *normal-form* if it takes the following shape

$$\bigwedge_{r \in R} \forall^\ell (\alpha_r \rightarrow \exists_{[n_r^{+p_r}]} \gamma_r) \wedge \bigwedge_{t \in T} \forall^\ell (\beta_t \rightarrow \neg \exists_{[n_t^{+p_t}]} \delta_t), \quad (\mathcal{FLPC}\text{-nmf})$$

where R, T are finite sets of indices, each α_r, β_t is a quantifier-free \mathcal{FLPC} -formula in ℓ variables, each γ_r, δ_t is a quantifier-free \mathcal{FLPC} -formula in $\ell+1$ variables, and each $n_r^{+p_r}, n_t^{+p_t}$ is a linear set. Using standard rewriting techniques we have:

Lemma 4.1.1. *Suppose φ is an $\mathcal{FLPC}^{\ell+1}$ -sentence. Then, we may compute, in polynomial time, an equisatisfiable normal-form $\mathcal{FLPC}^{\ell+1}$ -sentence ψ .*

Proof. We start by assuming that φ contains no existential or universal quantifiers. No loss of generality follows this supposition as every formula of the form $\exists \chi$ is equivalent to $\neg \forall \neg \chi$ whilst formulas $\forall \chi$ are equivalent to $\exists_{[0+0]} \neg \chi$. Writing $\varphi_0 := \varphi$, take any subformula $\theta := \exists_{[n+p]} \chi$ of φ_0 , where χ is quantifier-free. Supposing that $\theta \in \mathcal{FLPC}^{[k]}$, let \mathbf{p} be a fresh k -ary predicate. We write ψ_1 as:

$$\forall^\ell (\mathbf{p} \rightarrow \exists_{[n+p]} \chi) \wedge \forall^\ell (\neg \mathbf{p} \rightarrow \neg \exists_{[n+p]} \chi)$$

and define φ_1 to be φ_0 but with θ replaced by \mathbf{p} . Clearly, $\varphi_1 \wedge \psi_1 \models \varphi_0$. Conversely, if $\mathfrak{A} \models \varphi_0$, we may expand \mathfrak{A} to \mathfrak{A}' by setting $\bar{a} \in \mathbf{p}^{\mathfrak{A}'}$ if $\mathfrak{A}, \bar{a} \models \theta$ for each $\bar{a} \in A^k$. Then, $\mathfrak{A}' \models \varphi_1 \wedge \psi_1$ as required. Processing φ_1 and subsequent sentences in the same way, we are left with a sentence φ_m composed solely of proposition letters and sentences ψ_1, \dots, ψ_m . The formula $\psi_1 \wedge \dots \wedge \psi_m \wedge \forall^\ell (\top \rightarrow \exists_{[0+0]} \neg \varphi_m)$ is then of the required form. \square

Notice that the negation before the periodic counting quantifier in the second conjunct of $(\mathcal{FLPC}\text{-nmf})$ is not moved-inwards. This deliberate so as to avoid computing complements of linear sets, which may be of exponential size as a function of $\|\varphi\|$.

Now, let Σ again be some finite, constant- and function-free signature. Recall that a *fluted* $(\ell+1)$ -*type* is a maximal consistent set of possibly negated $(\ell+1)$ -fluted atoms with symbols from $\Sigma \cup \{=\}$. Given a Σ -structure \mathfrak{A} , each $(\ell+1)$ -tuple $a\bar{a}b$ over A realises a fluted $(\ell+1)$ -type denoted by $\text{ftp}^{\mathfrak{A}}[a\bar{a}b]$. Intuitively, for each $\mathbf{p} \in \Sigma$, we have that the $(\ell+1)$ -fluted atom $\mathbf{p}(x_k \cdots x_{\ell+1})$ is in $\text{ftp}^{\mathfrak{A}}[a\bar{a}b]$ if and only if $\mathfrak{A}, a\bar{a}b \models \mathbf{p}(x_k \cdots x_{\ell+1})$. We invite the reader to view the ℓ -tuple $a\bar{a}$ as *emitting* ξ , and the ℓ -tuple $\bar{a}b$ as *absorbing* ξ . Write $\xi|_{[2, \ell+1]}$ for the fluted ℓ -type obtained by deleting entries in ξ of arity greater than ℓ and decrementing variable indices by 1. We say that a fluted ℓ -type ζ is the endpoint of a fluted $(\ell+1)$ -type ξ if $\xi|_{[2, \ell+1]} = \zeta$. Again write $\text{FTP}_{\ell+1}^{\Sigma}$ for the set of all fluted $(\ell+1)$ -types over Σ .

A fluted ℓ -profile is a function mapping fluted $(\ell+1)$ -types to cardinal numbers. Each ℓ -tuple $a\bar{a}$ in a given Σ -structure realises a fluted ℓ -profile ρ denoted by $\text{fpr}^{\mathfrak{A}}[a\bar{a}]$. Formally, ρ is defined on $\xi \in \text{FTP}_{\ell+1}^{\Sigma}$ as follows:

$$\rho(\xi) = |\{b \in A \mid \text{ftp}_{i+1}^{\mathfrak{A}}[a\bar{a}b] = \xi\}|.$$

If ψ is a quantifier-free $\mathcal{FLPC}^{\ell+1}$ -formula, we write $\rho \models \exists_{[n+p]}\psi$ if and only if $\sum_{\xi \in \text{FTP}_{\ell+1}^{\Sigma}}^{\xi \models \psi} \rho(\xi) \in n+p$. Clearly, $\rho \models \exists_{[n+p]}\psi$ if and only if $\mathfrak{A}, a\bar{a} \models \exists_{[n+p]}\psi$.

It is easy to verify that every sentences of \mathcal{FLPC} is logically equivalent to a computable formula in $\mathcal{L}_{\omega_1, \omega}$. Since the downward Löwenheim-Skolem Theorem holds for computable sentences¹ of $\mathcal{L}_{\omega_1, \omega}$ (folklore, see [Keisler, 1971, p. 69]), we will implicitly assume that each structure given in the sequel is countable. Note that, in \mathcal{FLPC} – as opposed to \mathcal{FL} and even \mathcal{FLL} – the finite model property fails, as $\neg \exists_{[0+1]}\top$ is an axiom of infinity.

4.2 The Two-Variable Subfragment

In this section we restrict attention to the two-variable fluted fragment with periodic counting. To achieve decidability of (finite) satisfiability we first specify what kind of “nice” structures we will be looking for. For the rest of the section let us fix Σ to be a constant- and function-free signature. Now, take any Σ -structure \mathfrak{A} and $\zeta \in \text{FTP}_1^{\Sigma}$. For convenience, we write A_{ζ} for the set of all elements $a \in A$ with $\text{ftp}^{\mathfrak{A}}[a] = \zeta$. We say that ζ is *globally homogeneous* in \mathfrak{A} if $\text{fpr}^{\mathfrak{A}}[a] = \text{fpr}^{\mathfrak{A}}[b]$ for each $a, b \in A_{\zeta}$. That is to say, ζ is globally homogeneous in \mathfrak{A}

¹Here conjunctions and disjunctions are formed over recursively enumerable sets.

if, for each $\xi \in \text{FTP}_2^\Sigma$, the number of fluted 2-types ξ emitted is the same cardinal number regardless of the emitting element picked amongst A_ζ . The structure \mathfrak{A} is *globally homogeneous* if each $\zeta \in \text{FTP}_1^\Sigma$ is globally homogeneous in \mathfrak{A} .

For the remainder of the section fix some normal-form \mathcal{FLPC}^2 -sentence φ . We claim that if φ is satisfiable, then it is satisfiable in a globally homogeneous model. To see this, take some structure $\mathfrak{A} \models \varphi$ and a pair of elements $cd \in A^2$ that realised the fluted 2-type ξ in \mathfrak{A} . Since ξ consists of fluted formulas, it does not feature atoms of the form $\mathbf{p}(x_1)$ and $\mathbf{p}(x_2x_1)$. Referencing ξ only, we can deduce what formulas are satisfied by the pair dc in \mathfrak{A} if and only if $c=d$. Continuing with $c \neq d$, if we were to alter the fluted 2-type of cd in \mathfrak{A} , the set of quantifier-free \mathcal{FLPC}^2 -formulas satisfied by $c'd' \in A^2 \setminus \{cd\}$ in \mathfrak{A} would not change.

Taking a step back, pick some $\zeta \in \text{FTP}_1^\Sigma$ and recall that $A_\zeta \subseteq A$ is the set of all elements realising the 1-type ζ in \mathfrak{A} . Writing $\mathfrak{B} := \mathfrak{A}$ we will redefine 2-types emitted by elements of A_ζ in such a way that makes ζ globally homogeneous in \mathfrak{B} whilst maintaining the fluted 2-types between pairs $(A \setminus A_\zeta) \times A$ as in \mathfrak{A} . Let us fix any $a \in A_\zeta$ and write $\rho = \text{fpr}^\mathfrak{A}[a]$. The element a and profile ρ picked will be an example as to how the rest of A_ζ should form fluted 2-types with other elements of the model. Taking any $b \in A_\zeta \setminus \{a\}$ we allow b to impersonate a in \mathfrak{B} by rewiring the fluted 2-type $\text{ftp}^\mathfrak{B}[bc]$ to be $\text{ftp}^\mathfrak{A}[ac]$ for each $c \in A \setminus \{a, b\}$ and, additionally, by setting $\text{ftp}^\mathfrak{B}[ba]$ to be $\text{ftp}^\mathfrak{A}[ab]$ and $\text{ftp}^\mathfrak{B}[bb]$ to be $\text{ftp}^\mathfrak{A}[aa]$. Clearly, only fluted 2-types emitted by b were reconsidered in this procedure, thus pairs in $(A \setminus \{b\}) \times A$ satisfy the same quantifier-free \mathcal{FLPC}^2 -formulas as before. To see that \mathfrak{B} still models φ we need only show that b does not violate $\alpha_r \rightarrow \exists_{[n_r+p_r]}\gamma_r$ and $\beta_t \rightarrow \neg\exists_{[n_t+p_t]}\delta_t$ for each $r \in R$ and $t \in T$. By our rewiring procedure, we have that $\text{fpr}^\mathfrak{B}[a] = \rho = \text{fpr}^\mathfrak{B}[b]$. Thus, for each quantifier-free $\psi \in \mathcal{FLPC}^2$:

$$\mathfrak{A}, a \models \exists_{[n+p]}\psi \iff \rho \models \exists_{[n+p]}\psi \iff \mathfrak{B}, b \models \exists_{[n+p]}\psi$$

By our initial assumption that $\mathfrak{A} \models \varphi$, we have $\mathfrak{A}, a \models \alpha_r \rightarrow \exists_{[n_r+p_r]}\gamma_r$ and $\mathfrak{A}, a \models \beta_t \rightarrow \neg\exists_{[n_t+p_t]}\delta_t$ for each $r \in R$ and $t \in T$. Thus, $\mathfrak{B} \models \varphi$ as required.

Since only fluted 2-types emitted by b are considered, we can run this construction in parallel for each element in $A_\zeta \setminus \{a\}$. Clearly, this renders ζ globally homogeneous in the resulting model. Since elements in $A \setminus A_\zeta$ are left untouched by our rewiring, repeating the above for each $\zeta \in \text{FTP}_1^\Sigma$ yields:

Lemma 4.2.1. *Suppose φ is a satisfiable normal-form \mathcal{FLPC}^2 -sentence. Then,*

φ is satisfiable in a globally homogeneous model.

In globally homogeneous structures elements realising the same fluted 1-type are, in a sense, stripped away of their individuality as they all realise the same fluted 1-profile. When the globally homogeneous structure \mathfrak{A} is clear from context, we can unambiguously write ρ_ζ for the fluted 1-profile realised by each element of A_ζ (here $\zeta \in \text{FTP}_1^\Sigma$).

When considering the (finite) satisfiability problem for normal-form \mathcal{FLPC}^2 -sentences such as φ , we will confine ourselves to the search of globally homogeneous models. More precisely, we will produce a system of linear Diophantine inequations Ψ that has a solution over \mathbb{N}^* if and only if φ is satisfiable in a globally homogeneous model. For this purpose, let $(x_\zeta)_{\zeta \in \text{FTP}_1^\Sigma}$, $(y_{\zeta,\xi})_{\substack{\xi \in \text{FTP}_2^\Sigma \\ \zeta \in \text{FTP}_1^\Sigma}}$, $(i_{\zeta,r})_{\zeta \in \text{FTP}_1^\Sigma}^{r \in R}$, and $(j_{\zeta,t})_{\zeta \in \text{FTP}_1^\Sigma}^{t \in T}$ be sequences of variables. Intuitively, the value assigned to x_ζ will represent the number of elements realising the fluted 1-type ζ . The value of $y_{\zeta,\xi}$ is then the number of times the 2-type ξ is emitted by an element realising ζ . Lastly, $i_{\zeta,r}$ and $j_{\zeta,t}$ act as periodic counters for elements realising ζ when considering linear sets $n_r^{+p_r}$ and $n_t^{+p_t}$. To be more precise, we will build a system of equations Ψ satisfying the following:

- If Ψ has a satisfying assignment π , then there is a model $\mathfrak{A} \models \varphi$ such that $|A_\zeta| = \pi(x_\zeta)$ and $\rho_\zeta(\xi) = \pi(y_{\zeta,\xi})$ for each $\zeta \in \text{FTP}_1^\Sigma$ and $\xi \in \text{FTP}_2^\Sigma$, and
- If \mathfrak{A} is a globally homogeneous model of φ , then $\pi^\mathfrak{A}$ is a satisfying assignment for Ψ obtained by setting the following for all $\zeta \in \text{FTP}_1^\Sigma$, $\xi \in \text{FTP}_2^\Sigma$, $r \in R$, and $t \in T$:²

$$\begin{aligned} - \pi^\mathfrak{A}(x_\zeta) &:= |A_\zeta|, \\ - \pi^\mathfrak{A}(y_{\zeta,\xi}) &:= \rho_\zeta(\xi), \\ - \pi^\mathfrak{A}(i_{\zeta,r}) &:= \left(\sum_{\xi' \in \text{FTP}_2^\Sigma}^{\xi' \models \gamma_r} \rho_\zeta(\xi') - n_r \right) / p_r, \\ - \pi^\mathfrak{A}(j_{\zeta,t}) &:= \left\lfloor \left(\sum_{\xi' \in \text{FTP}_2^\Sigma}^{\xi' \models \delta_t} \rho_\zeta(\xi') - n_t \right) / p_t \right\rfloor. \end{aligned}$$

We proceed by showing that the latter assignment satisfies our (yet to be defined) system of inequations:

$$\Psi_1 \cup \dots \cup \Psi_6. \tag{\Psi}$$

²In case p_r (resp. p_t) is 0, we allow $i_{\zeta,r}$ (resp. $j_{\zeta,t}$) to take any integer value.

Given any $\mathfrak{A} \models \varphi$ we have that the domain $A = \bigcup_{\zeta \in \text{FTP}_1^\Sigma} A_\zeta$ is non-empty. The following singleton set is thus satisfied by the assignment $\pi^{\mathfrak{A}}(x_\zeta) := |A_\zeta|$:

$$\left\{ \sum_{\zeta \in \text{FTP}_1} x_\zeta \geq 1 \right\}. \quad (\Psi_1)$$

Additionally, picking any element $a \in A_\zeta$ (for any $\zeta \in \text{FTP}_1^\Sigma$) and any $\eta \in \text{FTP}_\eta^\Sigma$, we have that the number of fluted 2-types emitted by a to A_η must be $|A_\eta|$. Assuming that \mathfrak{A} is globally homogeneous, we may fixate on the fact that the shared profile ρ_ζ of A_ζ has exactly $|A_\eta|$ witnesses for fluted 2-types with the endpoint η , or, more formally, $\sum_{\xi \in \text{FTP}_2^\Sigma}^{\xi|_{[2,2]}=\eta} \rho_\zeta(\xi) = |A_\eta|$. Thus, the assignments $\pi^{\mathfrak{A}}(x_\eta) := |A_\eta|$ and $\pi^{\mathfrak{A}}(y_{\zeta,\xi}) := \rho_\zeta(\xi)$ satisfies the following set of inequations:

$$\left\{ x_\zeta \neq 0 \rightarrow \sum_{\xi \in \text{FTP}_2^\Sigma}^{\xi|_{[2,2]}=\eta} y_{\zeta,\xi} = x_\eta \mid \zeta, \eta \in \text{FTP}_1^\Sigma \right\}. \quad (\Psi_2)$$

Of course, under the supposition that $\zeta \models \alpha_r$ for some $r \in R$ and $\zeta \in \text{FTP}_1^\Sigma$, we have that $\rho_\zeta \models \exists_{[n_r+p_r]} \gamma_r$. Clearly, $k := \sum_{\xi \in \text{FTP}_2^\Sigma}^{\xi \models \gamma_r} \rho_\zeta(\xi)$ must be a member of the linear set $n_r^{+p_r}$. Thus, there is some $i_{\zeta,r} \in \mathbb{N}$ such that $k = n_r + i_{\zeta,r} p_r$. Turning the equation around we get that $i_{\zeta,r} = (k - n_r) / p_r$. Recalling that $\pi^{\mathfrak{A}}(y_{\zeta,\xi}) := \rho_\zeta(\xi)$ for all $\xi \in \text{FTP}_2^\Sigma$, we have that the following is satisfied by our assignments:

$$\left\{ x_\zeta \neq 0 \rightarrow \sum_{\xi \in \text{FTP}_2^\Sigma}^{\xi \models \gamma_r} y_{\zeta,\xi} = n_r + i_{\zeta,r} p_r \mid r \in R, \zeta \in \text{FTP}_1^\Sigma \text{ s.t. } \zeta \models \alpha_r \right\}. \quad (\Psi_3)$$

On the other hand, supposing $\zeta \models \beta_t$ for some $t \in T$, we have that $\rho_\zeta \not\models \exists_{[n_t+p_t]} \delta_r$, thus leaving $k := \sum_{\xi \in \text{FTP}_2^\Sigma}^{\xi \models \gamma_r} \rho_\zeta(\xi)$ outside the linear set $n_t^{+p_t}$. Notice that this happens when one of the following conditions is met:

1. $k < n_t$; or
2. $p_t \neq 0$ and $k > m$ for all $m \in n_t^{+p_t}$ (which only happens when $k = \aleph_0$); or
3. $p_t = 0$ and $k > n_t$; or
4. for some $m \in n_t^{+p_t}$ we have $m < k < m + p_t$.

Note that the listed conditions are exhaustive. We translate the requirements 1–4 into four functions $\Theta_1, \dots, \Theta_4$ which map fluted 1-types paired with indices

in T to linear equations. The functions are then defined as follows:

- $\Theta_1(\zeta, t) := \sum_{\xi \in \text{FTP}_2^\Sigma}^{\xi \models \delta_t} y_{\zeta, \xi} < n_t$,
- $\Theta_2(\zeta, t) := \sum_{\xi \in \text{FTP}_2^\Sigma}^{\xi \models \delta_t} y_{\zeta, \xi} = \aleph_0$,
- $\Theta_3(\zeta, t) := (p_t = 0) \wedge (n_t < \sum_{\xi \in \text{FTP}_2^\Sigma}^{\xi \models \delta_t} y_{\zeta, \xi})$,
- $\Theta_4(\zeta, t) := n_t + j_{\zeta, t} p_t < \sum_{\xi \in \text{FTP}_2^\Sigma}^{\xi \models \delta_t} y_{\zeta, \xi} < n_t + (j_{\zeta, t} + 1) p_t$.

Since k adheres to at least one of the four conditions, we write the following clauses for eligible fluted 1-types:

$$\left\{ x_\zeta \neq 0 \rightarrow \bigvee_{i \in [1,4]} \Theta_i(\zeta, t) \mid t \in T \text{ and } \zeta \in \text{FTP}_1^\Sigma \text{ such that } \zeta \models \beta_t \right\}. \quad (\Psi_4)$$

To verify that $\pi^{\mathfrak{A}}(j_{\zeta, \xi}) = \lfloor (k - n_t)/p_t \rfloor$ is indeed a satisfying assignment for Ψ_4 , we need only consider case 4. For this, assume that $m < k < m + p_t$ for some $m \in n_t^{+p_t}$. We can thus write $m = n_t + j p_t$. Since $\pi^{\mathfrak{A}}(j_{\zeta, \xi}) = \lfloor (k - n_t)/p_t \rfloor$, we conclude that $\pi^{\mathfrak{A}}(j_{\zeta, \xi}) = j$ thus satisfying $\Theta_4(\zeta, t)$.

Reflecting on the semantics of the equality predicate, we see that for any $\zeta \in \text{FTP}_1^\Sigma$ in any globally homogeneous model $\mathfrak{A} \models \varphi$ there is exactly one $\xi \in \text{FTP}_2^\Sigma$ featuring the non-negated equality symbol and having $\rho_\zeta(\xi) \neq 0$. More precisely, $\rho_\zeta(\xi) = 1$ and the endpoint of ξ is ζ . We respect this condition by writing:

$$\left\{ y_{\zeta, \xi} = 0 \mid \zeta \in \text{FTP}_1^\Sigma, \xi \in \text{FTP}_2^\Sigma \text{ s.t. } = \in \xi, \xi|_{[2,2]} \neq \zeta \right\} \cup \left\{ \sum_{\xi \in \text{FTP}_2^\Sigma}^{\equiv \xi} y_{\zeta, \xi} = 1 \mid \zeta \in \text{FTP}_1^\Sigma \right\}. \quad (\Psi_5)$$

Finally, as a technicality, we forbid the periodic counters $(i_{\zeta, r})_{\zeta \in \text{FTP}_1^\Sigma}^{r \in R}$ and $(j_{\zeta, t})_{\zeta \in \text{FTP}_1^\Sigma}^{t \in R}$ from taking the value \aleph_0 :

$$\left\{ i_{\zeta, r}, j_{\zeta, t} < \aleph_0 \mid \zeta \in \text{FTP}_1^\Sigma \text{ and } r \in R, t \in T \right\}. \quad (\Psi_6)$$

Putting everything together, we have engineered a system of equations $\Psi = \Psi_1 \cup \dots \cup \Psi_6$ that is satisfied by extracting relevant cardinalities (as was done with $\pi^{\mathfrak{A}}$) from homogeneous models of φ . Thus, the following has been proved:

Lemma 4.2.2. *Suppose $\mathfrak{A} \models \varphi$ is a globally homogeneous model. Then, Ψ has a satisfying assignment.*

We now move on to the converse direction:

Lemma 4.2.3. *Suppose Ψ has a satisfying assignment. Then, there is a globally homogeneous model $\mathfrak{A} \models \varphi$.*

Proof. Suppose that Ψ has a satisfying assignment π . We will build a globally homogeneous model \mathfrak{A} over the domain

$$A = \bigcup_{\zeta \in \text{FTP}_1^\Sigma} A'_\zeta, \text{ where each } A'_\zeta \text{ is a set of } \pi(x_\zeta) \text{ fresh elements.}$$

Intuitively, we wish that elements of A'_ζ realise the fluted 1-type ζ . We thus assign $\text{ftp}_1^{\mathfrak{A}}[a] := \zeta$ for all $a \in A'_\zeta$. By our notation, A_ζ is the set of all elements that realise the 1-type ζ in \mathfrak{A} . Clearly, $A_\zeta = A'_\zeta$ is of cardinality $\pi(x_\zeta)$. By Ψ_1 , the domain is non-empty.

Picking any $\zeta \in \text{FTP}_1^\Sigma$ and $a \in A_\zeta$ we now move on to the assignment of fluted 2-types. Take any $\eta \in \text{FTP}_1^\Sigma$ and let $S = \{\xi \in \text{FTP}_2^\Sigma \mid \xi|_{[2,2]} = \eta\}$, i.e. S is the set of all fluted 2-types containing the fluted 1-type η as an endpoint. By Ψ_2 and the assumption that A_ζ is non-empty, we have that $\sum_{\xi \in S} \pi(y_{\zeta,\xi}) = \pi(x_\eta) = |A_\eta|$. In case $\zeta \neq \eta$, equation Ψ_5 prohibits fluted 2-types that feature the (non-negated) equality literal. We set fluted 2-types between a and elements of A_η in any way that results in $|\{b \in A_\eta \mid \text{ftp}_2^{\mathfrak{A}}[ab] = \xi\}| = \pi(y_{\zeta,\xi})$ for each $\xi \in S$ (n.b. the exact configuration of fluted 2-types between a and elements of A_η is irrelevant as the fluted 2-type of ba for any $b \in A_\eta$ is not set in this process). In case $\zeta = \eta$ notice that by Ψ_5 there is exactly one $\xi^= \in S$ such that (i) $= \in \xi^=$, (ii) $\pi(y_{\zeta,\xi^=}) \geq 0$, and with (iii) $\xi^=|_{[2,2]} = \zeta$. By Ψ_5 again, we have that $\pi(y_{\zeta,\xi^=}) = 1$. We therefore set the fluted 2-types between a and $A_\zeta \setminus \{a\}$ for each $\xi \in S \setminus \{\xi^=\}$ as in the case before and, additionally, specify that $\text{ftp}_2^{\mathfrak{A}}[aa] := \xi^=$.

By repeating the fluted 2-type assignment for each element $a \in A$ and fluted 1-type $\eta \in \text{FTP}_1^\Sigma$ we are guaranteed that elements in A_ζ (where $\zeta = \text{ftp}_1^{\mathfrak{A}}[a]$) realise the fluted 1-profile $\rho_\zeta := \{\xi \mapsto \pi(y_{\zeta,\xi}) \mid \xi \in \text{FTP}_2^\Sigma\}$ thus producing a globally homogeneous structure.

We now claim that the resulting structure is a model of φ . Indeed, take any $a \in A$ with $\zeta = \text{ftp}_1^{\mathfrak{A}}[a]$ and suppose $\zeta \models \alpha_r$ for some $r \in R$. Let $S = \{\xi \in \text{FTP}_2^\Sigma \mid \xi \models \gamma_r\}$. By equations Ψ_3 and Ψ_6 , the sum $\sum_{\xi \in S} \pi(y_{\zeta,\xi})$ is a member of the linear

set $n_r^{+p_r}$. Since the element a is of fluted 1-type ζ , we have that it realises the profile ρ_ζ in \mathfrak{A} . By our construction, $\rho_\zeta(\xi) = \pi(y_{\zeta,\xi})$ for each $\xi \in \text{FTP}_2^\Sigma$. Thus, $\rho_\zeta \models \exists_{[n_r^{+p_r}]} \gamma_r$ which is equivalent to $\mathfrak{A}, a \models \exists_{[n_r^{+p_r}]} \gamma_r$ as required.

On the other hand, suppose $\zeta \models \beta_t$ for some $t \in T$. We claim that $\mathfrak{A}, a \not\models \exists_{[n_t^{+p_t}]} \delta_t$. To see this, let S be the set $\{\xi \in \text{FTP}_2^\Sigma \mid \xi \models \delta_t\}$. Writing $k = \sum_{\xi \in S} \pi(y_{\zeta,\xi})$ we take note of equations Ψ_4 and Ψ_6 , and conclude that one of the following conditions must be true:

1. k is smaller than the minimal element of $n_t^{+p_t}$; or
2. $n_t^{+p_t} \subseteq \mathbb{N}$ and $k = \aleph_0$; or
3. $p_t = 0$ and $k > n_t$; or
4. k is in between two consecutive elements of $n_t^{+p_t}$.

Whichever case it may be, we have that $k \notin n_t^{+p_t}$. Again, recalling that $\rho_\zeta(\xi) = \pi(y_{\zeta,\xi})$ for each $\xi \in \text{FTP}_2^\Sigma$, we conclude that $\rho_\zeta \not\models \exists_{[n_t^{+p_t}]} \delta_t$ which is equivalent to saying $\mathfrak{A}, a \not\models \exists_{[n_t^{+p_t}]} \delta_t$. \square

Given an \mathcal{FLPC}^2 -sentence φ we present a decision procedure for the (finite) satisfiability problem. Compute a normal-form formula ψ from φ as done in Lemma 4.1.1 and write the linear Diophantine equations Ψ (in regards to ψ). Now, guess a solution vector \bar{z} which can be done in non-deterministic polynomial time as a function of $\|\Psi\|$. If \bar{z} is indeed a solution for Ψ , accept, otherwise, reject. In the case of the finite satisfiability problem, prohibit \aleph_0 from being a solution in Ψ . Correctness of the procedure follows from the fact that, by Lemma 4.2.1, if ψ is satisfiable, then it is satisfiable in a globally homogeneous model. Combining this with Lemma 4.2.2 we have that if ψ is satisfiable, then Ψ has a solution. On the other hand, by Lemma 4.2.3, if Ψ has a solution, then ψ is satisfiable.

Noting that the satisfiability problem for \mathcal{FLC}^2 is NEXPTIME-hard [Pratt-Hartmann et al., 2019], and that $\|\Psi\|$ is bounded by a polynomial function on the number of different fluted 1- and 2-types (of which there are $2^{\|\varphi\|}$ many), we conclude the following:

Theorem 4.2.4. *The (finite) satisfiability problem of \mathcal{FLPC}^2 is NEXPTIME-complete.*

4.3 More Than Two Variables

We now generalise our results on homogeneity and decidability of satisfiability for higher-arity formulas of \mathcal{FLPC} . Thus, throughout this section, we will be working in the $(\ell+1)$ -variable sub-fragment of \mathcal{FLPC} , where $\ell \geq 2$ is fixed.

Firstly, we lift our homogeneity conditions to the multivariable setting. Suppose \mathfrak{A} is a Σ -structure and take any $(\ell-1)$ -tuple \bar{a} from A and $\zeta \in \text{FTP}_\ell^\Sigma$. Let $A_{\zeta \leftarrow \bar{a}}$ be the set $\{a \in A \mid \text{ftp}^\mathfrak{A}[a\bar{a}] = \zeta\}$. We say that ζ is \bar{a} -homogeneous in \mathfrak{A} if for each $a, a' \in A_{\zeta \leftarrow \bar{a}}$ and all $b \in A$ we have that $\text{ftp}^\mathfrak{A}[a\bar{a}b] = \text{ftp}^\mathfrak{A}[a'\bar{a}b]$. That is to say, $\bar{a}b$ absorbs the same fluted $(\ell+1)$ -type from each ℓ -tuple $a\bar{a}$ that realises the fluted ℓ -type ζ . If each $\zeta \in \text{FTP}_\ell^\Sigma$ is \bar{a} -homogeneous in \mathfrak{A} , then we say that the $(\ell-1)$ -tuple \bar{a} is *homogeneous* in \mathfrak{A} . Finally, if each $(\ell-1)$ -tuple \bar{a} is homogeneous in \mathfrak{A} , then we say that \mathfrak{A} is *locally ℓ -homogeneous*.

When considering satisfiable normal-form $\mathcal{FLPC}^{\ell+1}$ -sentences we can, without loss of generality, confine ourselves to locally ℓ -homogeneous structures. To see this, fix some normal-form $\mathcal{FLPC}^{\ell+1}$ -sentence φ and suppose $\mathfrak{A} \models \varphi$. Now, take $\bar{a} \in A^{\ell-1}$ and distinct elements $a, a' \in A_{\zeta \leftarrow \bar{a}}$, where $\zeta \in \text{FTP}_\ell^\Sigma$. Taking any $b \in A$ notice that fluted $(\ell+1)$ -type ξ of $a'\bar{a}b$ lacks atoms $\mathbf{p}(x_{\ell+1} \cdots x_k)$, where $k < \ell+1$. Thus, ξ determines the satisfaction of quantifier-free $\mathcal{FLPC}^{\ell+1}$ -formulas by $b\bar{a}a'$ if and only if $a'\bar{a}b = b\bar{a}a'$. It is immediate that redefining the fluted $(\ell+1)$ -type of $a'\bar{a}b$ will not alter the satisfaction of quantifier-free $\mathcal{FLPC}^{\ell+1}$ -formulas by $c\bar{c}d \in A^{\ell+1} \setminus \{a'\bar{a}b\}$. Letting $\mathfrak{B} := \mathfrak{A}$ we set $\text{ftp}^\mathfrak{B}[a'\bar{a}b] := \text{ftp}^\mathfrak{A}[a\bar{a}b]$ for each $b \in A$ whilst leaving $\text{ftp}^\mathfrak{B}[c\bar{c}d] := \text{ftp}^\mathfrak{A}[c\bar{c}d]$ for $c\bar{c} \in A^\ell \setminus \{a'\bar{a}\}$ and $d \in A$. Notice that the assignment $\text{ftp}^\mathfrak{B}[a'\bar{a}b] := \text{ftp}^\mathfrak{A}[a\bar{a}b]$ is well-defined even if $b \in \{a, a'\}$ as neither $x_1 = x_{\ell+1}$ nor $x_1 \neq x_{\ell+1}$ are fluted atoms. To verify that $\mathfrak{B} \models \varphi$ we need only consider the ℓ -tuple $a'\bar{a}$. Picking any $r \in R$ suppose $\mathfrak{B} \models \alpha_r[a'\bar{a}]$. But then, $\mathfrak{A} \models \alpha_r[a\bar{a}]$ and thus $\mathfrak{A}, a\bar{a} \models \exists_{[n_r+p_r]} \gamma_r$. Since $\text{ftp}^\mathfrak{B}[a'\bar{a}b] = \text{ftp}^\mathfrak{A}[a\bar{a}b]$ for each $b \in A$ we must then have that $\mathfrak{B}, a'\bar{a} \models \exists_{[n+p]} \gamma_r$ as required. The reasoning is analogous for the requirements indexed by T .

Since fluted $(\ell+1)$ -types of $c\bar{c}d \in A^{\ell+1}$ with $c\bar{c} \neq a'\bar{a}$ are left unaltered, we may run the above in parallel for all $a' \in A_{\zeta \leftarrow \bar{a}} \setminus \{a\}$, thus obtaining a structure that models φ and in which ζ is \bar{a} -homogeneous. Using the same reasoning, we can repeat the procedure for each $\zeta \in \text{FTP}_\ell^\Sigma$ and $\bar{a} \in A^{\ell-1}$ thus obtaining:

Lemma 4.3.1. *Suppose φ is a satisfiable normal-form $\mathcal{FLPC}^{\ell+1}$ -sentence. Then, φ is satisfiable in a locally ℓ -homogeneous model.*

Using local ℓ -homogeneity coupled with variable reduction techniques prevalent in studies of fluted logics (see [Pratt-Hartmann et al., 2019]), we will establish a decidability result for the (finite) satisfiability problem of $\mathcal{FLPC}^{\ell+1}$. Fixing a normal-form $\mathcal{FLPC}^{\ell+1}$ -sentence φ we will compute a normal-form \mathcal{FLPC}^{ℓ} -sentence ψ defined as:

$$\psi_1 \wedge \cdots \wedge \psi_4 \tag{\psi}$$

that is satisfiable in structures holding just enough information to build locally ℓ -homogeneous models for φ . To aid intuition, we fix \mathfrak{A} to be any locally ℓ -homogeneous model of φ . We construct ψ whilst also expanding \mathfrak{A} into $\mathfrak{A}' \models \psi$. Note that the construction depends exclusively on the syntactic properties of φ .

First, set $\mathfrak{A}' := \mathfrak{A}$ and take $(\mathbf{q}_\zeta)_{\zeta \in \text{FTP}_\ell^\Sigma}$ to be a sequence of fresh $(\ell-1)$ -ary predicate symbols. In \mathfrak{A}' we decorate $(\ell-1)$ -tuples \bar{a} over A with \mathbf{q}_ζ just in case $A_{\zeta \leftarrow \bar{a}} \neq \emptyset$. That is to say, $\mathbf{q}_\zeta^{\mathfrak{A}'}$ remembers which $(\ell-1)$ -tuples can be extended (by appending an element to the left) to realise the fluted ℓ -type ζ . It is clear that \mathfrak{A}' models the following:

$$\bigwedge_{\zeta \in \text{FTP}_\ell^\Sigma} \forall^\ell (\zeta \rightarrow \mathbf{q}_\zeta). \tag{\psi_1}$$

Proceeding similarly, let $(\mathbf{s}_{\zeta, \xi})_{\zeta \in \text{FTP}_\ell^\Sigma, \xi \in \text{FTP}_{\ell+1}^\Sigma}$ be a sequence of new ℓ -ary predicates. Intuitively, we will have $\bar{a}b \in \mathbf{s}_{\zeta, \xi}^{\mathfrak{A}'}$ if in \mathfrak{A} the ℓ -tuple $\bar{a}b$ absorbs the fluted $(\ell+1)$ -type ξ emitted from $a\bar{a}$ for some $a \in A_{\zeta \leftarrow \bar{a}}$. Notice that, by local ℓ -homogeneity, if $\bar{a}b$ absorbs ξ from some $a\bar{a}$ with $a \in A_{\zeta \leftarrow \bar{a}}$, then it absorbs ξ from $a'\bar{a}$ for all $a' \in A_{\zeta \leftarrow \bar{a}}$. Thus, by our construction, $\mathbf{s}_{\zeta, \xi}$ is the unique predicate amongst $(\mathbf{s}_{\zeta, \xi'})_{\xi' \in \text{FTP}_{\ell+1}^\Sigma}$ satisfied by $\bar{a}b$ in \mathfrak{A}' . Clearly, \mathfrak{A}' models:

$$\begin{aligned} & \bigwedge_{\zeta \in \text{FTP}_\ell^\Sigma} \bigwedge_{\xi \in \text{FTP}_{\ell+1}^\Sigma} \forall^\ell \left(\mathbf{s}_{\zeta, \xi} \rightarrow \xi|_{[2, \ell+1]} \right) \wedge \\ & \bigwedge_{\zeta \in \text{FTP}_\ell^\Sigma} \forall^{\ell-1} \left(\mathbf{q}_\zeta \rightarrow \forall \left(\bigvee_{\xi \in \text{FTP}_{\ell+1}^\Sigma} \mathbf{s}_{\zeta, \xi} \wedge \bigwedge_{\substack{\xi \neq \xi' \\ \xi, \xi' \in \text{FTP}_{\ell+1}^\Sigma}} (\neg \mathbf{s}_{\zeta, \xi} \vee \neg \mathbf{s}_{\zeta, \xi'}) \right) \right). \end{aligned} \tag{\psi_2}$$

Again taking $\bar{a} \in A^{\ell-1}$ and any $\zeta \in \text{FTP}_\ell^\Sigma$ suppose $\zeta \models \alpha_r$ for some $r \in R$. In case $A_{\zeta \leftarrow \bar{a}}$ is non-empty (thus guaranteeing $\bar{a} \in \mathbf{q}_\zeta^{\mathfrak{A}'}$), we pick any $a \in A_{\zeta \leftarrow \bar{a}}$ and write $S = \{b \in A \mid \mathfrak{A}, a\bar{a}b \models \gamma_r\}$. Since ζ is \bar{a} -homogeneous in \mathfrak{A} , the exact element in $A_{\zeta \leftarrow \bar{a}}$ we pick has no effect on S . By our construction, S is then exactly the set of elements $b \in A$ such that $\mathfrak{A}', \bar{a}b \models \bigvee_{\xi \in \text{FTP}_{\ell+1}^\Sigma}^{\xi \models \gamma_r} \mathbf{s}_{\zeta, \xi}$. Since $|S| \in n_r^{+p_r}$ it is

then immediate that \mathfrak{A}' models the following:

$$\bigwedge_{r \in R} \bigwedge_{\zeta \in \text{FTP}_{\ell}^{\Sigma}}^{\zeta \models \alpha_r} \forall^{\ell-1} \left(\mathbf{q}_{\zeta} \rightarrow \exists_{[n_r^{+pr}]} \bigvee_{\xi \in \text{FTP}_{\ell+1}^{\Sigma}}^{\xi \models \gamma_r} \mathbf{s}_{\zeta, \xi} \right). \quad (\psi_3)$$

Similar observations follow whenever $\zeta \models \beta_t$ for some $t \in T$. This time, however, the cardinality of $S = \{b \in A \mid \mathfrak{A}, a\bar{a}b \models \delta_r\}$ must be outside the set n_r^{+pr} . Clearly, \mathfrak{A}' models:

$$\bigwedge_{t \in T} \bigwedge_{\zeta \in \text{FTP}_{\ell}^{\Sigma}}^{\zeta \models \beta_t} \forall^{\ell-1} \left(\mathbf{q}_{\zeta} \rightarrow \neg \exists_{[n_t^{+pt}]} \bigvee_{\xi \in \text{FTP}_{\ell+1}^{\Sigma}}^{\xi \models \delta_t} \mathbf{s}_{\zeta, \xi} \right). \quad (\psi_4)$$

Recalling that $\psi = \psi_1 \wedge \dots \wedge \psi_4$ we have shown the following:

Lemma 4.3.2. *Suppose $\mathfrak{A} \models \varphi$ is a locally ℓ -homogeneous model. Then, \mathfrak{A} can be extended to a model \mathfrak{A}' of ψ .*

We will now show the converse:

Lemma 4.3.3. *Suppose $\mathfrak{A}' \models \psi$. Then, we can construct a locally ℓ -homogeneous model \mathfrak{A}^+ of φ over the same domain.*

Proof. Supposing ψ is satisfiable we take any model \mathfrak{A}' . Now, let \mathfrak{A}^- be the model \mathfrak{A}' but with the predicates $(\mathbf{q}_{\zeta})_{\zeta \in \text{FTP}_{\ell}^{\Sigma}}$ and $(\mathbf{s}_{\zeta, \xi})_{\substack{\xi \in \text{FTP}_{\ell+1}^{\Sigma} \\ \zeta \in \text{FTP}_{\ell}^{\Sigma}}}$ removed from the signature. We proceed by expanding \mathfrak{A}^- into a locally ℓ -homogeneous model \mathfrak{A}^+ of the original sentence φ .

Fix $\bar{a} \in A^{\ell-1}$ and take some $a \in A$. Supposing that $\text{ftp}_{\ell}^{\mathfrak{A}^-}[a\bar{a}] = \zeta$, by ψ_1 we have that $\mathfrak{A}' \models \mathbf{q}_{\zeta}[a\bar{a}]$. Taking any $b \in A$ we observe that the conjuncts of ψ_2 enforce the following:

- if $\mathfrak{A}' \models \mathbf{s}_{\zeta, \xi}[a\bar{a}b]$ for some $\xi \in \text{FTP}_{\ell+1}^{\Sigma}$, then $\bar{a}b$ can absorb the fluted $(\ell+1)$ -type ξ ,
- $\bar{a}b$ satisfies at least one of the predicates $(\mathbf{s}_{\zeta, \xi})_{\xi \in \text{FTP}_{\ell+1}^{\Sigma}}$, and
- $\bar{a}b$ satisfies at most one of the predicates $(\mathbf{s}_{\zeta, \xi})_{\xi \in \text{FTP}_{\ell+1}^{\Sigma}}$.

We can then safely set $\text{ftp}_{\ell+1}^{\mathfrak{A}^+}[a\bar{a}b] := \xi$ for each $b \in A$, where ξ is taken from the subscript of the unique $\mathbf{s}_{\zeta, \xi} \in (\mathbf{s}_{\zeta, \xi})_{\xi \in \text{FTP}_{\ell+1}^{\Sigma}}$ that $\bar{a}b$ satisfies in \mathfrak{A}' . By repeating

the above procedure for all $a \in A$ and tuples $\bar{a} \in A^{\ell-1}$ we will obtain the desired structure \mathfrak{A}^+ .

To verify that \mathfrak{A}^+ is a model of φ we first claim that

$$\mathfrak{A}^+ \models \bigwedge_{r \in R} \forall^\ell (\alpha_r \rightarrow \exists_{[n_r^{+pr}]} \gamma_r).$$

For this purpose, fix some $r \in R$ and $a\bar{a} \in A^\ell$, and suppose $\zeta \models \alpha_r$, where $\text{ftp}_\ell^{\mathfrak{A}^+}[a\bar{a}] = \zeta$. Recall that $\bar{a} \in \mathfrak{q}_\zeta^{\mathfrak{A}'}$ by ψ_1 . Then, ψ_3 gives us $\mathfrak{A}', \bar{a} \models \exists_{[n_r^{+pr}]} \bigvee_{\xi \in \text{FTP}_{\ell+1}^{\Sigma}}^{\xi \models \gamma_r} \mathfrak{s}_{\zeta, \xi}$. Taking any $b \in A$ we have, by our construction, that $\mathfrak{A}^+, a\bar{a}b \models \xi$ if and only if $\mathfrak{A}', \bar{a}b \models \mathfrak{s}_{\zeta, \xi}$. Thus, $\mathfrak{A}^+, a\bar{a} \models \exists_{[n_r^{+pr}]} \bigvee_{\xi \in \text{FTP}_{\ell+1}^{\Sigma}}^{\xi \models \gamma_r} \xi$, which is equivalent to saying $\mathfrak{A}^+, a\bar{a} \models \exists_{[n_r^{+pr}]} \gamma_r$. Repeating the argument for each $r \in R$ and $a\bar{a} \in A^\ell$ will yield the required result. To show

$$\mathfrak{A}^+ \models \bigwedge_{t \in T} \forall^\ell (\beta_t \rightarrow \neg \exists_{[n_t^{+pt}]} \delta_t)$$

we proceed analogously with ψ_4 in place of ψ_3 . □

Let us take stock of the previous three lemmas. Take φ_ℓ to be an \mathcal{FLPC}^ℓ -sentence over some signature Σ_ℓ . We may, using Lemma 4.1.1, assume that it is in normal-form. By Lemma 4.3.1, if φ is satisfiable, then it is satisfiable in a locally $(\ell-1)$ -homogeneous model. By computing the formula $\varphi_{\ell-1} := \psi$ over $\Sigma_{\ell-1}$ as described above, we have, by Lemmas 4.3.2 and 4.3.3, that $\varphi_{\ell-1}$ is (finitely) satisfiable if and only if φ is. We claim that $\|\varphi_{\ell-1}\|$ is exponentially larger than $\|\varphi_\ell\|$. To see this we need to count the number of members in $\text{FTP}_\ell^{\Sigma_\ell}$ and $\text{FTP}_{\ell-1}^{\Sigma_\ell}$. Notice that a single fluted ℓ - and $(\ell-1)$ -type is of size at most $|\Sigma_\ell|$. Thus, there are at most $2^{|\Sigma_\ell|}$ different fluted ℓ - and $(\ell-1)$ -types. Since $|\Sigma_\ell| \leq \|\varphi_\ell\|$ it is easy to verify that $\|\varphi_{\ell-1}\|$ is bounded by $2^{\|\varphi_\ell\|^{O(1)}}$.

We may repeat the variable reduction procedure on $\varphi_{\ell-1}$ and subsequent formulas thus obtaining a sequence $\varphi_\ell, \dots, \varphi_2$ of equisatisfiable sentences each being in one fewer variable than the last, but exponentially larger. By Theorem 4.2.4, we can determine the (finite) satisfiability status of φ_2 (and thus of φ_ℓ) in non-deterministic time as an exponential function on $\|\varphi_2\|$. Thus, to determine the computational resource needed for satisfiability checking, we need to calculate

$\|\varphi_2\|$ relative to $\|\varphi_\ell\|$. Recall the tower of exponentials function $\mathfrak{t} : \mathbb{N} \times \mathbb{N} \rightarrow \mathbb{N}$:

$$\begin{aligned}\mathfrak{t}(0, n) &:= n, \\ \mathfrak{t}(m+1, n) &:= 2^{\mathfrak{t}(m, n)}.\end{aligned}$$

An easy induction reveals that $\|\varphi_i\|$ is bounded by $\mathfrak{t}(\ell-i, \|\varphi_\ell\|^{O(1)})$ for all $i \in [2, \ell]$. Thus, $\|\varphi_2\|$ is at most $\mathfrak{t}(\ell-2, \|\varphi_\ell\|^{O(1)})$ revealing that it takes nondeterministic time bounded by $\mathfrak{t}(\ell-1, \|\varphi_\ell\|^{O(1)})$ to determine the (finite) satisfiability status of φ_ℓ . We are in a position to conclude the following:

Theorem 4.3.4. *The finite and general satisfiability problems for \mathcal{FLPC}^ℓ are in $(\ell-1)$ -NEXPTIME.*

Recalling that the satisfiability problems for \mathcal{FL}^ℓ are $\lfloor \ell/2 \rfloor$ -NEXPTIME-hard [Pratt-Hartmann et al., 2019], we see that no elementary function can encapsulate the computational resources needed for deciding (finite) satisfiability for \mathcal{FLPC} .

Theorem 4.3.5. *The (finite) satisfiability problem for \mathcal{FLPC} is TOWER-complete.*

4.4 Homogeneity in the Study of Flutedness

In this chapter we utilised the homogeneity property of satisfiable \mathcal{FLPC} sentences to establish a decision procedure for (finite) satisfiability of the new language. With this methodology we not only gained a better understanding of models of fluted formulas, but also managed to establish decidability of (finite) satisfiability using simpler methods when compared to Presburger quantifiers discussed in previous literature [Pratt-Hartmann, 2021].

Recent years have seen fluted languages in consideration with dedicated relations which are otherwise undefinable in \mathcal{FL} (or even \mathcal{FLPC}). The transitivity relation, studied in [Pratt-Hartmann and Tendera, 2019] without equality, [Pratt-Hartmann and Tendera, 2022] with equality, and [Pratt-Hartmann and Tendera, 2023] without equality but with counting quantifiers, has seen particular interest. A striking consequence of our work is that even semantically enhanced fragments (such as the ones with transitive relations) adhere to *some* homogeneity conditions. To see this recall that global homogeneity established in Lemma 4.2.1 is dependent on the semantics of the equality predicate. On the other hand, local

ℓ -homogeneity (for $\ell \geq 2$) is not. We are thus invited to make the following generalisation. Consider a signature that is split into symbols Σ^* with a constrained interpretation (e.g. transitive relation, reversed relation, etc.) and standard predicate symbols Σ with no predetermined meaning. Now, let k be the maximum arity of any symbol in Σ^* . Then, Lemma 4.3.1 implies:

Corollary 4.4.1. *Suppose φ is a fluted, normal-form, $(\ell+1)$ -variable sentence (possibly with periodic counting) over the signature $\Sigma \cup \Sigma^*$, and where $\ell \geq k$. Then, if φ is satisfiable, it is satisfiable in a locally ℓ -homogenous model.*

The same, however cannot be said about $(\ell+1)$ -variable sentences when $\ell < k$, and thus establishing an analogue to global homogeneity (as was done for \mathcal{FLPC}^2 with equality in Lemma 4.2.1) requires case-by-case consideration.

Nonetheless, we believe our approach could not only be used to simplify existing decidability procedures for satisfiability (e.g. for \mathcal{FL} with a transitive relation and counting [Pratt-Hartmann and Tendera, 2023]) but to also expand on expressiveness of fluted languages.

Chapter 5

The Adjacent Fragment of First-Order Logic

The following chapter reproduces a technical report [Bednarczyk et al., 2024] submitted to Journal of Logic and Computation. The technical report itself is an extension of the conference paper [Bednarczyk et al., 2023]. Both the technical report and conference paper is a result of equal collaboration between Bartosz Bednarczyk, Ian Pratt-Hartmann, and the Ph. d. candidate.

In the spirit of W. V. Quine [Quine, 1969] we push the boundary of decidability for logics obtained by restricting allowable variable sequences to its extremities. We begin by identifying a new language which we will call the *adjacent fragment* and denote it by \mathcal{AF} . This formalism not only subsumes the fluted [Purdy, 1996], ordered [Herzig, 1990] and forward [Bednarczyk, 2021] fragments, but also two-variable logic [Scott, 1962].

The *adjacency* constraint is perhaps best explained informally, from the comfort of the fluted fragment. Suppose that $\alpha := p(x_k \cdots x_\ell)$ is an atom appearing in some fluted formula ψ . Then, the formula obtained by replacing the arguments $x_k \cdots x_\ell$ of α in ψ by some word \bar{x} over $x_1 \cdots x_\ell$ is adjacent so long as neighbouring elements of \bar{x} do not differ in their indices by more than 1. Thus, a sentence such as “Some opera soloists would like to see Wozzeck played by themselves at every concert” can be translated to the adjacent setting as follows:

$$\exists x_1 \exists x_2 \forall x_3 \left(\text{wozzeck}(x_1) \wedge \text{soloist}(x_2) \wedge (\text{concert}(x_3) \rightarrow \text{person_wouldLikeCharacter_bePlayedBy_at}(x_2 x_1 x_2 x_3)) \right).$$

Additionally, as opposed to fluted, ordered and forward languages, the standard axioms of reflexivity and symmetry are in the adjacent fragment:

$$\forall x_1 \mathbf{r}(x_1x_1), \quad \forall x_1x_2 (\mathbf{r}(x_1x_2) \rightarrow \mathbf{r}(x_2x_1)).$$

It is worth noting that the standard axiom of transitivity remains outside our new paradigm as the conclusion $\mathbf{r}(x_1x_3)$ involves a skip in variable indices:

$$\forall x_1x_2x_3 \left((\mathbf{r}(x_1x_2) \wedge \mathbf{r}(x_2x_3)) \rightarrow \mathbf{r}(x_1x_3) \right).$$

We define the fragment formally in Section 5.1.

Our main result concerns the finite and general satisfiability problems of \mathcal{AF} . We will show that the fragment in question enjoys the finite model property, allowing us to conclude that $\text{FinSat}(\mathcal{AF}) = \text{Sat}(\mathcal{AF})$ is decidable. In particular, Section 5.4 will show that equality-free sentences of the ℓ -variable subfragment of \mathcal{AF} can be checked for satisfiability in $(\ell-2)$ -NEXPTIME for $\ell \geq 3$. In Section 5.5 we will show that, if equality is present, one can also check for satisfiability albeit in $(\ell-1)$ -NEXPTIME. This, combined with the fact that the satisfiability problem for the ℓ -variable fluted fragment is $\lfloor \ell/2 \rfloor$ -NEXPTIME-hard [Pratt-Hartmann et al., 2019] will allow us to conclude that $\text{Sat}(\mathcal{AF})$ is TOWER-complete.

We discuss the expressive power of \mathcal{AF} in Section 5.3. We note here, however, that the adjacent fragment is incomparable to the guarded fragment. In fact, $\text{Sat}(\mathcal{AF} \cup \mathcal{GF})$ is immediately undecidable as in \mathcal{AF} one can axiomatise universal relations $\mathbf{r}_1, \dots, \mathbf{r}_n$ of arities $1, \dots, n$, and thus unlock the full expressive power of \mathcal{FO}^n (over structures where \mathbf{r}_i is satisfied by every i -tuple) by writing \mathcal{GF} -sentences guarded by $\mathbf{r}_1, \dots, \mathbf{r}_n$. We will thus shift our attention to the *guarded adjacent fragment* $\mathcal{GA} := \mathcal{AF} \cap \mathcal{GF}$. The language \mathcal{GA} is of particular interest as it is a multi-variable extension of the family of description logics \mathcal{ALC} ; in particular, \mathcal{ALCHI} ; or: \mathcal{ALC} augmented with role hierarchies and role inverses. (See [Hustadt et al., 2004] for translations). In Section 5.5, we show that $\text{Sat}(\mathcal{GA})$ is no easier than $\text{Sat}(\mathcal{GF})$; i.e. 2-EXPTIME-complete. This contrasts with previously known EXPTIME-completeness results for the satisfiability problem of the guarded fluted fragment [Bednarczyk, 2021]. It is unclear if \mathcal{GA} can be extended further to generalise more expressive variants of \mathcal{ALCHI} such as \mathcal{ALCHIQ} (that is \mathcal{ALCHI} with cardinality restrictions).

Lastly, Section 5.7 will culminate in a result stating that the adjacency constraint cannot be relaxed without the loss of decidability for satisfiability. This renders the adjacent fragment maximal amongst variable-sequence fragments.

5.1 Preliminaries

We begin by formally introducing the adjacency constraint. Pick any $m, k \in \mathbb{N}$ and let $f : [1, m] \rightarrow [1, k]$ be a function. We say that f is a *walk* just in case $|f(i+1) - f(i)| \leq 1$ for each $i \in [1, m]$. Given a k -tuple $\bar{a} := a_1 \cdots a_k$ we define \bar{a}^f to be the m -tuple $a_{f(1)}a_{f(2)} \cdots a_{f(m)}$. We picture the operation \cdot^f on \bar{a} as *walking on \bar{a}* by starting at position $f(1)$ and, at step $i \in [1, m-1]$, taking a step left, right or staying put as dictated by $f(i+1) - f(i)$. In the sequel we will call tuples $\bar{b} := \bar{a}^f$ *walks on \bar{a}* , sometimes leaving the walk f implicit.

Now, let us write \mathbf{x}_ℓ for the word $x_1x_2 \cdots x_\ell$. Supposing \mathbf{p} is an m -ary predicate symbol, we define an *adjacent ℓ -atom* to be a formula of the form $\mathbf{p}(\mathbf{x}_\ell^f)$, where $f : [1, m] \rightarrow [1, \ell]$ is a walk. Alternatively, an adjacent ℓ -atom is an atomic formula over variables $x_1x_2 \cdots x_\ell$ in which adjacent arguments differ by no more than one in terms of their indices. We shall count proposition letters as adjacent 0-atoms. We are now in a position to define the adjacent fragment formally. First, let us inductively define the set of formulas $\mathcal{AF}^{[\ell]}$ for all $\ell \in \mathbb{N}$ as follows:

1. every adjacent ℓ -atom is in $\mathcal{AF}^{[\ell]}$;
2. $\mathcal{AF}^{[\ell]}$ is closed under Boolean combinations;
3. if φ is in $\mathcal{AF}^{[\ell+1]}$, then $\exists x_{\ell+1} \varphi$ in $\mathcal{AF}^{[k]}$ for all $k \geq \ell$.

Then, $\mathcal{AF} := \bigcup_{\ell \in \mathbb{N}} \mathcal{AF}^{[\ell]}$, whilst the ℓ -variable subfragment is defined as $\mathcal{AF}^\ell := \mathcal{AF} \cap \mathcal{FO}^\ell$. For convenience, we will continue to write $\forall x \psi$ in place of $\neg \exists x \neg \psi$. When given any first-order formula $\varphi(x_1 \cdots x_n)$ it will be convenient to write φ^{-1} for $\varphi(x_n \cdots x_1)$ and $\widehat{\varphi}$ for $\varphi \wedge \varphi^{-1}$.

An \mathcal{AF} -sentence is in *normal-form* if it takes the following shape:

$$\forall \mathbf{x}_{\ell+1} \beta \wedge \bigwedge_{t \in T} \forall \mathbf{x}_\ell \exists x_{\ell+1} \gamma_t, \quad (\mathcal{AF}\text{-nmf})$$

where β and γ_t are quantifier-free $\mathcal{AF}^{\ell+1}$ -formulas indexed by a finite set T . Each \mathcal{AF} -formula can be put into normal-form by introducing new predicate symbols as evident by the following:

Lemma 5.1.1. *Suppose φ is an $\mathcal{AF}^{\ell+1}$ -sentence. Then, we may compute, in polynomial time, a normal-form $\mathcal{AF}^{\ell+1}$ -sentence ψ that is satisfiable over the same domains as φ . Moreover, ψ is equality-free if and only if φ is.*

Proof. If the sentence φ is quantifier-free, then it is a formula of the propositional calculus, and the result is easily obtained by adding vacuous quantification. Otherwise, write $\varphi_0 := \varphi$, and let $\theta := Qx_{k+1} \chi$ be a subformula of φ , where $Q \in \{\forall, \exists\}$ and χ is quantifier-free. Writing $\widehat{\exists} := \forall$ and $\widehat{\forall} := \exists$, let \mathbf{p} be a new predicate of arity k . Define φ_1 to be the result of replacing θ in φ_0 by the atom $\mathbf{p}(\mathbf{x}_k)$, and let ψ_1 be the formula

$$\forall \mathbf{x}_k Qx_{k+1} (\mathbf{p}(\mathbf{x}_k) \rightarrow \chi) \wedge \forall \mathbf{x}_k \widehat{Q}x_{k+1} (\chi \rightarrow \mathbf{p}(\mathbf{x}_k)).$$

It is immediate that $\varphi_1 \wedge \psi_1 \models \varphi_0$. Conversely, if $\mathfrak{A} \models \varphi_0$, then we may expand \mathfrak{A} to a model \mathfrak{A}' of $\varphi_1 \wedge \psi_1$ by taking $\mathbf{p}^{\mathfrak{A}'}$ to be the set of k -tuples \bar{a} such that $\mathfrak{A} \models \theta[\bar{a}]$. Evidently, φ_1 is a sentence of $\mathcal{AF}^{\ell+1}$. Processing φ_1 in the same way, and proceeding similarly, we obtain a set of formulas $\varphi_2, \dots, \varphi_m$ and ψ_2, \dots, ψ_m , with φ_m quantifier-free and φ_0 satisfiable over the same domains as $\psi_1 \wedge \dots \wedge \psi_m \wedge \varphi_m$. Since φ_m is a sentence, it is a formula of the propositional calculus. By moving φ_m inside one of the quantified formulas, re-indexing variables and re-ordering conjuncts, we obtain a formula ψ of the form (\mathcal{AF} -nmf). \square

We generalise ideas from the fluted fragment to fit the adjacent setting. Take Σ to be some finite, constant- and function-free signature. Recall that an atom is an adjacent $(\ell+1)$ -atom if its argument sequence is a walk on $\mathbf{x}_{\ell+1}$. An *adjacent $(\ell+1)$ -literal* is an expression of the form $\pm \mathbf{p}(\bar{x})$, where $\mathbf{p}(\bar{x})$ is an adjacent $(\ell+1)$ -atom and \pm is either the negation symbol or the empty word. An *adjacent $(\ell+1)$ -type* is then a maximal consistent set of adjacent $(\ell+1)$ -literals over $\Sigma \cup \{=\}$. Given a Σ -structure \mathfrak{A} , each $(\ell+1)$ -tuple $a\bar{a}b$ realises a unique adjacent $(\ell+1)$ -type denoted by $\mathbf{atp}^{\mathfrak{A}}[a\bar{a}b]$. An adjacent $(\ell+1)$ -literal α is in $\mathbf{atp}^{\mathfrak{A}}[a\bar{a}b]$ if and only if $\mathfrak{A} \models \alpha[a\bar{a}b]$. Taking an adjacent ℓ -type η , we will write η^+ for the formula obtained by incrementing all variable indices in η by 1. Lastly, we define $\mathbf{ATP}_{\ell+1}^{\Sigma}$ as the set of all adjacent $(\ell+1)$ -types over $\Sigma \cup \{=\}$.

Let us keep the signature Σ as before. We define an *incremental $(\ell+1)$ -type* ι to be a maximal consistent of adjacent $(\ell+1)$ -literals of the form $\pm \mathbf{p}(\mathbf{x}_{\ell+1}^f)$ with $f : [1, m] \rightarrow [1, \ell+1]$ being a *surjective* walk. Alternatively, we think of an incremental $(\ell+1)$ -type as a collection adjacent $(\ell+1)$ -literals that feature *all*

variables from $\mathbf{x}_{\ell+1}$ as arguments. Given a Σ -structure \mathfrak{A} and $\bar{a} \in A^{\ell+1}$ we denote the unique incremental $(\ell+1)$ -type realised by \bar{a} as $\text{itp}^{\mathfrak{A}}[\bar{a}]$. Every adjacent $(\ell+1)$ -type ξ can be thought of as the union of ζ, η^+ and ι , for some unique $\zeta, \eta \in \text{ATP}_{\ell}^{\Sigma}$ and a unique incremental $(\ell+1)$ -type ι . We write $\partial\xi$ for the unique incremental $(\ell+1)$ -type included in ξ .

5.2 Primitive generators of words

Our results on decidability of satisfiability for the adjacent fragment depend on the following observations on the combinatorics of words, presented in [Pratt-Hartmann, 2024]. Say, that a k -tuple \bar{c} generates an m -tuple \bar{a} if there is a surjective walk $f : [1, m] \rightarrow [1, k]$ such that $\bar{c}^f = \bar{a}$. We again invite the reader to regard \cdot^f applied to \bar{c} as f walking on \bar{c} by first visiting position $f(1)$ of \bar{c} , then, at step $i \in [1, m-1]$, taking a step left, right or staying put as dictated by $f(i+1) - f(i)$. Notice that the surjectiveness criterion requires f to visit every position of \bar{c} . Clearly, if \bar{c} generates \bar{a} , then \bar{c} does as well. Additionally, generation is transitive; that is, if \bar{c} generates \bar{a} and \bar{a} generates \bar{b} , then \bar{c} generates \bar{b} . We say that a word \bar{c} is *primitive* if \bar{c} and \bar{c} are the only words generating it. It is obvious that every word has a primitive generator; what might come as some surprise is that it is unique up to reversal:

Lemma 5.2.1. [Pratt-Hartmann, 2024, Theorem 1] *The primitive generator of any word is unique up to reversal.*

Whilst primitive generators are unique, modes of generation are not. As an example, take $\bar{c} := \text{cabad}$. Clearly, \bar{c} generates $\bar{a} := \text{cababad}$, however there are two surjective walks that do so. We provide them as the course of values $(f(1) \cdots f(7))$ and $(g(1) \cdots g(7))$ below:

$$\begin{aligned} f &:= (1 \ 2 \ 3 \ 2 \ 3 \ 4 \ 5), \\ g &:= (1 \ 2 \ 3 \ 4 \ 3 \ 4 \ 5). \end{aligned}$$

In short, f begins generation at position 1, takes two steps right until position 3 and then takes a single step left before continuing rightwards without any further hesitations. The walk g is almost identical; the only difference is that the ‘‘hesitation’’ present in f is different for a single step. Nonetheless, $\bar{c}^f = \bar{a} =$

\bar{c}^g . The ambiguity in generation arises from the presence of palindromes in the primitive generator. In the case of \bar{c} , this is $c_2c_3c_4 = aba$. Given any primitive word \bar{d} , we say that $\langle i, j \rangle$ is a *defect* (or: *defective*) just in case $d_i \cdots d_j$ is a palindrome. Each primitive string \bar{d} gives rise to a relation:

$$D := \{\langle i, j \rangle \mid 1 \leq i \leq j \leq |\bar{d}|, \langle i, j \rangle \text{ is defective in } \bar{d}\}$$

which we will call *the defect set of \bar{d}* . Note that, for $i > 1$ and $j < |\bar{d}|$, neither $\langle 1, i \rangle$, nor $\langle j, |\bar{d}| \rangle$ are defects of primitive words \bar{d} . Given a set of defects D and two walks f, g over the same domain and co-domain, we write $f \stackrel{D}{=} g$ just in case the pair $\langle f(i), g(i) \rangle$ is in the equivalence closure of D for each $i \in \text{dom}(f)$. The following equivalence will help us deal with ambiguous generation in the sequel:

Lemma 5.2.2. [Pratt-Hartmann, 2024, Theorem 4]¹ *Let \bar{c} be a primitive word of length k with the defect set D , and let $f, g : [1, m] \rightarrow [1, k]$ be surjective walks. Then $\bar{c}^f = \bar{c}^g$ if and only if $f \stackrel{D}{=} g$.*

Returning to the adjacent fragment, take any formula adjacent formula ψ in the free variables \mathbf{x}_ℓ . If $f : [1, \ell] \rightarrow [1, k]$ is some walk, it will be convenient to write ψ^f for the formula $\psi(\mathbf{x}_k^f)$. Under the supposition that f is surjective, the formula ψ^f is in free variables \mathbf{x}_k . The following lemma ties the syntactic notion of walking on variables with the semantics of a formula being satisfied by a tuple.

Lemma 5.2.3. *Take a quantifier-free \mathcal{AF}^ℓ -formula ψ , and let $f : [1, \ell] \rightarrow [1, k]$ be a walk. For any structure \mathfrak{A} and any $\bar{c} \in A^k$, we have $\mathfrak{A} \models \psi^f[\bar{c}]$ iff $\mathfrak{A} \models \psi[\bar{c}^f]$.*

Proof. Suppose that ψ is atomic. Then $\psi = \mathbf{p}(\mathbf{x}_\ell^g)$, where $g : [1, m] \rightarrow [1, \ell]$ is a walk and \mathbf{p} is m -ary. By definition, $\psi^f = \psi(x_{f(1)} \cdots x_{f(\ell)}) = \mathbf{p}(x_{f(g(1))} \cdots x_{f(g(m))})$. Thus, $\mathfrak{A} \models \psi^f[\bar{c}]$ is equivalent to $\mathfrak{A} \models \mathbf{p}[c_{f(g(1))} \cdots c_{f(g(m))}]$. On the other hand, suppose $\mathfrak{A} \models \psi[\bar{c}^f]$. We have $\mathfrak{A} \models \psi[c_{f(1)} \cdots c_{f(\ell)}]$, thus $\mathfrak{A} \models \mathbf{p}[c_{f(g(1))} \cdots c_{f(g(m))}]$. Since the two final expressions coincide, we will have the required result by routine structural induction on ψ . \square

An immediate consequence of Lemma 5.2.3 is that, in any structure \mathfrak{A} , the adjacent type of any tuple \bar{a} is determined by the adjacent type of its primitive generator \bar{c} . More precisely, if f is a surjective walk such that $\bar{c}^f = \bar{a}$, then $\alpha \in \text{atp}^\mathfrak{A}[\bar{c}^f]$ if and only if $\alpha^f \in \text{atp}^\mathfrak{A}[\bar{c}]$.

¹The Ph. d. candidate contributed essentially to the formulation of [Pratt-Hartmann, 2024, Theorem 4].

Say that a surjective walk $f : [1, m] \rightarrow [1, \ell]$ is *terminal* if $f(m) = \ell$. Still keeping f as before, denote by f^+ the function $f \cup \{m+1 \mapsto \ell+1\}$. Clearly, f^+ is a walk as long as f is a terminal walk. The following is an easy observation about non-primitive strings that will be useful shortly.

Lemma 5.2.4. *Take \bar{a} to be an m -tuple and b to be an element that does not appear in \bar{a} . If $\bar{a}b$ is not primitive, then neither is \bar{a} . In fact, \bar{a} is generated by a word \bar{c} of length less than m and a surjective terminal walk f .*

Proof. Let us take \bar{a} and b as described above. Since $\bar{a}b$ is non-primitive, it is generated by a primitive $(k+1)$ -tuple \bar{d} as a surjective walk $f : [1, m+1] \rightarrow [1, k+1]$. Clearly, $k < m$. Since b does not appear in \bar{a} we have that either $d_1 = b$ or $d_{k+1} = b$. Reversing \bar{d} and f if needed, suppose it is the latter. Then, f is terminal. Defining $g := f \setminus \{k+1 \mapsto m+1\}$, we have that g is a surjective terminal walk that generates \bar{a} from \bar{d} . Since, b does not appear in \bar{a} we, in fact, have that the co-domain of g is $[1, k]$ and thus $\bar{c}^g = \bar{a}$, where $\bar{c} := d_1 \cdots d_k$. \square

We now identify some logical consequences of normal-form $\mathcal{AF}^{\ell+1}$ -sentences such as φ . Let us write $\mathbf{A}_k^{\ell+1}$ for the set of walks $f : [1, \ell+1] \rightarrow [1, k]$, and $\vec{\mathbf{A}}_k^\ell$ for the set of all terminal walks $f : [1, \ell] \rightarrow [1, k]$. We define the *adjacent closure* of φ , denoted $\text{acl}(\varphi)$, to be the following \mathcal{AF}^ℓ -sentence:

$$\bigwedge_{k=1}^{k \leq \ell} \bigwedge_{g \in \mathbf{A}_k^{\ell+1}} \forall \mathbf{x}_k \beta^g \wedge \bigwedge_{t \in T} \bigwedge_{k=1}^{k \leq \ell-1} \bigwedge_{f \in \vec{\mathbf{A}}_k^\ell} \forall \mathbf{x}_k \exists x_{k+1} \gamma_t^{f^+}. \quad (\text{acl}(\varphi))$$

Lemma 5.2.5. *Let $\varphi \in \mathcal{AF}^{\ell+1}$ be in normal-form. Then $\varphi \models \text{acl}(\varphi)$.*

Proof. Taking any $\mathfrak{A} \models \varphi$ we show that $\mathfrak{A} \models \text{acl}(\varphi)$. To this end fix some $t \in T$, $k \in [1, \ell-1]$ and a terminal walk $f : [1, \ell] \rightarrow [1, k]$. Taking any $\bar{a} \in A^k$ we claim that $\mathfrak{A}, \bar{a} \models \exists x_{k+1} \gamma_t^{f^+}$. To see this, let $\bar{c} := \bar{a}^f$. Since \bar{c} is an ℓ -tuple we have, by φ , that $\mathfrak{A}, \bar{c} \models \exists x_{\ell+1} \gamma_t$. Let $d \in A$ be the witness to this existential requirement in regard to \bar{c} . Since $\bar{c} = \bar{a}^f$ and f is terminal, we have that $\bar{c}d = (\bar{a}d)^{f^+}$. Thus, d is also a witness to $\exists x_{k+1} \gamma_t^{f^+}$ in regard to \bar{a} . Showing that \mathfrak{A} models the universal requirements of $\text{acl}(\varphi)$ is similar. \square

The formula $\text{acl}(\varphi)$ will play an important role in variable reductions to come. The task of $\text{acl}(\varphi)$ is to ensure that non-primitive $(\ell+1)$ -tuples follow the universal requirements of φ . Similarly, $\text{acl}(\varphi)$ provides witnesses for the existential

requirements of φ to non-primitive ℓ -tuples that can be generated by a terminal walk on a strictly shorter tuple. It will become apparent that (for the most part) only primitive $(\ell+1)$ -tuples require additional attention.

5.3 Expressive Power

We study the expressive power of languages by introducing a notion of *bisimulations* suitable for the adjacent fragment. Let \mathfrak{A} be any structure and take some $\bar{a} \in A^k$. We call the pair (\mathfrak{A}, \bar{a}) a *pointed structure*. For convenience let us fix two such structure (\mathfrak{A}, \bar{a}) and (\mathfrak{B}, \bar{b}) for the rest of the section. We will write $(\mathfrak{A}, \bar{a}) \simeq^{\mathcal{AF}} (\mathfrak{B}, \bar{b})$ just in case the following holds:

1. **Harmony:** $\text{atp}^{\mathfrak{A}}[\bar{a}] = \text{atp}^{\mathfrak{B}}[\bar{b}]$,
2. **Forth:** for each² $0 \leq i \leq j \leq |\bar{a}|$ and $c \in A$ there is some $d \in B$ such that $(\mathfrak{A}, a_i \cdots a_j c) \simeq^{\mathcal{AF}} (\mathfrak{B}, b_i \cdots b_j d)$,
3. **Back:** for each² $0 \leq i \leq j \leq |\bar{b}|$ and $d \in B$ there is some $c \in A$ such that $(\mathfrak{A}, a_i \cdots a_j c) \simeq^{\mathcal{AF}} (\mathfrak{B}, b_i \cdots b_j d)$.

We say that (\mathfrak{A}, \bar{a}) and (\mathfrak{B}, \bar{b}) are \mathcal{AF} -*bisimilar* just in case $(\mathfrak{A}, \bar{a}) \simeq^{\mathcal{AF}} (\mathfrak{B}, \bar{b})$.

Alternatively, the conditions above can be formulated as in terms of Ehrenfeucht–Fraïssé-style games, where, broadly speaking, the first player, nicknamed the *spoiler*, tries to show that the two pointed structures are different, whilst the second player, the *duplicator*, matches the spoiler’s moves in an attempt to show that the two pointed structures are similar. More formally, the games are played as follows. The two players are handed two pointed structures (\mathfrak{A}, \bar{a}) and (\mathfrak{B}, \bar{b}) . The spoiler picks some structure, say \mathfrak{A} , indices $0 \leq i \leq j \leq |\bar{a}|$, and an element $c \in A$. The duplicator then chooses some element $d \in B$. We say that the spoiler wins if by repeating the process above for a finite number of steps on $(\mathfrak{A}, a_i \cdots a_j c)$ and $(\mathfrak{B}, b_i \cdots b_j d)$ (and subsequent structures thereafter), two pointed structures (\mathfrak{A}, \bar{c}) and (\mathfrak{B}, \bar{d}) having the property $\text{atp}^{\mathfrak{A}}[\bar{c}] \neq \text{atp}^{\mathfrak{B}}[\bar{d}]$ are obtained. The duplicator wins if such pointed structures are not reached.

Now, let us write $(\mathfrak{A}, \bar{a}) \equiv^{\mathcal{AF}} (\mathfrak{B}, \bar{b})$ just in case $\mathfrak{A} \models \psi[\bar{a}] \Leftrightarrow \mathfrak{B} \models \psi[\bar{b}]$ for every \mathcal{AF} -formula ψ . That is to say, $(\mathfrak{A}, \bar{a}) \equiv^{\mathcal{AF}} (\mathfrak{B}, \bar{b})$ just in case the two pointed structures satisfy the same adjacent formulas. The following links \mathcal{AF} -bisimulations to the notion of structures satisfying the same \mathcal{AF} -formulas.

²In the context of \bar{a} and \bar{b} we will assume that a_0 and b_0 are empty words.

Lemma 5.3.1. $(\mathfrak{A}, \bar{a}) \simeq^{\mathcal{AF}} (\mathfrak{B}, \bar{b})$ implies $(\mathfrak{A}, \bar{a}) \equiv^{\mathcal{AF}} (\mathfrak{B}, \bar{b})$. The converse is true over ω -saturated structures.

Proof. We show the first statement by structural induction on formulas $\psi \in \mathcal{AF}$. Firstly, if ψ is atomic, then $\mathfrak{A} \models \psi[\bar{a}] \Leftrightarrow \text{atp}^{\mathfrak{A}}[\bar{a}] \models \psi$ and $\mathfrak{B} \models \psi[\bar{b}] \Leftrightarrow \text{atp}^{\mathfrak{B}}[\bar{b}] \models \psi$. By harmony, $\text{atp}^{\mathfrak{A}}[\bar{a}] = \text{atp}^{\mathfrak{B}}[\bar{b}]$ thus securing the required property. The cases where $\psi := \neg\theta$ and $\psi := \chi \wedge \theta$ are trivial. Thus, we are only left to deal with $\psi := \exists x_{j+1} \theta$. Keeping $(\mathfrak{A}, \bar{a}) \simeq^{\mathcal{AF}} (\mathfrak{B}, \bar{b})$ suppose that $\mathfrak{A} \models \psi[\bar{a}]$. Thus, there is some $c \in A$ such that $\mathfrak{A} \models \theta[a_1 \cdots a_j c]$. By forth we may find $d \in B$ such that $(\mathfrak{A}, a_1 \cdots a_j c) \simeq^{\mathcal{AF}} (\mathfrak{B}, b_1 \cdots b_j d)$. By our inductive hypothesis, $\mathfrak{B} \models \theta[b_1 \cdots b_j d]$ as required. The case where $\mathfrak{B} \models \psi[\bar{b}]$ is handled similarly.

For the second statement assume that \mathfrak{A} and \mathfrak{B} are ω -saturated and that $(\mathfrak{A}, \bar{a}) \equiv^{\mathcal{AF}} (\mathfrak{B}, \bar{b})$. We claim that the set

$$\mathbf{S} := \bigcup_{n < \omega} \{(\bar{c}, \bar{d}) \in A^n \times B^n \mid (\mathfrak{A}, \bar{c}) \equiv^{\mathcal{AF}} (\mathfrak{B}, \bar{d})\}$$

is an \mathcal{AF} -bisimulation between the two pointed structures. Clearly, the harmony condition is satisfied by pairs in \mathbf{S} . To show the back and forth properties take some $(\bar{c}, \bar{d}) \in \mathbf{S}$. Let first us focus on the structure (\mathfrak{A}, \bar{c}) . Picking positions $0 \leq i \leq j \leq |\bar{c}|$ to keep and an element $c' \in A$ let us form the pointed structure $(\mathfrak{A}, c_i \cdots c_j c')$. We show that the set of formulas

$$\Gamma := \{\psi \in \mathcal{AF} \mid \mathfrak{A} \models \psi[c_i \cdots c_j c']\}$$

is realised by $d_i \cdots d_j d'$ in \mathfrak{B} for some $d' \in B$. We first show that the claim holds for every finite subset $\gamma \subseteq \Gamma$. Let us identify γ as the conjunction of the formulas it contains. Clearly, $\mathfrak{A}, c_i \cdots c_j \models \exists x_{j+1} \gamma$. Since $(\mathfrak{A}, \bar{c}) \equiv^{\mathcal{AF}} (\mathfrak{B}, \bar{d})$ we have that $\mathfrak{B}, d_i \cdots d_j \models \exists x_{j+1} \gamma$. Writing

$$\Gamma' := \{\psi \in \mathcal{AF} \mid \mathfrak{B} \models \psi[d_i \cdots d_j]\}$$

we have that $\Gamma' \cup \gamma$ is consistent. By compactness, the same is true for $\Gamma' \cup \Gamma$. Since \mathfrak{B} is ω -saturated, we may find an element $d' \in B$ such that $\mathfrak{B} \models \Gamma[d_i \cdots d_j d']$. Clearly, $(\mathfrak{A}, c_i \cdots c_j c') \equiv^{\mathcal{AF}} (\mathfrak{B}, d_i \cdots d_j d')$ thus securing $(c_i \cdots c_j c', d_i \cdots d_j d') \in \mathbf{S}$. The case where (\mathfrak{B}, \bar{d}) is chosen is similar. \square

Let us take any \mathcal{FO} -formula ψ . We say that ψ is *invariant under \mathcal{AF} -bisimulations* just in case, for all pointed pairs of structures (\mathfrak{C}, \bar{c}) and (\mathfrak{D}, \bar{d}) , we have that $(\mathfrak{C}, \bar{c}) \simeq^{\mathcal{AF}} (\mathfrak{D}, \bar{d})$ implies $\mathfrak{C} \models \psi[\bar{c}] \Leftrightarrow \mathfrak{D} \models \psi[\bar{d}]$. The following lemma attests that our notion of \mathcal{AF} -bisimulations captures which \mathcal{FO} -definable properties are in the adjacent fragment.

Lemma 5.3.2. *Suppose that $\psi \in \mathcal{FO}$ is invariant under \mathcal{AF} -bisimulations. Then ψ is logically equivalent to an \mathcal{AF} -formula.*

Proof. In case ψ is inconsistent, then the required formula is \perp . Thus let us suppose the opposite. Defining Ψ to be the \mathcal{AF} consequences of ψ ; i.e.:

$$\Psi := \{\theta \in \mathcal{AF} \mid \psi \models \theta\}.$$

We show that $\Psi \models \psi$. By our assumption that ψ is consistent so is Ψ . Thus, let (\mathfrak{C}, \bar{c}) be any pointed structure satisfying Ψ and write

$$\Gamma := \{\theta \in \mathcal{AF} \mid \mathfrak{C} \models \theta[\bar{c}]\}.$$

Clearly, $\Gamma \cup \{\psi\}$ is consistent as otherwise, by compactness, there is a finite subset $\gamma \subseteq \Gamma$ such that $\gamma \models \neg\psi$. Identifying γ as the conjunction of its (finitely many) formulas we have $\psi \models \neg\gamma$. Then, by definition, $\neg\gamma \in \Psi$ thus contradicting $\mathfrak{C} \models \gamma[\bar{c}]$. With that we may find a pointed structure (\mathfrak{D}, \bar{d}) such that $\mathfrak{D} \models \Gamma[\bar{d}]$ and $\mathfrak{D} \models \psi[\bar{d}]$. Clearly, $(\mathfrak{C}, \bar{c}) \equiv^{\mathcal{AF}} (\mathfrak{D}, \bar{d})$. Thus, writing \mathfrak{C}' and \mathfrak{D}' for the ω -saturated extensions of \mathfrak{C} and \mathfrak{D} respectively, we have, by Lemma 5.3.1, that $(\mathfrak{C}', \bar{c}) \simeq^{\mathcal{AF}} (\mathfrak{D}', \bar{d})$. By our assumption of invariance for ψ , we have $\mathfrak{C}' \models \psi[\bar{c}]$ and thus $\mathfrak{C} \models \psi[\bar{c}]$. Since (\mathfrak{C}, \bar{c}) was chosen arbitrarily, we conclude that $\Psi \models \psi$.

To see that ψ is equivalent to a formula in \mathcal{AF} we need only apply compactness. Indeed, we have that there is a finite subset $\psi' \subseteq \Psi$ such that $\psi' \models \psi$. On the other hand, $\psi \models \psi'$ by definition. Thus, the conjunction of elements in ψ' is then the required \mathcal{AF} -formula. \square

We are now at a position to compare the expressive power of \mathcal{AF} to that of \mathcal{FO}^2 . In particular, we will show the following:

Theorem 5.3.3. *Let ψ be an \mathcal{FO} -formula with no more than 2 free variables. If ψ is in \mathcal{FO}^2 , then it is logically equivalent to an \mathcal{AF} -formula. If ψ is in \mathcal{AF} and over a signature of nullary, unary and binary predicates, then it is logically equivalent to an \mathcal{FO}^2 -formula.*

The proof of the above theorem is split into the following two lemmas. To show the claim that \mathcal{FO}^2 -formulas are logically equivalent to those in \mathcal{AF} , consider any $\psi \in \mathcal{FO}^2$. By Lemma 5.3.2 we need only guarantee that ψ is invariant under \mathcal{AF} -bisimulations. Since ψ is invariant under \mathcal{FO}^2 -bisimulations (Lemma 2.3.2), this amounts to showing that \mathcal{AF} -bisimilar structures are also \mathcal{FO}^2 -bisimilar. To see why this condition is sufficient we proceed by contradiction. Suppose $(\mathfrak{A}, \bar{a}) \simeq^{\mathcal{AF}} (\mathfrak{B}, \bar{b})$, but $\mathfrak{A} \models \psi[\bar{a}]$ whilst $\mathfrak{B} \not\models \psi[\bar{b}]$. Then, our supposition that \mathcal{AF} -bisimulations imply \mathcal{FO}^2 -bisimulations will result in ψ not being invariant under \mathcal{FO}^2 -bisimulations. Thus, by Lemma 2.3.1, $\psi \notin \mathcal{FO}^2$ – a contradiction to our initial assumption. We proceed by proving the missing implication:

Lemma 5.3.4. *If $(\mathfrak{A}, \bar{a}) \simeq^{\mathcal{AF}} (\mathfrak{B}, \bar{b})$ with $|\bar{a}| = |\bar{b}| \leq 2$, then $(\mathfrak{A}, \bar{a}) \simeq^{\mathcal{FO}^2} (\mathfrak{B}, \bar{b})$.*

Proof. Taking the premises above we claim that

$$\mathbf{S} := \bigcup_{\ell \in [0,2]} \{(\bar{c}, \bar{d}) \in A^\ell \times B^\ell \mid (\mathfrak{A}, \bar{c}) \simeq^{\mathcal{AF}} (\mathfrak{B}, \bar{d})\}$$

is an \mathcal{FO}^2 -bisimulation between the two pointed structures called to question. To see this take any $(\bar{c}, \bar{d}) \in \mathbf{S}$. By definition, $(\mathfrak{A}, \bar{c}) \simeq^{\mathcal{AF}} (\mathfrak{B}, \bar{d})$. By harmony, $\text{atp}^{\mathfrak{A}}[\bar{c}] = \text{atp}^{\mathfrak{B}}[\bar{d}]$. Since $|\bar{c}| = |\bar{d}| \leq 2$ this implies $\text{tp}^{\mathfrak{A}}[\bar{c}] = \text{tp}^{\mathfrak{B}}[\bar{d}]$.

To see that \mathbf{S} obeys back and forth conditions we resort to the \mathcal{FO}^2 -variant of Ehrenfeucht–Fraïssé games. In particular, we will show that the duplicator can always match the spoiler on \mathbf{S} . Assume that the spoiler chose to play on \mathfrak{A} by keeping the element c_k for $k \in \{1, 2\}$ whilst replacing c_{3-k} by $c' \in A$. By our initial assumption, $(\mathfrak{A}, \bar{c}) \simeq^{\mathcal{AF}} (\mathfrak{B}, \bar{d})$. We may thus, by forth, find some $d' \in A$ such that $(\mathfrak{A}, c_k c') \simeq^{\mathcal{AF}} (\mathfrak{B}, d_k d')$. Clearly, $(c_k c', d_k d') \in \mathbf{S}$. The reasoning is similar if the spoiler plays on \mathfrak{B} . \square

We now show the second statement of Theorem 5.3.3, i.e. that \mathcal{AF} -formulas over binary signatures and with at most 2 free variables are logically equivalent to \mathcal{FO}^2 -formulas. Following a similar approach as before, we show that an \mathcal{FO}^2 -bisimulation over a structure with a binary signature implies an \mathcal{AF} -bisimulation. This will allow us to argue that if an \mathcal{AF} -formula ψ is not invariant under \mathcal{FO}^2 -bisimulations, then it is not invariant under \mathcal{AF} -bisimulations either thus, by Lemma 5.3.1, contradicting the supposition that ψ is in \mathcal{AF} .

Lemma 5.3.5. *If $(\mathfrak{A}, \bar{a}), (\mathfrak{B}, \bar{b})$ with $|\bar{a}| = |\bar{b}| \leq 2$ are over a binary signature, then $(\mathfrak{A}, \bar{a}) \simeq^{\mathcal{FO}^2} (\mathfrak{B}, \bar{b})$ implies $(\mathfrak{A}, \bar{a}) \simeq^{\mathcal{AF}} (\mathfrak{B}, \bar{b})$.*

Proof. Take $(\mathfrak{A}, \bar{a}), (\mathfrak{B}, \bar{b})$ to be pointed structures as described above. For any $\ell < \omega$ say that a tuple $(\bar{c}, \bar{d}) \in A^\ell \times B^\ell$ is a *path* if $(\mathfrak{A}, c_\ell) \simeq^{\mathcal{FO}^2} (\mathfrak{B}, b_\ell)$ and $(\mathfrak{A}, c_i c_{i+1}) \simeq^{\mathcal{FO}^2} (\mathfrak{B}, d_i d_{i+1})$ for each $i \in [1, \ell-1]$. We claim that

$$\mathbf{S} := \bigcup_{\ell < \omega} \{(\bar{c}, \bar{d}) \in A^\ell \times B^\ell \mid (\bar{c}, \bar{d}) \text{ is a path}\}$$

is the required \mathcal{AF} -bisimulation. We proceed by showing harmony by induction on the length $i = |\bar{c}| = |\bar{d}|$ of tuples $(\bar{c}, \bar{d}) \in \mathbf{S}$.

For the base case $i \in \{0, 1, 2\}$, take $(\bar{c}, \bar{d}) \in \mathbf{S}$, where $|\bar{c}| = |\bar{d}| = i$. By definition, $(\mathfrak{A}, \bar{c}) \simeq^{\mathcal{FO}^2} (\mathfrak{B}, \bar{d})$. Thus, $\text{tp}^{\mathfrak{A}}[\bar{c}] = \text{tp}^{\mathfrak{B}}[\bar{d}]$. Clearly, this is equivalent to $\text{atp}^{\mathfrak{A}}[\bar{c}] = \text{atp}^{\mathfrak{B}}[\bar{d}]$ thus securing harmony.

For the inductive step $i+1 \geq 3$ let us take $(\bar{c}', \bar{d}') \in \mathbf{S}$, where $|\bar{c}'| = |\bar{d}'| = i+1$. Since (\bar{c}', \bar{d}') is a path, we have that $(\mathfrak{A}, c_i c'_i) \simeq^{\mathcal{FO}^2} (\mathfrak{B}, d_i d'_i)$ and, by harmony, $\text{tp}^{\mathfrak{A}}[c_i c'_i] = \text{tp}^{\mathfrak{B}}[d_i d'_i]$. Write $\eta := \text{tp}^{\mathfrak{A}}[c_i c'_i]$. By the inductive hypothesis, $\text{atp}^{\mathfrak{A}}[\bar{c}] = \text{atp}^{\mathfrak{B}}[\bar{d}]$. Thus, $\text{atp}^{\mathfrak{A}}[\bar{c} c'_i] = \text{atp}^{\mathfrak{A}}[\bar{c}] \wedge \eta(x_i x_{i+1}) = \text{atp}^{\mathfrak{B}}[\bar{d} d'_i]$ as required.

To see that \mathbf{S} has back and forth properties we will show that the duplicator can answer any move made by the spoiler in the \mathcal{AF} -variant of Ehrenfeucht–Fraïssé games. Take $(\bar{c}, \bar{d}) \in \mathbf{S}$ and suppose that the spoiler chose the structure \mathfrak{A} . Additionally, let the spoiler keep positions $0 \leq i \leq j \leq |\bar{c}|$ and take the element $c' \in A$. Clearly, \mathbf{S} is closed under taking subpaths, thus $(c_i \cdots c_j, d_i \cdots d_j) \in \mathbf{S}$. By definition, $(\mathfrak{A}, c_j) \simeq^{\mathcal{FO}^2} (\mathfrak{B}, d_j)$. By forth (of \mathcal{FO}^2 -bisimulations), the duplicator can find some $d' \in B$ such that $(\mathfrak{A}, c_j c') \simeq^{\mathcal{FO}^2} (\mathfrak{B}, d_j d')$. Then, $(c_i \cdots c_j c', d_i \cdots d_j d')$ is a path and thus a member of \mathbf{S} as required. The case where \mathfrak{B} is chosen is similar. \square

We finish the section with an observation about the length of words and satisfaction of sentences. For this let us write $\mathfrak{A} \simeq_\ell^{\mathcal{AF}} \mathfrak{B}$ if the spoiler cannot win in the first ℓ turns of the \mathcal{AF} -variant of the Ehrenfeucht–Fraïssé games. Given Lemma 5.3.1, proving the following lemma is an easy exercise:

Lemma 5.3.6. *\mathfrak{A} and \mathfrak{B} satisfy the same \mathcal{AF} -sentences of quantifier rank at most ℓ if and only if $\mathfrak{A} \simeq_\ell^{\mathcal{AF}} \mathfrak{B}$.*

Say that a tuple \bar{a} has *primitive length* ℓ just in case \bar{c} is a primitive generator of \bar{a} and $|\bar{c}| = \ell$. A straightforward application of Lemmas 5.2.3 and 5.3.6 result in the following:

Lemma 5.3.7. *Take \mathfrak{A} and \mathfrak{A}' to be structures over a common domain A and suppose that $\text{atp}^{\mathfrak{A}}[\bar{c}] = \text{atp}^{\mathfrak{A}'}[\bar{c}]$ for each primitive ℓ -tuple $\bar{c} \in A^\ell$. Then, $\mathfrak{A} \simeq_{\ell}^{\mathcal{AF}} \mathfrak{A}'$.*

Thus, when considering sentences of quantifier rank at most ℓ , we may, without loss of generality, draw attention to structures \mathfrak{A} which have $\text{atp}^{\mathfrak{A}}[\bar{a}]$ undefined for tuples \bar{a} over A of primitive length greater than ℓ . We will say that such structures are of *primitive height* ℓ .

5.4 Upper bounds for \mathcal{AF}^ℓ without equality

For the following two sections let us fix $\ell \geq 2$. We will now establish an upper-bound for the satisfiability problem of the adjacent fragment. At the top level, our decision procedure is as follows. Given an $\mathcal{AF}^{\ell+1}$ -sentence we will compute an exponentially larger (in terms of formula size and signature) \mathcal{AF}^ℓ -formula that is satisfiable if and only if the original formula is. By repeating this reduction on output formulas, we will eventually reach an \mathcal{FO}^2 -formula which we can then check for satisfiability as given in Theorem 2.2.2. Our approach will not only guarantee that the satisfiability problem of \mathcal{AF}^ℓ (with equality) is in $(\ell-1)$ -NEXPTIME, but also that \mathcal{AF} has the finite model property.

To streamline the main ideas of our procedure, we will, in the current section, focus on the equality-free subfragment of \mathcal{AF} . To avoid confusion, we will refer to this subfragment by lowercase cursive letters – af . Our simplified approach is not without its own merit. As will become apparent in Section 5.4.1, the satisfiability problem for af^3 has a satisfiability problem that is in NEXPTIME. Thus, by ending our variable reduction a step earlier, we will be able to show that the satisfiability problem for af^ℓ is in $(\ell-2)$ -NEXPTIME for $\ell \geq 3$. As an aid we provide recapitulated ideas of previous and current sections in the form of Table 5.4.1. (There is no need to read them just yet!)

Let us, for the rest of the section, fix some normal-form $\mathit{af}^{\ell+1}$ -formula φ over a function- and constant-free signature Σ . Furthermore, when speaking about adjacent types, let us disregard the equality predicate. As promised, we will produce an af^ℓ -sentence ψ of the form

$$\text{acl}(\varphi) \wedge \psi_1 \wedge \psi_2 \wedge \psi_3 \wedge \psi_4 \tag{\psi}$$

that is satisfiable if and only if φ is. In fact, we will have:

Functions $f : [1, m] \rightarrow [1, k]$ and tuples $\bar{a} = a_1 a_2 \cdots a_k$	
\mathbf{A}_k^m	the set of all walks $f : [1, m] \rightarrow [1, k]$
$\bar{\mathbf{A}}_k^m$	the set of all final walks $f : [1, m] \rightarrow [1, k]$
f^+	$f \cup \{m+1 \mapsto k+1\}$ only if $f \in \bar{\mathbf{A}}_k^m$
\mathbf{x}_k	$x_1 x_2 \cdots x_k$
$\tilde{\bar{a}}$	reversal of \bar{a}
\bar{a}^f	$a_{f(1)} a_{f(2)} \cdots a_{f(m)}$
Formulas $\chi \in \mathcal{AF}^{[m]}$ with $f : [1, m] \rightarrow [1, k]$	
$\text{sig}(\chi)$	the signature of χ
χ^{-1}	$\chi(x_m x_{m-1} \cdots x_1)$
$\hat{\chi}$	$\chi \wedge \chi^{-1}$
χ^f	$\chi(\mathbf{x}_k^f)$
Normal-form φ (in $(\ell+1)$ -variables)	
φ	$\forall \mathbf{x}_{\ell+1} \beta \wedge \bigwedge_{t \in T} \forall \mathbf{x}_\ell \exists x_{\ell+1} \gamma_t$
$\text{acl}(\varphi)$	$\bigwedge_{k=1}^\ell \bigwedge_{g \in \mathbf{A}_k^{\ell+1}} \forall \mathbf{x}_k \beta^g \wedge \bigwedge_{t \in T} \bigwedge_{k=1}^{\ell-1} \bigwedge_{f \in \bar{\mathbf{A}}_k^\ell} \forall \mathbf{x}_k \exists x_{k+1} \gamma_t^{f^+}$
Adjacent k -types χ and the k -type of a k -tuple \bar{a} in \mathfrak{A}	
ATP_k^Σ	the set of all adjacent k -types over Σ
$\text{atp}_k^\mathfrak{A}[\bar{a}]$	the adjacent k -type of \bar{a} in \mathfrak{A}
$\text{itp}_k^\mathfrak{A}[\bar{a}]$	the incremental k -type of \bar{a} in \mathfrak{A}
$\partial\chi$	the incremental k -type included in χ
χ^+	$\chi(x_2 x_3 \cdots x_{k+1})$
Surjective $f, g \in \mathbf{A}_k^m$, incremental k -types ι , quantifier-free $\chi \in \mathcal{AF}^{[k]}$	
$\mathbf{D}_{\ell-1}^\circ$	all pairs $\langle i, j \rangle$ for $1 \leq i \leq j \leq \ell-1$ with $j-i+1$ odd
D^+	$\{\langle i+1, j+1 \rangle \mid \langle i, j \rangle \in D\}$
R^*	the equivalence closure of a binary relation R
$f \stackrel{D}{=} g$	$\langle f(i), g(i) \rangle \in D^*$ for all $i \in [1, m]$
ι is D -compatible	$f \stackrel{D}{=} g$ implies $\iota \models p(\mathbf{x}_k^f) \leftrightarrow p(\mathbf{x}_k^g)$ for all $p \in \text{sig}(\chi)$, all $f, g \in \mathbf{A}_k^m$
χ is D -consistent	there exists an $\xi \in \text{ATP}_k^\Sigma$ s.t. $\xi \models \chi$ and $\partial\xi$ is D -compatible
Construction of ψ from φ	
$d_k(\mathbf{x}_k)$	atom implying that \mathbf{x}_k is a palindrome
$\delta_D(\mathbf{x}_\ell)$	$\bigwedge \{d_{j-i+1}(x_i \cdots x_j) \mid \langle i, j \rangle \in D\}$
$p_\zeta(\mathbf{x}_{\ell-1})$	atom implying there is some x s.t. $x\mathbf{x}_{\ell-1}$ realizes ζ
$\mathfrak{B} \times H$	the model obtained by cloning elements of \mathfrak{B} for each $h \in H$

Table 5.4.1: Quick reference guide for Sec 5.1, 5.2 and 5.4.

- if $\mathfrak{A} \models \varphi$, then \mathfrak{A} can be expanded into a model \mathfrak{A}^+ of ψ .
- if $\mathfrak{B} \models \psi$, then $\mathfrak{B} \times T \times [1, \ell^2 + \ell + 1]^{\ell+1}$ can be transformed into $\mathfrak{C}^+ \models \varphi$.

Fix $\mathfrak{A} \models \varphi$. The ensuing definition of ψ depends on the syntactic properties φ and is independent of \mathfrak{A} . We proceed by defining ψ_1, \dots, ψ_4 and expanding \mathfrak{A} into a model $\mathfrak{A}^+ \models \psi$.

For each $s \in \mathbb{N}$ satisfying $1 \leq 2s+1 \leq \ell-1$ let \mathbf{d}_{2s+1} be a fresh $(2s+1)$ -ary predicate. We identify palindromic words of odd length by setting $\mathfrak{A}^+ \models \mathbf{d}_{2s+1}[\bar{a}]$ for all $1 \leq 2s+1 \leq \ell-1$ and all $\bar{a} \in A^s$ that are palindromes. Clearly, \mathfrak{A}^+ models the following \mathbf{af}^ℓ -sentence:

$$\bigwedge_{s=1}^{2s+1 \leq \ell-1} \forall \mathbf{x}_{s+1} \mathbf{d}_{2s+1}(x_1 \cdots x_s x_{s+1} x_s \cdots x_1). \quad (\psi_1)$$

To see the significance of the above formula take any $\mathfrak{B} \models \psi_1$ and primitive $\bar{c} \in B^{\ell+1}$. Suppose, now, that $\langle i, j \rangle$ is a defect of \bar{c} ; that is to say, $c_i \cdots c_j$ is palindromic. Clearly, neither $i = 1$, nor $j = \ell+1$ as that would contradict primitiveness. Thus, writing $s = j - i + 1$, we have that $s \leq \ell-1$ and, by ψ_1 , that $\mathfrak{B} \models \mathbf{d}_s[c_i \cdots c_j]$. That is to say, ψ_1 is an $\mathbf{af}^{\ell-1}$ -sentence, which identifies the defects of primitive $(\ell+1)$ -tuples in its models.

Now, for each $\zeta \in \text{ATP}_\ell^\Sigma$, let us introduce a fresh $(\ell-1)$ -ary predicate \mathbf{p}_ζ . In \mathfrak{A}^+ we then set $\mathbf{p}_\zeta^{\mathfrak{A}^+} := \{\bar{a} \in A^{\ell-1} \mid \mathfrak{A} \models \zeta[a\bar{a}]\}$ for some $a \in A$ for each $\zeta \in \text{ATP}_\ell^\Sigma$. That is, \mathbf{p}_ζ remembers which $(\ell-1)$ -tuples can be extended (to the left) to realise the adjacent ℓ -type ζ . It should be clear that \mathfrak{A}^+ is then a model of the following:

$$\bigwedge_{\zeta \in \text{ATP}_\ell^\Sigma} \forall \mathbf{x}_\ell (\zeta \rightarrow \mathbf{p}_\zeta(x_2 \cdots x_\ell)). \quad (\psi_2)$$

The definition of the final two formulas requires additional observations regarding adjacent $(\ell+1)$ -types. Suppose we are given some Σ -structure \mathfrak{B} and a primitive $(\ell+1)$ -tuple \bar{a} with $\mathbf{atp}^\mathfrak{B}[a_1 \cdots a_\ell] = \zeta$ and $\mathbf{atp}^\mathfrak{B}[a_2 \cdots a_{\ell+1}] = \eta$. The following question then arises: what adjacent $(\ell+1)$ -types can \bar{a} realise in \mathfrak{B} ? Clearly, any adjacent $(\ell+1)$ -type ξ attributed to \bar{a} in \mathfrak{B} should satisfy the following: if $\zeta' \in \text{ATP}_\ell^\Sigma$ is such that $\xi \models \zeta'$, then $\zeta' = \zeta$. A similar observation holds for η^+ . Contrary to what was stated in [Bednarczyk et al., 2023, Section 4], these conditions are insufficient. The counter-example relies on the already mentioned fact that, whilst primitive generators are unique (up to reversal), modes

of generation are not. Recall the following two walks from Section 5.2:

$$\begin{aligned} f &:= (1 \ 2 \ 3 \ 2 \ 3 \ 4 \ 5), \\ g &:= (1 \ 2 \ 3 \ 4 \ 3 \ 4 \ 5). \end{aligned}$$

If \bar{a} is the primitive string $cabad$, then $\bar{a}^f = \bar{a}^g = cababad$. Now, let us take an adjacent 5-type ξ that contains both $\mathbf{r}(\mathbf{x}_5^f)$ and $\neg\mathbf{r}(\mathbf{x}_5^g)$. Notice that the two literals do not contradict one another as $\mathbf{x}_5^f = x_1x_2x_3x_2x_3x_4x_5$, whilst $\mathbf{x}_5^g = x_1x_2x_3x_4x_3x_4x_5$. But, if we were to have $\text{atp}^{\mathfrak{B}}[\bar{a}] = \xi$, then, by Lemma 5.2.3, $\mathfrak{B} \models \mathbf{r}[cababad]$ and $\mathfrak{B} \models \neg\mathbf{r}[cababad]$ – a contradiction.

We return to the general setting and address this issue as follows. Let us write $\mathbf{D}_{\ell-1}^\circ$ for the set $\{(i, j) \mid 1 \leq i \leq j \leq \ell-1 \text{ with } j-i+1 \text{ odd}\}$. Now, let $D \subseteq \mathbf{D}_{\ell-1}^\circ$ be some set of defects. Recall that $f \stackrel{D}{=} g$ for two walks $f, g : [1, m] \rightarrow [1, \ell+1]$ just in case, for each $i \in [1, m]$, we have that $\langle f(i), g(i) \rangle$ is in the equivalence closure of D . Taking an incremental $(\ell+1)$ -type ι we say that ι is *D-compatible* just in case $f \stackrel{D}{=} g$ implies $\iota \models \mathbf{r}(\mathbf{x}_{\ell+1}^f) \leftrightarrow \mathbf{r}(\mathbf{x}_{\ell+1}^g)$ for all m -ary symbols \mathbf{r} featured in ι and all surjective walks $f, g : [1, m] \rightarrow [1, \ell+1]$. The following will help us identify which adjacent $(\ell+1)$ -types can be attributed to particular $(\ell+1)$ -tuples:

Lemma 5.4.1. *Let \bar{c} be a primitive $(\ell+1)$ -tuple over B , and D the defect set of \bar{c} . If \mathfrak{B} is a structure with domain B , then $\iota = \text{itp}^{\mathfrak{B}}[\bar{c}]$ is D -compatible. Conversely, if ι is a D -compatible incremental $(\ell+1)$ -type over some signature Σ , then there is a Σ -structure \mathfrak{B} over B such that $\text{itp}^{\mathfrak{B}}[\bar{c}] = \iota$.*

Proof. The first statement of the lemma is almost immediate. Fix a predicate \mathbf{r} of arity m interpreted by \mathfrak{B} , and suppose $f, g : [1, m] \rightarrow [1, \ell+1]$ are surjective walks. We must show that $f \stackrel{D}{=} g$ implies $\iota \models \mathbf{r}(\mathbf{x}_{\ell+1}^f) \leftrightarrow \mathbf{r}(\mathbf{x}_{\ell+1}^g)$. But by Lemma 5.2.2, if $f \stackrel{D}{=} g$, then $\bar{c}^f = \bar{c}^g$, hence, by Lemma 5.2.3, $\mathfrak{B} \models \mathbf{r}[\bar{c}^f] \Leftrightarrow \mathfrak{B} \models \mathbf{r}[\bar{c}^g]$. Thus, $\iota \models \mathbf{r}(\mathbf{x}_{\ell+1}^f) \leftrightarrow \mathbf{r}(\mathbf{x}_{\ell+1}^g)$.

For the second statement, define \mathfrak{B} over the domain B by setting, for any predicate \mathbf{r} of Σ having arity, say, m :

$$\mathbf{r}^{\mathfrak{B}} := \{\bar{c}^f \mid \text{there is a surjective walk } f : [1, m] \rightarrow [1, \ell+1] \text{ s.t. } \iota \models \mathbf{r}(\mathbf{x}_{\ell+1}^f)\}.$$

To show that $\mathfrak{B} \models \iota[\bar{c}]$, fix any $\mathbf{r} \in \Sigma$ with arity m . If ι contains the atom $\alpha := \mathbf{r}(\mathbf{x}_{\ell+1}^f)$, then it is immediate from the construction of \mathfrak{B} that $\mathfrak{B} \models \alpha[\bar{c}]$. It remains to show that if ι contains the negated atom $\nu := \neg\mathbf{r}(\mathbf{x}_{\ell+1}^f)$, then $\mathfrak{B} \models \nu[\bar{c}]$.

Suppose otherwise. From the construction of \mathfrak{B} , we have $\iota \models \mathbf{r}(\mathbf{x}_{\ell+1}^g)$ for some surjective walk $g: [1, m] \rightarrow [1, \ell+1]$ such that $\bar{c}^f = \bar{c}^g$. By Lemma 5.2.2, $f \stackrel{D}{=} g$. Yet ι is by assumption D -compatible, whence $\iota \models \mathbf{r}(\bar{x}^f) \leftrightarrow \mathbf{r}(\bar{x}^g)$, contradicting the fact that ξ contains both $\neg \mathbf{r}(\mathbf{x}_{\ell+1}^f)$ and $\mathbf{r}(\mathbf{x}_{\ell+1}^g)$. \square

We return to the definition of the final two sentences of ψ . When considering any subset $D \subseteq \mathbf{D}_{\ell-1}^\circ$ let us write D^+ for the result of incrementing both components of every pair in D by 1, i.e. $D^+ := \{\langle i+1, j+1 \rangle \mid \langle i, j \rangle \in D\}$. Now, let us write δ_D for the formula:

$$\bigwedge \{d_{j-i+1}(x_i \cdots x_j) \mid \langle i, j \rangle \in D\}. \quad (\delta_D)$$

Clearly, if $\mathfrak{B} \models \psi_1$, then each primitive $(\ell+1)$ -tuple \bar{c} over B satisfies $\mathfrak{B} \models \delta_D[\bar{c}]$, where D is the set of defects of \bar{c} . Thus, we are now in a position to test if the defects of a given tuple are contained in a given set. Take a quantifier-free $\mathbf{af}^{\ell+1}$ -formula θ over Σ and a set of defects $D \subseteq \mathbf{D}_{\ell-1}^\circ$. We say that θ is D -consistent if there is an adjacent $(\ell+1)$ -type $\xi \models \theta$ such that $\partial\xi$ is D -compatible. With that, we define the final two formulas as follows:

$$\bigwedge_{t \in T} \bigwedge_{\zeta \in \text{ATP}_\ell^\Sigma} \bigwedge_{D \subseteq \mathbf{D}_{\ell-1}^\circ} \bigwedge_{\forall \mathbf{x}_{\ell-1}} \exists x_\ell \left((\delta_D \wedge \mathbf{p}_\zeta(\mathbf{x}_{\ell-1})) \rightarrow \bigvee \{ \eta \in \text{ATP}_\ell^\Sigma \mid (\zeta \wedge \widehat{\beta} \wedge \gamma_t \wedge \eta^+) \text{ is } D^+\text{-consistent} \} \right) \quad (\psi_3)$$

$$\bigwedge_{\zeta \in \text{ATP}_\ell^\Sigma} \bigwedge_{D \subseteq \mathbf{D}_{\ell-1}^\circ} \bigwedge_{\forall \mathbf{x}_\ell} \left((\delta_D \wedge \mathbf{p}_\zeta(\mathbf{x}_{\ell-1})) \rightarrow \bigvee \{ \eta \in \text{ATP}_\ell^\Sigma \mid (\zeta \wedge \widehat{\beta} \wedge \eta^+) \text{ is } D^+\text{-consistent} \} \right). \quad (\psi_4)$$

To see that $\mathfrak{A}^+ \models \psi_3$ take any $\bar{a} \in A^{\ell-1}$ and fix $t \in T$. Suppose that $\mathfrak{A}^+ \models \delta_D[\bar{a}]$, and $\mathfrak{A}^+ \models \mathbf{p}_\zeta[\bar{a}]$ for some $D \subseteq \mathbf{D}_{\ell-1}^\circ$ and $\zeta \in \text{ATP}_\ell^\Sigma$. By our construction, $\text{atp}^{\mathfrak{A}}[a\bar{a}] = \zeta$ for some $a \in A$. Since $\mathfrak{A} \models \varphi$, there is an element $b \in A$ such that $\mathfrak{A} \models \gamma_t[a\bar{a}b]$. Taking $\xi := \text{atp}^{\mathfrak{A}}[a\bar{a}b]$ and $\eta := \text{atp}^{\mathfrak{A}}[\bar{a}b]$, we have $\xi \models \zeta \wedge \widehat{\beta} \wedge \gamma_t \wedge \eta^+$. On the other hand, since $\mathfrak{A}^+ \models \delta_D[\bar{a}]$, we have, by construction of \mathfrak{A}^+ , that the defect set of \bar{a} is D . Hence, the defect set of $a\bar{a}b$ is D^+ . By the first statement of Lemma 5.4.1 we have that $\partial\xi$ is D^+ -compatible. Hence, b is a witness for \bar{a} in regard to the appropriate existential requirement in ψ_3 . Using similar reasoning, we have $\mathfrak{A}^+ \models \psi_4$.

To get a better view of what these formulas capture, take some model \mathfrak{B} of ψ . Since ψ is of quantifier rank ℓ , we may assume, by Lemma 5.3.7, that \mathfrak{B} is

of primitive height ℓ . Thus, when taking a primitive $(\ell+1)$ -tuple $a\bar{a}b$ over B we may freely take $\mathbf{atp}^{\mathfrak{B}}[a\bar{a}b]$ to be undefined. We sketch how one can define the adjacent $(\ell+1)$ -type of $a\bar{a}b$ consistently with the universal requirements of φ . Let $D := \{\langle i, j \rangle \mid 1 \leq i \leq j \leq \ell-1, \mathfrak{B} \models \mathbf{d}_{j-i+1}[a_i \cdots a_j]\}$. It is easy to see, by ψ_1 , that the defects of $a\bar{a}b$ reside in D^+ . By ψ_2 , $\mathfrak{B} \models \mathbf{p}_\zeta[\bar{a}]$, where $\zeta := \mathbf{atp}^{\mathfrak{B}}[a\bar{a}]$. That is, with the help of \mathbf{p}_ζ we are able to remember the adjacent ℓ -type that $a\bar{a}$ realises with only $\ell-1$ variables. Writing $\eta := \mathbf{atp}^{\mathfrak{B}}[\bar{a}b]$ we have, by ψ_4 , that $\zeta \wedge \widehat{\beta} \wedge \eta^+$ is D^+ -consistent. Thus, there is an adjacent $(\ell+1)$ -type ξ such that $\xi \models \zeta \wedge \widehat{\beta} \wedge \eta^+$ and $\partial\xi$ is D^+ -compatible. By the second statement of Lemma 5.4.1, we may consistently assign $\mathbf{itp}^{\mathfrak{B}}[a\bar{a}b] := \partial\xi$ thus securing $\mathfrak{B} \models \widehat{\beta}[a\bar{a}b]$. The conjunct ψ_3 acts similarly with the added benefit of identifying witnesses for $t \in T$. At the moment we are somewhat ill-equipped to deal with the witnesses provided by ψ_3 . Thus, we briefly pause to restate what has already been proved:

Lemma 5.4.2. *Suppose $\mathfrak{A} \models \varphi$. Then we can expand \mathfrak{A} to a model $\mathfrak{A}^+ \models \psi$.*

Before proceeding with the converse direction we will establish two technical lemmas. Recall that if \mathfrak{B} is a Σ -structure, and H a non-empty set of indices, then $\mathfrak{B} \times H$ is a Σ -structure over $B \times H$ defined as follows: for any $\mathbf{r} \in \Sigma$ of arity k , any k -tuples $b_1 \cdots b_k$ over B , and $h_1 \cdots h_k$ over H , set $\mathfrak{B} \times H \models \mathbf{r}[\langle b_1, h_1 \rangle \cdots \langle b_k, h_k \rangle]$ if and only if $\mathfrak{B} \models \mathbf{r}[b_1 \cdots b_k]$. Lemma 3.2.2 states that $\mathfrak{B} \models \chi[\bar{b}]$ if and only if $\mathfrak{B} \models \chi[\langle b_1, h_1 \rangle \cdots \langle b_k, h_k \rangle]$ for each $\bar{b} \in B^k$, $\bar{h} \in H^k$, and first-order formula without equality χ . The following lemma will be used in conjunction with Lemma 3.2.2.

Lemma 5.4.3. *Suppose \bar{b} is a k -tuple over some set B and H is a non-empty set. Then, for any $h_1, \dots, h_k \in H$, if \bar{c} is the tuple $\langle b_1, h_1 \rangle \cdots \langle b_k, h_k \rangle$ over $B \times H$, we have that every defect of \bar{c} is a defect of \bar{b} .*

Proof. Pick some defect $\langle i, j \rangle$ of \bar{c} . Thus, $c_i \cdots c_j$ is a palindrome, or, in other words, $\langle b_{i+m}, h_{i+m} \rangle = \langle b_{j-m}, h_{j-m} \rangle$ for all $0 \leq m \leq \lfloor (j-i)/2 \rfloor$. But then $b_i \cdots b_j$ is certainly a palindrome, and hence $\langle i, j \rangle$ a defect of \bar{b} as required. \square

Finally, the following combinatorial lemma allows us to extend the circular witnessing technique used in Lemma 2.2.1 to multivariable setting.

Lemma 5.4.4. *There is a set H with $|H| = (\ell^2 + \ell + 1)^{\ell+1}$ and a function $g: H^\ell \rightarrow H$ such that, for any tuple $\bar{t} \in H^\ell$ consisting of the elements t_1, \dots, t_ℓ in some order: (i) $g(\bar{t})$ is not in \bar{t} ; (ii) if $\bar{t}' \in H^\ell$ consists of the elements $\{t_2, \dots, t_\ell, g(\bar{t})\}$ in some order, then $g(\bar{t}')$ is not in \bar{t} either.*

Proof. Let $z = \ell^2 + \ell + 1$ and $H = [1, z]^{\ell+1}$. Thus, $|H|$ is as required. Writing any element of H as the word $i\bar{s}$, where $i \in [1, z]$ and $\bar{s} \in [1, z]^\ell$, let g be defined by $g(i_1\bar{s}_1, \dots, i_\ell\bar{s}_\ell) = i_0i_1 \cdots i_\ell$, where i_0 is the smallest positive integer not in the set $S = \{i_1, \dots, i_\ell\} \cup \bar{s}_1 \cup \dots \cup \bar{s}_\ell$. (For brevity, we are here identifying words over the integers with the sets of their members.) Note that $i_0 \in [1, z]$ since $|S| < z$. Now let some tuple $\bar{t} \in H^\ell$ be given, consisting of words t_1, \dots, t_ℓ in some order, and write $t_h = i_h\bar{s}_h$ (for all $1 \leq h \leq \ell$), where $i_h \in [1, z]$ and $\bar{s}_h \in [1, z]^\ell$. Thus, $t = g(\bar{t})$ is a word of the form $i_0\bar{s}$, where \bar{s} consists of i_1, \dots, i_ℓ in some order, and i_0 does not occur anywhere in \bar{t} . Condition (i) is then immediate because of the choice of i_0 . For condition (ii), we observe that $t' = g(\bar{t}')$ is a word of the form $i'\bar{s}'$, where \bar{s}' consists of i_0, i_2, \dots, i_ℓ in some order, and i' does not occur in any of the words in \bar{t}' . By the choice of i' , it is immediate that $t' \notin \{t_2, \dots, t_\ell\}$. But the value i_1 (the first letter of t_1) occurs in $g(\bar{t})$ (which belongs to the tuple \bar{t}'), whence $i' \neq i_1$. It follows that $t' \neq t_1$, whence t' is not in \bar{t} , as required. \square

With the above, we are now ready to show the converse of Lemma 5.4.2:

Lemma 5.4.5. *Suppose $\mathfrak{B} \models \psi$. Then there is a model $\mathfrak{C}^+ \models \varphi$ such that $|C^+|/|B| \leq |T| \cdot (\ell^2 + \ell + 1)^{\ell+1}$.*

Proof. Recall that φ is a Σ -sentence of the form (**FLPC-nmf**) recapitulated here:

$$\forall \mathbf{x}_{\ell+1} \beta \wedge \bigwedge_{t \in T} \forall \mathbf{x}_\ell \exists x_{\ell+1} \gamma_t.$$

Writing \mathfrak{B}^- for the Σ -reduct of \mathfrak{B} , we define \mathfrak{C} to be $\mathfrak{B}^- \times T \times H$, where H is some set of cardinality $(\ell^2 + \ell + 1)^{\ell+1}$. It should be clear from Lemma 3.2.2 that $\mathfrak{C} \models \text{acl}(\varphi)$. Assuming, as allowed by Lemma 5.3.7, that \mathfrak{C} is of primitive height ℓ , we lift \mathfrak{C} to \mathfrak{C}^+ by providing incremental $(\ell+1)$ -types for primitive $(\ell+1)$ -tuples.

Let us consider first the model \mathfrak{B} and identifying potential witnesses for the existential requirements $t \in T$. To this end, take some $a\bar{a} \in B^\ell$ and denote the defect set of \bar{a} by D . Writing $\zeta := \text{atp}^{\mathfrak{B}^-}[a\bar{a}]$ we have $\mathfrak{B} \models \delta_D[\bar{a}]$ by ψ_1 , and $\mathfrak{B} \models \text{p}_\zeta[\bar{a}]$ by ψ_2 . Then, by ψ_3 , we may identify, for each $t \in T$, an element $b_{a\bar{a},t}$ having $\eta := \text{atp}^{\mathfrak{B}^-}[\bar{a}b_{a\bar{a},t}]$ and such that $\zeta \wedge \hat{\beta} \wedge \gamma_t \wedge \eta^+$ is D^+ -consistent. By D^+ -consistency, there is an adjacent $(\ell+1)$ -type $\xi_{a\bar{a},t}$ such that $\partial\xi_{a\bar{a},t}$ is D^+ -compatible and $\xi_{a\bar{a},t} \models \zeta \wedge \hat{\beta} \wedge \gamma_t \wedge \eta^+$. If it should transpire that $a\bar{a}b_{a\bar{a},t}$ is primitive, we could, by Lemma 5.4.1, assign $a\bar{a}b_{a\bar{a},t}$ the incremental $(\ell+1)$ -type $\partial\xi_{a\bar{a},t}$ and thus provide a witness for $a\bar{a}$ in regard to the existential requirement $\exists x_{\ell+1} \gamma_t$ of φ .

Such a hasty decision, however, might invite contradictions down the line as, for instance, we could have that $b_{a\bar{a},t} = b_{a\bar{a},t'}$, but $\xi_{a\bar{a},t} \neq \xi_{a\bar{a},t'}$ for distinct $t, t' \in T$. Moreover, writing $b' = b_{a\bar{a},t}$, it might be the case that $a = b_{b'\bar{a},t'}$, but $\xi_{a\bar{a},t} \neq \xi_{b'\bar{a},t'}^{-1}$ for some $t, t' \in T$. That is to say, when considering the tuple $b_{a\bar{a},t}\bar{a} = b'\bar{a}$ the element a might be picked as a witness but without guarantee that the adjacent $(\ell+1)$ -types $\xi_{a\bar{a},t}$ and $\xi_{b'\bar{a},t'}$ are inverses of one another. To address these issues we shift our focus to the “blown-up” structure \mathfrak{C} .

We proceed by introducing a series of witnessing functions $w_t : C^\ell \rightarrow C$ satisfying the following properties for all $\bar{c} \in C^\ell$:

1. $w_t(\bar{c}) \neq w_{t'}(\bar{c})$ for distinct $t, t' \in T$,
2. $w_t(\bar{c})$ is not amongst the elements of \bar{c} , and
3. if $\bar{c}' \in C^\ell$ is a word over $\{c_2, \dots, c_\ell, w_t(\bar{c})\}$, then $w_t(\bar{c}')$ is not an element of \bar{c} .

To define w_t take $\bar{c} \in C^\ell$. Recalling that $C = B \times T \times H$ we may write $c_i = (a_i, t_i, h_i)$ for $i \in [1, \ell]$. For convenience, we will write \bar{a} for $a_1 \cdots a_\ell$ and \bar{h} for $h_1 \cdots h_\ell$. Since $|H| = (\ell^2 + \ell + 1)^{\ell+1}$, let $g : H^\ell \rightarrow \ell$ be the function guaranteed by Lemma 5.4.4. Then, for all $t \in T$, we define $w_t(\bar{c}) := (b_{\bar{a},t}, t, g(\bar{h}))$. Clearly, condition 1 is satisfied as the second component of $w_t(\bar{c})$ and $w_{t'}(\bar{c})$ only coincide when $t = t'$. Conditions 2 and 3 are satisfied by the properties of g in Lemma 5.4.4.

We are now in a position to provide witnesses for ℓ -tuples in \mathfrak{C}^+ in regard to the existential requirements of φ . Let us fix some $t \in T$ and take $\bar{c} \in C^\ell$. We will write $d = w_t(\bar{c})$, $c_i = (a_i, t_i, h_i)$ and $\bar{a} = a_1 \cdots a_\ell$. Let us first suppose that $\bar{c}d$ is primitive. By the definition of w_t , the element d is the tuple $(b_{\bar{a},t}, t, h)$. Writing $\zeta := \text{atp}^{\mathfrak{B}^-}[\bar{a}]$, $\eta := \text{atp}^{\mathfrak{B}^-}[a_2 \cdots a_\ell b_{\bar{a},t}]$, and D for the defects of $a_2 \cdots a_\ell$, let us recall that $\xi_{\bar{a},t}$ is an adjacent $(\ell+1)$ -type that entails $\zeta \wedge \hat{\beta} \wedge \gamma_t \wedge \eta^+$ and is D^+ -consistent. Notice that, by Lemma 3.2.2, $\text{atp}^{\mathfrak{C}}[\bar{c}] = \zeta$ and $\text{atp}^{\mathfrak{C}}[c_2 \cdots c_\ell d] = \eta$. By Lemma 5.4.3 the defects of $c_2 \cdots c_\ell$ are included in D . By primitiveness of $\bar{c}d$, neither $\langle 1, n \rangle$, nor $\langle j, \ell+1 \rangle$ are defects of \bar{c} for any $i > 1$ and $j < \ell+1$. Thus, the defects of $\bar{c}d$ are included in D^+ . Thus, by the second statement of Lemma 5.4.1, the tuple $\bar{c}d$ can be assigned the incremental $(\ell+1)$ -type $\partial\xi_{\bar{a},t}$ without introducing contradictions. We do just that in \mathfrak{C}^+ . Setting $\text{itp}^{\mathfrak{C}^+}[\bar{c}d] := \partial\xi_{\bar{a},t}$ we secure $\text{atp}^{\mathfrak{C}^+}[\bar{c}d] = \zeta \cup \eta^+ \cup \partial\xi_{\bar{a},t}$ and thus $\mathfrak{C}^+, \bar{c} \models \exists x_{\ell+1} \gamma_t$.

On the other hand, still keeping $\bar{c} \in C^\ell$, $t \in T$ and $d = w_t(\bar{c})$ as before, let us assume that $\bar{c}d$ is not primitive. But then, by Lemma 5.2.4, neither is \bar{c} as

$w_t(\bar{c}) = d$ is not amongst the elements of \bar{c} . In fact, by Lemma 5.2.4, \bar{c} is generated by some tuple $\bar{e} \in C^k$ and a terminal walk $f : [1, \ell] \rightarrow [1, k]$ with $k < \ell$. But then, by $\text{acl}(\varphi)$, we have that $\mathfrak{C}^+, \bar{e} \models \exists x_{k+1} \gamma_t^{f+}$. Picking $e \in C$ to be the witness for the existential quantifier we have, by Lemma 5.2.3, that $\mathfrak{C} \models \gamma_t[(\bar{e}e)^{f+}]$. That is to say $\mathfrak{C} \models \gamma_t[(\bar{c}e)]$. Thus, the tuple \bar{c} already has a witness for $\exists x_{\ell+1} \gamma_t$ in \mathfrak{C} leaving us with nothing to do.

Let us briefly analyse the operations carried out on \bar{c} . Notice that condition 1 of w_t promises unique witnesses for all $t \in T$. Thus, we may repeat the above procedure on \bar{c} for all $t \in T$ without fear of clashes. In the process of giving an incremental $(\ell+1)$ -type to $\bar{c}d$ notice that we have also provided one for the tuple $d\bar{c}$. That is to say, $\text{itp}^{\mathfrak{C}^+}[d\bar{c}] = (\text{itp}^{\mathfrak{C}^+}[\bar{c}d])^{-1}$. We have to be mindful of this assignment if we wish to provide witnesses for $dc_\ell \cdots c_2$. But, since the witness for $t \in T$ is given by $d' = w_t(dc_\ell \cdots c_2)$, we have, by condition 3, that $d' \neq c_1$. (In fact, d' is not equal to any of the elements in $\bar{c}d$). Thus, we are free to assign an incremental $(\ell+1)$ -type to the tuple $dc_\ell \cdots c_2d'$ so long as it is primitive. Putting the two observations together we may, using the procedure above, provide witnesses for all $\bar{c} \in C^\ell$ and all $t \in T$ in regard to existential requirement $\exists x_{\ell+1} \gamma_t$ of φ .

Since each adjacent $(\ell+1)$ -type assigned to primitive $(\ell+1)$ -tuples entails $\widehat{\beta}$, we have that no newly introduced adjacent $(\ell+1)$ -type violates the universal requirements β of φ . Notice that, by $\text{acl}(\varphi)$, non-primitive $(\ell+1)$ -tuples satisfy β as well. We thus need only deal with primitive $(\ell+1)$ -tuples $\bar{c}d \in C^{\ell+1}$ that have not been considered in the construction above. But this is simple with reference to ψ_4 . Let a_1, \dots, a_ℓ, b be the projection of c_1, \dots, c_ℓ, d to their first component. Writing $\zeta := \text{atp}^{\mathfrak{B}^-}[\bar{a}]$, $\eta := \text{atp}^{\mathfrak{B}^-}[a_2 \cdots a_\ell b]$, and D for the defects of $a_2 \cdots a_\ell$, notice that, by ψ_4 , the formula $\zeta \wedge \widehat{\beta} \wedge \eta^+$ is D^+ -consistent. Thus, let us write ξ for the adjacent $(\ell+1)$ -type entailing the above and with $\partial\xi$ being D^+ -compatible. Notice, again by Lemma 3.2.2, that $\text{atp}^{\mathfrak{C}^+}[\bar{c}] = \zeta$ and $\text{atp}^{\mathfrak{C}^+}[c_2 \cdots c_\ell d] = \eta$. By Lemma 5.4.3 the defects of $c_2 \cdots c_\ell$ are included in D^+ . Thus, the defects of the primitive tuple $\bar{c}d$ are included in D^+ . Lemma 5.4.1 then allows us to set $\text{itp}^{\mathfrak{C}^+}[\bar{c}d] := \partial\xi_{\bar{a},t}$ and conclude that $\text{atp}^{\mathfrak{C}^+}[\bar{c}d] = \xi$. Clearly, then, $\mathfrak{C}^+ \models \widehat{\beta}[\bar{c}d]$. Repeating this assignment for each primitive $(\ell+1)$ -tuple without an incremental $(\ell+1)$ -type will result in $\mathfrak{C}^+ \models \forall \mathbf{x}_{\ell+1} \beta$ and hence $\mathfrak{C}^+ \models \varphi$. \square

Let us take stock. Above we provided a procedure that, when given a normal-form af^ℓ -sentence φ_ℓ over some signature Σ_ℓ , produces an $\text{af}^{\ell-1}$ -sentence $\varphi_{\ell-1} :=$

ψ . By Lemma 5.4.2, every model of φ_ℓ can be expanded into a model of $\varphi_{\ell-1}$. Lemma 5.4.5 establishes the converse; i.e. each model \mathfrak{B} of $\varphi_{\ell-1}$ can be “blown-up” to form a model $\mathfrak{C}^+ \models \varphi_\ell$ over the domain $C := B \times T \times H$, where $|H|$ is at most $(\ell^2 + \ell + 1)^{\ell+1}$. It is easy to verify that $\|\psi_1\|$ is a constant whilst $\|\text{acl}(\varphi)\|$ is bounded by $O(\|\varphi_\ell\|)$. The sizes of ψ_2, ψ_3, ψ_4 are a polynomial function on the cardinality of $\text{ATP}_{\ell-1}^{\Sigma_\ell}$. On first sight, $|\text{ATP}_{\ell-1}^{\Sigma_\ell}|$ is bounded by $2^{2^{O(\|\varphi_\ell\|)}}$ as there are at most $2^{O(\|\varphi_\ell\|)}$ adjacent atoms. We can, however, do better. For this purpose let us take ψ to be any quantifier-free subformula of φ_ℓ and ξ to be an adjacent k -type over Σ_ℓ for some $k \leq \ell$. Clearly, the atoms of ξ that are not mentioned in φ_ℓ play no role in determining if ξ entails ψ . Writing $\text{atoms}(\varphi_\ell)$ for the atoms in φ_ℓ we define the *reduced* adjacent k -type of ξ to be

$$\xi^\downarrow := \{\pm \mathbf{p}(\bar{x}) \in \xi \mid \mathbf{p}(\bar{x}) \in \text{atoms}(\varphi_\ell)^f \text{ for some walk } f : [1, \ell] \rightarrow [1, \ell]\},$$

where $\text{atoms}(\varphi_\ell)^f := \{\alpha^f \mid \alpha \in \text{atoms}(\varphi_\ell)\}$. Clearly, $|\xi^\downarrow|$ is bounded by $O(\|\varphi_\ell\|)$. The following is then almost immediate:

Lemma 5.4.6. *Let ψ be a quantifier-free subformula of φ_ℓ , and take ξ to be an adjacent k -type. Then, for all walks $f : [1, \ell] \rightarrow [1, \ell]$, $\xi \models \psi^f$ iff $\xi^\downarrow \models \psi^f$.*

Proof. The “if” direction of the implication is immediate by the fact that $\xi^\downarrow \subseteq \xi$. For the converse direction assume, without loss of generality, that ψ is in disjunctive normal-form. Supposing, first, that $\xi \models \psi^f$ for some walk $f : [1, \ell] \rightarrow [1, \ell]$, we may find some disjunct θ of ψ such that $\xi \models \theta^f$. We claim that $\xi^\downarrow \models \theta^f$. To see this take any literal α of θ^f . Clearly, $\alpha \in \xi$. But since α is a part of ψ^f , we have, by the definition of reduced adjacent k -types, that $\alpha \in \xi^\downarrow$. \square

Restricting our attention to the reduced adjacent $(\ell-1)$ -types over Σ_ℓ when constructing $\varphi_{\ell-1}$ results in a $2^{\|\varphi_\ell\|^{O(1)}}$ bound on the sizes of ψ_2, ψ_3, ψ_4 . Thus, $2^{\|\varphi_\ell\|^{O(1)}}$ is bound on $\|\varphi_{\ell-1}\|$.

Recall the tower of exponentials function $\mathbf{t} : \mathbb{N} \times \mathbb{N} \rightarrow \mathbb{N}$:

$$\begin{aligned} \mathbf{t}(0, n) &:= n, \\ \mathbf{t}(m+1, n) &:= 2^{\mathbf{t}(m, n)}. \end{aligned}$$

By repeating the reduction on $\varphi_{\ell-1}$ and subsequent formulas we will eventually reach a formula φ_k over Σ_k in an “easy enough” fragment of \mathcal{af} that we can check for satisfiability directly. An easy inductive argument shows that $\|\varphi_k\|$ is bounded

by $\mathfrak{t}(\ell-k, \|\varphi_\ell\|^{O(1)})$. The question is at what step k do we stop? An obvious answer is $k = 2$. This would leave us with $\varphi_2 \in \mathcal{FO}^2$ which we can check for satisfiability in non-deterministic time $2^{\|\varphi_2\|^{O(1)}}$ as per Theorem 2.2.2. But this would leave $\text{Sat}(\mathbf{af}^\ell)$ in $(\ell-1)$ -NEXPTIME – a whole exponential above what was advertised. Anticipating Section 5.4.1 we mention now that satisfiable \mathbf{af}^3 -sentences have models of exponential size in regard to the input. Thus, stopping at $k = 3$ will shave-off an exponential in time complexity making $\text{Sat}(\mathbf{af}^\ell)$ decidable in $(\ell-2)$ -NEXPTIME. In addition, an easy inductive argument then yields that satisfiable \mathbf{af}^ℓ -sentences have a model of size $\mathfrak{t}(\ell-2, \|\varphi\|^{O(1)})$. Noting that $\text{Sat}(\mathbf{af})$ inherits TOWER-hardness from $\text{Sat}(fl)$ we conclude the following:

Theorem 5.4.7. *If an \mathbf{af}^ℓ -sentence is satisfiable, then it is satisfiable in a structure of size at most $\mathfrak{t}(\ell-2, \|\varphi\|^{O(1)})$. Hence, $\text{Sat}(\mathbf{af}^\ell)$ is in $(\ell-2)$ -NEXPTIME for $\ell \geq 3$ and $\text{Sat}(\mathbf{af})$ is TOWER-complete.*

5.4.1 Excursion: the equality-free 3-variable case

As mentioned in the previous section, we can achieve better satisfiability upper-bounds by tackling $\text{Sat}(\mathbf{af}^3)$ directly. We adapt the connector-type approach given in [Pratt-Hartmann et al., 2019] for the fluted fragment to fit the adjacent setting, thus establishing NEXPTIME-completeness for the satisfiability problem of \mathbf{af}^3 . (Readers uninterested in this optimisation can safely skip this subsection).

Let Σ be a function- and constant-free signature. We remind the reader that the equality symbols is not featured any of the following adjacent formulas or types. If π is a 1-type over Σ , we define the 2-type π^2 , over the same signature, to be $\{\alpha \mid \alpha \text{ is a literal in } \mathbf{af}^2 \text{ s.t. } \alpha(x_1x_1) \in \pi\}$. The intuition here is that if π is the type of an element a in some structure, then π^2 is the type of the pair aa . A *star* (over Σ) is a set σ of 2-types over Σ subject to the condition that there exists some 1-type π over Σ such that $\pi^2 \in \sigma$ and $\zeta \models \pi$ for all $\zeta \in \sigma$. This 1-type π is clearly unique, and we denote it by $\text{tp}(\sigma)$. If \mathfrak{A} is any structure interpreting Σ and $a \in A$, then a defines a star σ over Σ as given by $\sigma := \{\text{tp}^{\mathfrak{A}}[ab] \mid b \in A\}$. We refer to σ as *the star of a in \mathfrak{A}* , and denote it $\text{str}^{\mathfrak{A}}[a]$. It follows immediately from the above definitions that $\text{tp}(\text{str}^{\mathfrak{A}}[a]) = \text{tp}^{\mathfrak{A}}[a]$. When speaking of stars, we suppress reference to Σ if irrelevant or clear from context.

Let φ be a normal-form formula of \mathbf{af}^3 , as given in (\mathcal{AF} -nmf), with $\ell = 2$, $\Sigma = \text{sig}(\varphi)$, and without the equality predicate. In the sequel we refer freely

to the subformulas $(\gamma_t)_{t \in T}$ and β of φ . Say that a star σ is φ -compatible if the following conditions hold:

$L\exists_1$: for all $t \in T$, there exists $\eta \in \sigma$ s.t. $\eta \models \gamma_t(x_1x_1x_2)$.

$L\exists_2$: for all ζ such that $\zeta^{-1} \in \sigma$ and all $t \in T$, there exists $\eta \in \sigma$ such that the \mathbf{af}^3 -formula $\zeta \wedge \eta^+ \wedge \gamma_t \wedge \hat{\beta}$ is consistent;

$L\forall_1$: for all $\eta \in \sigma$ and all walks $f : [1, 3] \rightarrow [1, 2]$ we have $\eta \models \beta^f$;

$L\forall_2$: for all ζ such that $\zeta^{-1} \in \sigma$ and all $\eta \in \sigma$, the \mathbf{af}^3 -formula $\zeta \wedge \eta^+ \wedge \hat{\beta}$ is consistent.

Lemma 5.4.8. *If φ is a normal-form \mathbf{af}^3 -formula, $\mathfrak{A} \models \varphi$ and $b \in A$, then $\mathbf{str}^{\mathfrak{A}}[b]$ is compatible with φ .*

Proof. Suppose $b \in A$, and let $\mathbf{str}^{\mathfrak{A}}[b] = \sigma$. For $L\exists_1$, consider $t \in T$. Since $\mathfrak{A} \models \varphi$, there exists $c \in A$ such that $\mathfrak{A} \models \gamma_t[bbc]$. Then $\eta = \mathbf{tp}^{\mathfrak{A}}[bc]$ is as required. For $L\exists_2$, suppose $\zeta^{-1} \in \sigma$; thus, there exists $a \in A$ such that $\mathfrak{A} \models \zeta[ab]$. Now consider $t \in T$. Since $\mathfrak{A} \models \varphi$, there exists $c \in A$ such that $\mathfrak{A} \models \gamma_t[abc]$. Then $\eta = \mathbf{tp}^{\mathfrak{A}}[bc]$ is as required. The conditions $L\forall_1$, $L\forall_2$ are established similarly. \square

A set \mathfrak{c} of stars is said to be *coherent* if the following conditions hold:

$G\exists$: for all $\sigma \in \mathfrak{c}$ and all $\zeta \in \sigma$, there exists $\sigma' \in \mathfrak{c}$ such that $\zeta^{-1} \in \sigma'$;

$G\forall$: for all $\sigma, \sigma' \in \mathfrak{c}$, there exists a 2-type ζ such that $\zeta \in \sigma$ and $\zeta^{-1} \in \sigma'$.

Lemma 5.4.9. *Let \mathfrak{A} be a structure. Then $\mathfrak{c} = \{\mathbf{str}^{\mathfrak{A}}[a] \mid a \in A\}$ is coherent.*

Proof. For $G\exists$, suppose $\sigma \in \mathfrak{c}$, and let $a \in A$ be s.t. $\mathbf{str}^{\mathfrak{A}}[a] = \sigma$. If $\zeta \in \sigma$, there exists $b \in A$ s.t. $\mathbf{tp}^{\mathfrak{A}}[ab] = \zeta$. Then $\sigma' = \mathbf{str}^{\mathfrak{A}}[b]$ is as required. $G\forall$ is similar. \square

Define a *certificate* for φ to be a non-empty, coherent set of stars, all of which are compatible with φ .

Lemma 5.4.10. *Any satisfiable normal-form \mathbf{af}^3 -formula has a certificate \mathfrak{c} such that both $|\mathfrak{c}|$ and $|\bigcup \mathfrak{c}|$ are $2^{O(\|\varphi\|)}$.*

Proof. Suppose $\mathfrak{A} \models \varphi$. We first create a model \mathfrak{A}' of φ over the same domain as \mathfrak{A} in which the number of different 2-types is bounded by $2^{O(\|\varphi\|)}$. To this end, take a 2-type ζ , and recall that the reduced adjacent 2-type of ζ is:

$$\zeta^\downarrow := \{\pm \mathbf{p}(\bar{x}) \in \zeta \mid \mathbf{p}(\bar{x}) \in \text{atoms}(\varphi)^f \text{ for some walk } f : [1, 3] \rightarrow [1, 3]\}.$$

We define \mathfrak{A}' by providing an interpretation for each $\mathbf{p} \in \Sigma$ of arity m :

- $a^m \in \mathbf{p}^{\mathfrak{A}'}$ if and only if $a^m \in \mathbf{p}^{\mathfrak{A}}$, for each $a \in A$,
- $\bar{a}^f \in \mathbf{p}^{\mathfrak{A}'}$ if and only if $\mathbf{p}(\mathbf{x}_2^f) \in (\mathbf{tp}^{\mathfrak{A}}[\bar{a}])^\perp$, for each primitive $\bar{a} \in A^2$ and each surjective walk $f : [1, m] \rightarrow [1, 2]$, and
- $\bar{a}^f \in \mathbf{p}^{\mathfrak{A}'}$ if and only if $\bar{a}^f \in \mathbf{p}^{\mathfrak{A}}$, for each primitive $\bar{a} \in A^3$ and each surjective walk $f : [1, m] \rightarrow [1, 3]$.

By Lemma 5.4.6, no meaningful information is lost in regard to entailment of quantifier-free subformulas of φ . Thus, $\mathfrak{A}' \models \varphi$ as required. Noting that $|\zeta^\perp|$ is bounded $O(\|\varphi\|)$ we have at most $2^{O(\|\varphi\|)}$ different 2-types in \mathfrak{A}' .

Now, let $\mathbf{c}' = \{\text{str}^{\mathfrak{A}'}[a] \mid a \in A\}$. By Lemmas 5.4.8 and 5.4.9, \mathbf{c}' is a certificate for φ . We shall construct a certificate $\mathbf{c} \subseteq \mathbf{c}'$ satisfying the size bound of the lemma. Pick any $\sigma \in \mathbf{c}'$, initialize $C := \{\sigma\}$, and $S := \{\zeta^{-1} \mid \zeta \in \sigma\}$. We shall add stars to the set C and 2-types to the set S (all of them realized in \mathfrak{A}'), maintaining the invariant $S = \{\zeta^{-1} \mid \zeta \in \sigma \text{ for some } \sigma \in C\}$. We call a star σ in C *satisfied* if $\sigma \subseteq S$. Now execute the following procedure until C contains no unsatisfied stars. Pick some unsatisfied $\sigma \in C$ and some $\zeta \in \sigma \setminus S$. By $\text{G}\exists$, there exists a star $\sigma' \in \mathbf{c}'$ such that $\zeta^{-1} \in \sigma'$. Set $C := C \cup \{\sigma'\}$ and $S := S \cup \{\zeta^{-1} \mid \zeta \in \sigma'\}$. These assignments maintain the invariant on C and S . Clearly, this process terminates after $2^{O(\|\varphi\|)}$ steps, since $|S|$ increases by at least one at each step, and $|C|$ by exactly 1. On termination, C is globally coherent, and $|C| \leq 2^{O(\|\varphi\|)}$. Set \mathbf{c} to be the final value of C . \square

We are now in a position to bound the size of models of \mathbf{af}^3 -sentences.

Lemma 5.4.11. *Let φ be a normal-form \mathbf{af}^3 -formula over a signature Σ . If φ is satisfiable, then it has a model of size $2^{O(\|\varphi\|)}$.*

Proof. We may assume without loss of generality that Σ features no proposition letters. Let φ be as given by (\mathcal{AF} -nmf). By Lemma 5.4.10, φ has a certificate \mathbf{c} of cardinality at most $2^{O(\|\varphi\|)}$; moreover the set of 2-types S occurring anywhere in \mathbf{c} is $2^{O(\|\varphi\|)}$. Let $J = \{0, 1, 2\}$, let T be the index set occurring in φ , let H be a set of cardinality $343 = 7^3$, and let $g: H^2 \rightarrow H$ be a function satisfying the conditions of Lemma 5.4.4 with $\ell = 2$. Defining $A = \mathbf{c} \times S \times J \times T \times H$, we see that $|A|$ is $2^{O(\|\varphi\|)}$, as required by the lemma. We write any element $a \in A$ as (σ, ζ, j, t, h) . We shall construct a model $\mathfrak{A} \models \varphi$ of primitive height 3 over this

domain, proceeding layer by layer. In the sequel, bear in mind that a pair or triple of elements is primitive if and only if those elements are distinct.

Primitive Height 1.

We set the 1-type of any $a = (\sigma, \zeta, j, t, h)$ to be $\mathbf{tp}^{\mathfrak{A}}[a] := \mathbf{tp}(\sigma)$. Clearly, all these determinations can be made independently, since Σ features no proposition letters. At this point, we have a structure of primitive height 1.

Primitive Height 2.

Now consider any $a = (\sigma, \zeta, j, t, h) \in A$ and any $\eta \in \sigma$. By (G \exists), there exists $\sigma_\eta \in \mathfrak{c}$ such that $\eta^{-1} \in \sigma_\eta$. For each $t' \in T$ and $h' \in H$ set $\mathbf{tp}^{\mathfrak{A}}[aa_{t',h'}] := \eta$, where $a_{t',h'}$ denotes the element $(\sigma_\eta, \eta, j+31, t', h')$, where $+3$ is taken to be addition modulo 3. The index η ensures that the $a_{t',h'}$ are chosen to be distinct for distinct $\eta \in \sigma$. Moreover, the index $j+31$ ensures that this process can be carried out for every $a \in A$ without danger of clashes. Finally, suppose $a = (\sigma, \zeta, j, t, h)$ and $a' = (\sigma', \zeta', j', t', h')$ are distinct elements of A for which $\mathbf{tp}^{\mathfrak{A}}[aa']$ has not yet been defined. By (G \forall), there exists $\eta \in \sigma$ such that $\eta^{-1} \in \sigma'$, and we set $\mathbf{tp}^{\mathfrak{A}}[aa'] := \eta$. At the end of this process, all 1- and 2-types have been defined, and we thus have a structure of primitive height 2. From the foregoing construction, if $a = (\sigma, \zeta, j, t, h) \in A$ and $\eta \in \sigma$, then there exists a star σ' such that $\mathbf{tp}^{\mathfrak{A}}[ab] = \eta$ for each $b \in A$ of the form $(\sigma', \eta, j+31, t', h')$ (where $t' \in T, h' \in H$); moreover, for all $a = (\sigma, \zeta, j, t, h)$ and $b = (\sigma', \zeta', j', t', h')$ with $\mathbf{tp}^{\mathfrak{A}}[ab] = \eta$, we are guaranteed that $\eta \in \sigma$ and $\eta^{-1} \in \sigma'$. We remark that, in particular, $\mathbf{str}^{\mathfrak{A}}[a] = \sigma$. It follows from L \exists_1 that, for every $a \in A$ and every $t \in T$, there exists $b \in A$ such that $\mathfrak{A} \models \gamma_t[aab]$. Another way of saying this is that, for every pair of elements $a_1 a_2$ whose primitive length is 1 (i.e. $a_1 = a_2$), \mathfrak{A} provides a witness for the formula $\exists x_3 \gamma_t$. Likewise, it follows from L \forall_1 that, for every triple \bar{a} whose primitive height is either 1 or 2, $\mathfrak{A} \models \beta[\bar{a}]$. Indeed, if $\bar{a} = \bar{b}^f$ where $|\bar{b}| \leq 2$, we have $\mathfrak{A} \models \beta^f[\bar{b}]$, whence $\mathfrak{A} \models \beta[\bar{a}]$ by Lemma 5.2.3.

Primitive Height 3.

We now increment the primitive height of \mathfrak{A} to 3 by setting the adjacent 3-types of all primitive triples in \mathfrak{A} . Fix any pair of distinct elements $a = (\sigma, \tau, h, i, j)$ and $a' = (\sigma', \tau', h', i', j')$. Let us write $\zeta := \mathbf{tp}^{\mathfrak{A}}[aa']$, so that, by construction of

\mathfrak{A} in the previous stage, $\zeta \in \sigma$ and $\zeta^{-1} \in \sigma'$. By $(L\exists_2)$, there exists some $\eta \in \sigma'$ such that the \mathbf{af}^3 -formula $\psi := \zeta \wedge \eta^+ \wedge \gamma_t \wedge \hat{\beta}$ is consistent; let ξ_t be an adjacent 3-type entailing this formula. By the construction of the previous stage again, we can find an element $b_t := (\sigma'', \eta, j' +_3 1, t, g(h, h')) \in A$ such that $\mathbf{tp}^{\mathfrak{A}}[a'b_t] = \eta$. We shall set $\mathbf{atp}^{\mathfrak{A}}[aa'b_t] := \xi_t$ for all $t \in T$. From the index t , the elements b_t are distinct, and so these assignments do not clash with each other. Since ξ_t entails $\zeta \wedge \eta^+$, they do not clash with the 2-types assigned so far. Since ξ_t entails γ_t , the pair aa' now has a witness in respect of the formula $\exists x_3 \gamma_t$. From property (i) of g secured by Lemma 5.4.4, the triple $aa'b_t$ is primitive; hence the only primitive triples whose adjacent types are thereby defined are $aa'b_t$ and $b_t a' a$. But since ξ_t entails $\hat{\beta}$, neither of these triples violates β . Now repeat this construction for all pairs of distinct elements $a = (\sigma, \tau, j, t, h)$ and $a' = (\sigma', \tau', j', t', h')$. We claim that no tuple \bar{c} is assigned to the extensions of any predicates twice in this process. Since \bar{c} must have some primitive generators $a_1 a_2 a_3$ and $a_3 a_2 a_1$, the only possibility for double assignment of \bar{c} is if a_3 is chosen as some witness for the pair $a_1 a_2$, and a_1 is chosen as some witness for the pair $a_3 a_2$. Remembering that a_1, a_2 and a_3 are actually quintuples, let their final components be, respectively, h, h', h'' . By the choice of witnesses, $h'' = g(h, h')$ and $h = g(h'', h')$. But this contradicts property (ii) of g secured by Lemma 5.4.4, thus establishing the claim that no primitive triple is assigned to extensions of predicates twice. At this point, for every pair of elements $a_1 a_2$ (of primitive length either 1 or 2) and every $t \in T$, \mathfrak{A} provides a witness for the formula $\exists x_3 \gamma_t$. Moreover, no adjacent 3-type so-far assigned violates β . To complete the extension of \mathfrak{A} to primitive height 3, it remains only to assign adjacent types to all remaining primitive triples without violating β . Suppose, then $aa'a''$ are distinct elements whose adjacent type in \mathfrak{A} has not yet been defined. Let $\zeta = \mathbf{tp}^{\mathfrak{A}}[a_1 a_2]$ and $\eta = \mathbf{tp}^{\mathfrak{A}}[a_2 a_3]$. By the previous stage, $\zeta \wedge \eta^+ \wedge \hat{\beta}$ is consistent, so let ξ be an adjacent 3-type entailing this formula, and set $\mathbf{atp}^{\mathfrak{A}}[a_1 a_2 a_3] := \xi$. Observe that we are also thereby assigning the adjacent 3-type of $\mathbf{atp}^{\mathfrak{A}}[a_3 a_2 a_1]$, but are assigning no other adjacent 3-types. Since ξ entails $\zeta \wedge \eta^+$, this assignment does not clash with the assignments of the previous step. Since ξ entails $\hat{\beta}$, no newly assigned triple violates β . This completes the construction of the model \mathfrak{A} . \square

5.5 Adding Equality

In this section, we establish decidability of satisfiability for each of the fragments \mathcal{AF}^ℓ , for $\ell \geq 2$, using the same strategy as employed in Section 5.4 for the equality-free sub-fragments. The difference is that, when equality is present, Lemma 3.2.2 becomes invalid. This lemma was used in the proof of Lemma 5.4.5, which constructed a model of the equality-free normal-form $\mathcal{af}^{\ell+1}$ -formula φ from a model of the equality-free \mathcal{af}^ℓ -formula ψ by means of duplicating elements. With duplication we could easily avoid clashes when selecting elements to serve as witnesses for the existential requirements of φ . When equality is present, such duplication is no longer available, thus necessitating additional observations about the structure of models. Table 5.5.1 provides a guide to important notions in this section. (Do note that they are, as of yet, undefined).

Suppose $\mathfrak{A} \models \varphi$, where φ is a normal-form $\mathcal{AF}^{\ell+1}$ -formula ($\ell \geq 2$) over some signature Σ . We fix φ and Σ for the remainder of this section, writing the former as in (\mathcal{AF} -nmf), again repeated here for convenience:

$$\forall \mathbf{x}_{\ell+1} \beta \wedge \bigwedge_{t \in T} \forall \mathbf{x}_\ell \exists x_{\ell+1} \gamma_t.$$

We employ the letters ℓ , T , β and γ_i with these denotations, and we assume without loss of generality that T is non-empty. We proceed to define an expansion \mathfrak{A}^+ of \mathfrak{A} , and simultaneously, a normal-form \mathcal{AF}^ℓ -formula ψ over the expanded signature, such that $\mathfrak{A}^+ \models \psi$; we later show that any layered model of ψ (having height ℓ) has a Σ -reduct that can be elevated to a model of φ .

Construction of ψ from φ in the presence of =

$G_{\bar{a}}$	directed graph $(A, D_{\bar{a}} \cup E_{\bar{a}})$ of witnesses in \mathfrak{A} around \bar{a}
σ	a star; a function mapping witness $(\ell + 1)$ -types to colours C
$\text{atp}(\sigma)$	the underlying ℓ -type of σ
STR	the set of all stars $\sigma : \text{ATP}_{\ell+1}^\Sigma \leftrightarrow C$
$\text{tl}(\xi)$	the ℓ -type η s.t. $\xi \models \eta^+$
$s_\sigma(\mathbf{x}_\ell)$	atom implying \mathbf{x}_ℓ realises σ
$q_{\sigma,c}(\mathbf{x}_{\ell-1})$	atom implying there is some x s.t. $x\mathbf{x}_{\ell-1}$ realise σ and colour c
$r_{\sigma,c,\xi}(\mathbf{x}_\ell)$	atom implying x_ℓ is the ξ -witness for $x\mathbf{x}_{\ell-1}$ realising σ and c

Table 5.5.1: Quick reference guide for Sec 5.5.

We take ψ to have the form

$$\text{acl}(\varphi) \wedge \psi_0 \wedge \cdots \wedge \psi_5, \quad (\psi)$$

where $\text{acl}(\varphi)$ is the adjacent closure of φ (featured in Lemma 5.2.5), and ψ_0, \dots, ψ_5 are \mathcal{AF}^ℓ -formulas over an expanded signature, defined below. We proceed to consider the conjuncts ψ_0, \dots, ψ_5 in turn.

The conjunct ψ_0

The initial step in this process is rather elaborate, and has no analogue in Section 5.4. For any element $a \in A$ and any tuple $\bar{a} \in A^{\ell-1}$, let $B_{a\bar{a}}$ be a minimal set such that, for each $t \in T$, there is some $b \in B_{a\bar{a}}$ with $\mathfrak{A} \models \gamma_t[a\bar{a}b]$. Since $\mathfrak{A} \models \varphi$, such a set exists, and moreover $|B_{a\bar{a}}| \leq |T|$. We call the elements of $B_{a\bar{a}}$ the *witnesses with respect to $a\bar{a}$* . By assumption of minimality of $B_{a\bar{a}}$, there are no two distinct elements $b, b' \in B_{a\bar{a}}$ such that $\text{atp}_{\ell+1}^{\mathfrak{A}}[a\bar{a}b] = \text{atp}_{\ell+1}^{\mathfrak{A}}[a\bar{a}b']$. Thus, given an adjacent $(\ell+1)$ -type ξ such that $\text{atp}_{\ell+1}^{\mathfrak{A}}[a\bar{a}b] = \xi$ for some $b \in B_{a\bar{a}}$, we will call this b the ξ -*witness with respect to $a\bar{a}$* . (Of course, this notion depends on our particular choice of the set $B_{a\bar{a}}$).

Having chosen the witnesses with respect to the various ℓ -tuples over A , consider now any $(\ell-1)$ -tuple \bar{a} over A . Defining the sets of ordered pairs over A

$$\begin{aligned} D_{\bar{a}} &:= \{(a, b) \mid a \neq b \text{ and } b \in B_{a\bar{a}}\} \\ E_{\bar{a}} &:= \{(a, a') \mid a \neq a' \text{ and there is some } b \in B_{a\bar{a}} \text{ s.t. } a' \in B_{b\bar{a}}\} \end{aligned}$$

we let $G_{\bar{a}}$ be the directed graph with vertices A and edges $D_{\bar{a}} \cup E_{\bar{a}}$. Thus, in this directed graph, there is an edge from a to every witness b with respect to $a\bar{a}$ (except a itself), and an edge from any element a to any other element a' if there exists a witness b with respect to $a\bar{a}$ such that a' is a witness with respect to $b\bar{a}$. Since no ℓ -tuple has more than $|T|$ witnesses, the out-degree of any vertex in $G_{\bar{a}}$ is at most $|T|^2 + |T|$. Recall that a k -*colouring* of a directed graph $G_{\bar{a}} = (V, E)$ is a function $f: V \rightarrow [0, k-1]$ satisfying $f(u) \neq f(v)$ for all $(u, v) \in E$. It is well-known that any directed graph with maximum out-degree d has a $(2d+1)$ -colouring (see e.g. [Pratt-Hartmann, 2023, p. 612]). Thus, we may colour the directed graph $G_{\bar{a}}$ with colours from a set C of cardinality $2(|T|^2 + |T|) + 1$. For every $(\ell-1)$ -tuple \bar{a} , then, let some such colouring of $G_{\bar{a}}$ be fixed.

The importance of this colouring will become clearer as the proof unfolds. For the present, however, we note the following: (i) if b is one of the witnesses with respect to $a\bar{a}$, and is distinct from a , then a and b are differently coloured in $G_{\bar{a}}$; (ii) if, in addition, a' is one of the witnesses with respect to $b\tilde{a}$, and is distinct from a , then a and a' are differently coloured in $G_{\bar{a}}$. It does not follow that the various witnesses with respect to $a\bar{a}$ will be differently coloured from each other in $G_{\bar{a}}$. Notice also that, for distinct $(\ell-1)$ -tuples \bar{a} and \bar{a}' , an element a might be coloured differently in the graphs $G_{\bar{a}}$ and $G_{\bar{a}'}$; this is true even if $\bar{a}' = \tilde{\bar{a}}$.

Now, for each $c \in C$ let us identify c as an ℓ -ary predicate, which we interpret in our expansion \mathfrak{A}^+ of \mathfrak{A} . Specifically, for all $a \in A$, $\bar{a} \in A^{\ell-1}$ and $c \in C$, we declare that $\mathfrak{A}^+ \models c[a\bar{a}]$ just in case a is assigned the colour c in the colouring of $G_{\bar{a}}$, and we write $\text{col}^{\mathfrak{A}^+}[a\bar{a}]$ to denote c . Note that, defining

$$\forall \mathbf{x}_\ell \left(\bigvee_{c \in C} c(\mathbf{x}_\ell) \wedge \bigwedge_{c, c' \in C}^{c \neq c'} \left(\neg c(\mathbf{x}_\ell) \vee \neg c'(\mathbf{x}_\ell) \right) \right), \quad (\psi_0)$$

we have $\mathfrak{A}^+ \models \psi_0$. Conversely, if \mathfrak{B} is any model of ψ_0 , every ℓ -tuple over B is assigned a unique colour from C in the obvious way.

The conjunct ψ_1

Here we can simply repeat material from Section 5.4. For each s satisfying $1 \leq 2s+1 \leq \ell-1$ we introduce a fresh $(2s+1)$ -ary predicate d_{2s+1} , and declare that a $(2s+1)$ -tuple over A satisfies d_{2s+1} in the expansion \mathfrak{A}^+ just in case it is a palindrome. In addition, we define the \mathcal{AF}^ℓ -formula

$$\bigwedge_{2 < 2s+1 \leq \ell} \forall \mathbf{x}_{s+1} d_{2s+1}(x_1 \cdots x_s x_{s+1} x_s \cdots x_1). \quad (\psi_1)$$

Thus, $\mathfrak{A}^+ \models \psi_1$, and, conversely, if \mathfrak{B} is any structure such that $\mathfrak{B} \models \psi_1$, and $\bar{c} \in B^{2s+1}$ is a palindrome ($1 \leq 2s+1 \leq \ell$), then $\mathfrak{B} \models d_{2s+1}[\bar{c}]$. As before, we write $\delta_D := \bigwedge \{d_{j-i+1}(x_i \cdots x_j) \mid \langle i, j \rangle \in D\}$ where D is any set of pairs of integers $\langle i, j \rangle$ for $1 \leq i \leq j \leq \ell$ with $j-i+1$ odd.

The conjunct ψ_2

The treatment of this conjunct is more elaborate than in Section 5.4, and involves the colouring of ℓ -tuples encountered above. However, the essential function of

securing witnesses and imposing universal constraints is the same. We remind the reader that, if \bar{a} is a tuple, then \tilde{a} denotes its reversal. Observe that each ℓ -tuple $a\bar{a}$ over A gives rise to a partial function $\sigma: \text{ATP}_{\ell+1}^{\Sigma} \hookrightarrow C$ mapping adjacent $(\ell+1)$ -types over Σ to colours, namely,

$$\sigma(\xi) := \begin{cases} \text{col}^{\mathfrak{A}^+}[b\tilde{a}] & \text{if } b \text{ is the (unique) } \xi\text{-witness in } \mathfrak{A} \text{ with respect to } a\bar{a}, \\ \text{undefined} & \text{if } a\bar{a} \text{ does not have a } \xi\text{-witness in } \mathfrak{A}. \end{cases}$$

Thus, for every adjacent $(\ell+1)$ -type ξ for which $a\bar{a}$ has a ξ -witness b in \mathfrak{A} , the value $\sigma(\xi)$ tells us the colour of the ‘reversed’ ℓ -tuple $b\tilde{a}$ in \mathfrak{A}^+ . We denote the domain of σ (i.e. the set of $\xi \in \text{ATP}_{\ell+1}^{\Sigma}$ for which $\sigma(\xi)$ is defined) by $\text{dom}(\sigma)$. Since T is non-empty, the ℓ -tuple $a\bar{a}$ has at least one witness, whence $\text{dom}(\sigma)$ is also non-empty. We call σ the *star of $a\bar{a}$* in \mathfrak{A}^+ , and denote it $\text{str}^{\mathfrak{A}^+}[a\bar{a}]$. Always remember that the elements of $\text{ATP}_{\ell+1}^{\Sigma}$ are adjacent types over the *original* signature Σ of \mathfrak{A} , not over any expanded signature.

A simple check with reference to the formula φ given above verifies that σ satisfies the following conditions:

1. there exists a (unique) adjacent ℓ -type ζ over Σ such that, for all $\xi \in \text{dom}(\sigma)$, we have $\xi \models \zeta$.
2. for every $\xi \in \text{dom}(\sigma)$ there exists $t \in T$ such that $\xi \models \gamma_t$;
3. for every $t \in T$ there is exactly one $\xi \in \text{dom}(\sigma)$ such that $\xi \models \gamma_t$; and
4. for every $\xi \in \text{dom}(\sigma)$, $\xi \models \widehat{\beta}$ (remember that $\widehat{\beta} := \beta \wedge \beta^{-1}$).

Condition 1 is verified by setting $\zeta := \text{atp}^{\mathfrak{A}^+}[a\bar{a}]$; it is obvious that this is the unique adjacent ℓ -type with the required properties; we call ζ the *underlying adjacent type* of σ and denote it $\text{atp}(\sigma)$. The remaining conditions are immediate from the properties of witnesses. Accordingly, we call any partial function $\sigma: \text{ATP}_{\ell+1}^{\Sigma} \hookrightarrow C$ satisfying conditions 1–4 above a *star*, and we denote the set of all such stars as **STR**. Thus, for any $(\ell+1)$ -tuple over A , we have $\text{str}^{\mathfrak{A}^+}[a\bar{a}] \in \text{STR}$. We remark that the notion of a star depends on the formula φ (and on the parameters ℓ , C , β , T , γ_t and Σ associated with φ). Since, however, φ may be considered fixed for the present, we suppress these parameters to avoid notational clutter.

Now we are ready to fix some additional predicates in the expansion \mathfrak{A}^+ . Since we have settled the interpretations of the predicates C in \mathfrak{A}^+ , the adjacent star-type $\text{str}^{\mathfrak{A}^+}[a\bar{a}] = \sigma$ of any ℓ -tuple $a\bar{a}$ is unaffected by these additional predicates:

in particular, the domain of σ consists of adjacent $(\ell+1)$ -types ξ over the *original signature* Σ ; and the values $\sigma(\xi)$ are determined by the interpretations of the predicates C in \mathfrak{A}^+ . With this in mind, for every $\zeta \in \text{ATP}_\ell^\Sigma$, we introduce the $(\ell-1)$ -ary predicate \mathbf{p}_ζ familiar from Section 5.4, declaring $\mathfrak{A}^+ \models \mathbf{p}_\zeta[\bar{a}]$ just in case, for some $a \in A$, $\text{atp}^\mathfrak{A}[a\bar{a}] = \zeta$. Thus, $\mathfrak{A}^+ \models \psi_{2,0}$, where $\psi_{2,0}$ is

$$\bigwedge_{\zeta \in \text{ATP}_\ell^\Sigma} \forall \mathbf{x}_\ell (\zeta \rightarrow \mathbf{p}_\zeta(x_2 \cdots x_\ell)). \quad (\psi_{2,0})$$

In addition, for every $\sigma \in \text{STR}$, we introduce a new ℓ -ary predicate \mathbf{s}_σ , and declare $\mathfrak{A}^+ \models \mathbf{s}_\sigma[a\bar{a}]$ if and only if $\text{str}^{\mathfrak{A}^+}[a\bar{a}] = \sigma$, for any ℓ -tuple $a\bar{a}$ over A . It is then easy to verify that $\mathfrak{A}^+ \models \psi_{2,1} \wedge \psi_{2,2}$, where the formulas are defined respectively

$$\forall \mathbf{x}_\ell \bigvee_{\sigma \in \text{STR}} \mathbf{s}_\sigma(\mathbf{x}_\ell) \wedge \forall \mathbf{x}_\ell \bigwedge_{\sigma, \sigma' \in \text{STR}}^{\sigma \neq \sigma'} (\neg \mathbf{s}_\sigma(\mathbf{x}_\ell) \vee \neg \mathbf{s}_{\sigma'}(\mathbf{x}_\ell)), \quad (\psi_{2,1})$$

$$\forall \mathbf{x}_\ell \bigwedge_{\sigma \in \text{STR}} (\mathbf{s}_\sigma(\mathbf{x}_\ell) \rightarrow \text{atp}(\sigma)). \quad (\psi_{2,2})$$

Indeed, the first of these formulas states that any ℓ -tuple satisfies exactly one of the predicates \mathbf{s}_σ , and the second, that its adjacent type in \mathfrak{A} is given by $\text{atp}(\sigma)$. (Recall in this connection that $\text{atp}(\sigma)$ is an ℓ -type, and thus a formula with variables \mathbf{x}_ℓ). Further, for every $\sigma \in \text{STR}$ and every $c \in C$, we introduce a new $(\ell-1)$ -ary predicate $\mathbf{q}_{\sigma,c}$, and declare $\mathfrak{A}^+ \models \mathbf{q}_{\sigma,c}[\bar{a}]$ just in case there is some $a \in A$ such that $\text{str}^{\mathfrak{A}^+}[a\bar{a}] = \sigma$ and $\text{col}^{\mathfrak{A}^+}[a\bar{a}] = c$. Thus, $\mathbf{q}_{\sigma,c}$ identifies tails of ℓ -tuples whose star in \mathfrak{A}^+ is σ and whose colour in \mathfrak{A}^+ is c . It is thus immediate that $\mathfrak{A}^+ \models \psi_{2,3}$, where $\psi_{2,3}$ is

$$\forall \mathbf{x}_\ell \bigwedge_{\sigma \in \text{STR}} \bigwedge_{c \in C} ((\mathbf{s}_\sigma(\mathbf{x}_\ell) \wedge \mathbf{c}(\mathbf{x}_\ell)) \rightarrow \mathbf{q}_{\sigma,c}(x_2 \cdots x_\ell)). \quad (\psi_{2,3})$$

Still proceeding with the construction of ψ_2 , for every $\sigma \in \text{STR}$, every $c \in C$, and every $\xi \in \text{dom}(\sigma)$, we introduce a new ℓ -ary predicate $\mathbf{r}_{\sigma,c,\xi}$, and fix its interpretation in \mathfrak{A}^+ as follows. Take any $(\ell-1)$ -tuple \bar{a} over A . If, on the one hand, $\mathfrak{A}^+ \models \mathbf{q}_{\sigma,c}[\bar{a}]$, then select some $a \in A$ such that $\text{str}^{\mathfrak{A}^+}[a\bar{a}] = \sigma$ and $\text{col}^{\mathfrak{A}^+}[a\bar{a}] = c$. By the interpretation of $\mathbf{q}_{\sigma,c}$ in \mathfrak{A}^+ , this is possible. Now, for each $\xi \in \text{dom}(\sigma)$, let b be the ξ -witness for $a\bar{a}$, and set $\mathfrak{A}^+ \models \mathbf{r}_{\sigma,c,\xi}[a\bar{b}]$. If, on the other hand, $\mathfrak{A}^+ \not\models \mathbf{q}_{\sigma,c}[\bar{a}]$, then do nothing in respect of the tuple \bar{a} . By carrying out this procedure for every $(\ell-1)$ -tuple \bar{a} over A , we thus fix the extensions of

the predicates $\mathbf{r}_{\sigma,c,\xi}$ in \mathfrak{A}^+ . Informally, it helps to read the atom $\mathbf{r}_{\sigma,c,\xi}(\mathbf{x}_\ell)$ as “ x_ℓ wants to be the ξ -witness with respect to the tuple $x'\mathbf{x}_{\ell-1}$ for a particular x' such that $x'\mathbf{x}_{\ell-1}$ has star-type σ and colour c ”. Since $\xi \in \text{dom}(\sigma)$ implies that a ξ -witness exists, we have $\mathfrak{A}^+ \models \psi_{2,4}$, where $\psi_{2,4}$ is

$$\forall \mathbf{x}_{\ell-1} \exists x_\ell \bigwedge_{\sigma \in \text{STR}} \bigwedge_{c \in C} \bigwedge_{\xi \in \text{dom}(\sigma)} (\mathbf{q}_{\sigma,c}(\mathbf{x}_{\ell-1}) \rightarrow \mathbf{r}_{\sigma,c,\xi}(\mathbf{x}_\ell)). \quad (\psi_{2,4})$$

Returning to our ℓ -tuple $a\bar{a}$ with $\text{str}^{\mathfrak{A}^+}[a\bar{a}] = \sigma$, $\text{col}^{\mathfrak{A}^+}[a\bar{a}] = c$ and ξ -witness b . Let $\text{tl}(\xi)$ denote the unique adjacent ℓ -type η over Σ such that $\xi \models \eta^+$. Thus, $\text{atp}^{\mathfrak{A}^+}[\bar{a}b] = \text{tl}(\xi)$. Moreover, from the fact that $\text{str}^{\mathfrak{A}^+}[a\bar{a}] = \sigma$ and b is the ξ -witness with respect to $a\bar{a}$, we have $\text{col}^{\mathfrak{A}^+}[b\bar{a}] = \sigma(\xi)$. Thus, $\mathfrak{A}^+ \models \psi_{2,5}$, where $\psi_{2,5}$ is

$$\forall \mathbf{x}_\ell \bigwedge_{\sigma \in \text{STR}} \bigwedge_{c \in C} \bigwedge_{\xi \in \text{dom}(\sigma)} \left(\mathbf{r}_{\sigma,c,\xi}(\mathbf{x}_\ell) \rightarrow (\text{tl}(\xi) \wedge (\sigma(\xi))(x_\ell \cdots x_1)) \right). \quad (\psi_{2,5})$$

Recall in this regard that $\text{tl}(\xi)$ is an adjacent ℓ -type (hence a formula with free variables \mathbf{x}_ℓ), and $\sigma(\xi) \in C$ is an ℓ -ary predicate of Σ .

More can be said considering the ℓ -tuple $\bar{a}b$ and its reversal, $b\bar{a}$. This latter tuple has some star in \mathfrak{A}^+ , say $\text{str}^{\mathfrak{A}^+}[b\bar{a}] = \sigma'$. Consider, then, any adjacent $(\ell+1)$ -type $\xi' \in \text{dom}(\sigma')$. Thus, $b\bar{a}$ has a ξ' -witness in \mathfrak{A} , say b' . If $b' = a$, then $\text{atp}^{\mathfrak{A}^+}[a\bar{a}b]$ and $\text{atp}^{\mathfrak{A}^+}[b\bar{a}b']$ are mutually inverse, and we have $\xi' = \xi^{-1}$. If, on the other hand, $b' \neq a$, then (a, b') is an edge of the directed graph $G_{\bar{a}}$. Thus, $\text{col}^{\mathfrak{A}^+}[a\bar{a}] \neq \text{col}^{\mathfrak{A}^+}[b'\bar{a}]$, by construction of \mathfrak{A}^+ . Thus, $\mathfrak{A}^+ \models \psi_{2,6}$, where $\psi_{2,6}$ is

$$\forall \mathbf{x}_\ell \bigwedge_{\sigma \in \text{STR}} \bigwedge_{c \in C} \bigwedge_{\xi \in \text{dom}(\sigma)} \left(\mathbf{r}_{\sigma,c,\xi}(\mathbf{x}_\ell) \rightarrow \bigvee \{ \mathbf{s}_{\sigma'}(x_\ell \cdots x_1) \mid \right. \\ \left. \sigma' \in \text{STR} \text{ and, for all } \xi' \in \text{dom}(\sigma'), \xi' \neq \xi^{-1} \Rightarrow \sigma'(\xi') \neq c \} \right). \quad (\psi_{2,6})$$

That is to say: if b is the ξ -witness with respect to $a\bar{a}$ and b' the ξ' -witness with respect to $b\bar{a}$, then either ξ and ξ' are mutually inverse types in \mathfrak{A} or else the tuples $a\bar{a}$ and $b'\bar{a}$ are differently coloured in \mathfrak{A}^+ .

We complete the construction of ψ_2 with a simple constraint concerning the predicates $\mathbf{r}_{\sigma,c,\xi}$. Consider again any $(\ell-1)$ -tuple \bar{a} over A , and suppose that, for some $\sigma \in \text{STR}$, some $c \in C$ and some pair of distinct $(\ell+1)$ -types $\xi, \xi' \in \text{dom}(\sigma)$, there exist elements b and b' such that $\mathfrak{A}^+ \models \mathbf{r}_{\sigma,c,\xi}[\bar{a}b]$ and $\mathfrak{A}^+ \models \mathbf{r}_{\sigma,c,\xi'}[\bar{a}b']$. From the construction of \mathfrak{A}^+ , b and b' must have been chosen as ξ - and ξ' -witnesses,

respectively, with respect to a tuple $a\bar{a}$ for some particular element a . It follows that b and b' must be distinct. That is, $\mathfrak{A}^+ \models \psi_{2,7}$, where $\psi_{2,7}$ is

$$\forall \mathbf{x}_\ell \bigwedge_{\sigma \in \text{STR}} \bigwedge_{c \in C} \bigwedge_{\substack{\xi \neq \xi' \\ \xi, \xi' \in \text{dom}(\sigma)}} (\neg r_{\sigma, c, \xi}(\mathbf{x}_\ell) \vee \neg r_{\sigma, c, \xi'}(\mathbf{x}_\ell)). \quad (\psi_{2,7})$$

Conversely, if \mathfrak{B} is any structure making $\psi_{2,7}$ true, and \bar{a} an $(\ell-1)$ -tuple over B , then we cannot have the same element b such that $\mathfrak{B} \models r_{\sigma, c, \xi}[\bar{a}b]$ and $\mathfrak{B} \models r_{\sigma, c, \xi'}[\bar{a}b]$ for different ξ and ξ' in the domain of σ . Setting ψ_2 for

$$\psi_{2,0} \wedge \cdots \wedge \psi_{2,7} \quad (\psi_2)$$

we have thus established that $\mathfrak{A}^+ \models \psi_2$.

The conjuncts ψ_3 and ψ_4

Here, we can again recapitulate the ideas of Section 5.4, though in a slightly different guise. As before, we take \mathbf{D}_k^o to denote the set of all pairs $\langle i, j \rangle$ for $1 \leq i \leq j \leq k$ such that $j-i+1$ is odd. Fix some subset $D \subseteq \mathbf{D}_{\ell-1}^o$, and suppose \bar{a} is an $(\ell-1)$ -tuple over A such that $\mathfrak{A}^+ \models \delta_D[\bar{a}]$. By construction of \mathfrak{A}^+ , the defect set of \bar{a} includes D , whence, for any elements $a, b \in A$, the defect set of $a\bar{a}b$ certainly includes D^+ . It follows that, if $\sigma \in \text{STR}$, $c \in C$ and $\xi \in \text{dom}(\sigma)$, but with $\partial\xi$ not D^+ -compatible, then there cannot exist $b \in A$ such that $\mathfrak{A} \models r_{\sigma, c, \xi}[\bar{a}b]$. For otherwise, by construction of \mathfrak{A}^+ , there exists a such that $\text{str}^{\mathfrak{A}^+}[a\bar{a}] = \sigma$, and b is the ξ -witness with respect to $a\bar{a}$, whence $\text{atp}^{\mathfrak{A}}[a\bar{a}b] = \xi$, contradicting the first statement of Lemma 5.4.1. We have thus proved $\mathfrak{A}^+ \models \psi_3$, where ψ_3 is

$$\bigwedge_{D \subseteq \mathbf{D}_{\ell-1}^o} \forall \mathbf{x}_\ell (\delta_D \rightarrow \bigwedge \{ \neg r_{\sigma, c, \xi}(\mathbf{x}_\ell) \mid \sigma \in \text{STR}, c \in C, \xi \in \text{dom}(\sigma), \partial\xi \text{ not } D^+\text{-compatible} \}). \quad (\psi_3)$$

For ψ_4 , we recapitulate the formula from Section 5.4, namely

$$\bigwedge_{\zeta \in \text{ATP}_\ell^\Sigma} \bigwedge_{D \subseteq \mathbf{D}_{\ell-1}^o} \forall \mathbf{x}_\ell \left((\delta_D \wedge \mathbf{p}_\zeta(\mathbf{x}_{\ell-1})) \rightarrow \bigvee \{ \eta \in \text{ATP}_\ell^\Sigma \mid (\zeta \wedge \widehat{\beta} \wedge \eta^+) \text{ is } D^+\text{-consistent} \} \right). \quad (\psi_4)$$

And by the same reasoning as in Section 5.4, we have $\mathfrak{A}^+ \models \psi_4$.

The conjunct ψ_5

Our final conjunct again has no analogue in Section 5.4, and concerns extra conditions we need to impose on certain *non-primitive* $(\ell+1)$ -tuples in models of ψ . We need to ensure that, when reconstructing a model of φ from such structures, we do not attempt to assign these tuples incompatible adjacent types. We require the following simple lemma regarding words. Recall that a walk $f : [1, m] \rightarrow [1, k]$ is terminal if $f(m) = k$. We say that an ℓ -tuple \bar{b} is *terminal* if it can be written $\bar{b} = \bar{d}^f$ for some k -tuple \bar{d} with $k < \ell$, and some terminal walk $f : [1, \ell] \rightarrow [1, k]$. A *leg*, for any walk $f : [1, \ell] \rightarrow [1, k]$, is a maximal interval $[i, j] \subseteq [1, \ell]$ such that $f(h+1) - f(h)$ is constant for all h ($i \leq h < j$).

Lemma 5.5.1. *Let \bar{b} be an ℓ -tuple and b an element. Then at least one of the following holds: (i) $\bar{b}b$ is primitive; (ii) b is the last element of \bar{b} ; (iii) $\bar{b}b$ has a suffix which is an odd-length, non-trivial palindrome; (iv) \bar{b} is terminal.*

Proof. Suppose that $\bar{b}b$ is not primitive. If \bar{b} has an immediately repeated letter, then it is certainly terminal: indeed $\bar{b} = \bar{a}c\bar{c}\bar{d}$ is generated from $\bar{a}c\bar{d}$ via a terminal function $f : [1, \ell] \rightarrow [1, \ell-1]$ which pauses for one step on the letter c . Hence we may assume that there are no immediately repeated letters in \bar{b} . Furthermore, if b is not the last element of \bar{b} , then the whole of $\bar{b}b$ has no immediately repeated letters. Since $\bar{b}b$ is not primitive, we have $\bar{b}b = \bar{c}^f$ for some k -tuple \bar{c} and some walk $f \in [1, \ell] \rightarrow [1, k]$ with at least two legs (i.e. maximal strictly ascending or descending intervals). If the final leg of f is shortest, then some suffix of $\bar{b}b$ is a non-trivial palindrome, and this palindrome must have odd length, since there are no immediately repeated letters. Otherwise $\bar{b}b$ has either of the forms $c\bar{d}\bar{d}\tilde{c}\bar{c}b$ or $\bar{a}c\bar{d}\bar{d}\tilde{c}\bar{c}\bar{d}\bar{e}b$, depending on whether the shortest leg is initial or internal. In the former case, $\bar{b} = (d\tilde{d}\bar{c})^g$ for some final walk g ; in the latter, $\bar{b} = (\bar{a}c\bar{d}\bar{d}\bar{e})^h$ for some final walk h . In both cases, \bar{b} is terminal. \square

We remark that, in the case where $\bar{b}b$ has a suffix which is an odd-length non-trivial palindrome, that palindrome may be the whole of $\bar{b}b$. The cases of Lemma 5.5.1 correspond to conditions on the adjacent type of the $(\ell+1)$ -tuple in question. Say that an adjacent $(\ell+1)$ -type ξ is *palindromic* if, for any $\mathcal{AF}^{\ell+1}$ -atom α , $\xi \models \alpha$ implies $\xi \models \alpha(x_{\ell+1} \cdots x_1)$. Evidently, if an $(\ell+1)$ -tuple is a palindrome, then its adjacent type in any structure is palindromic. Say that ξ is *blunt* if, for any $\mathcal{AF}^{\ell+1}$ -atom α , we have that $\xi \models \alpha$ implies $\xi \models \alpha(\mathbf{x}_\ell x_\ell)$. If an $(\ell+1)$ -tuple has the same last two elements, then its adjacent type in any structure is blunt.

Say that ξ is s -hooked (for s satisfying $2 < 2s + 1 < \ell + 1$) if, for any $\mathcal{AF}^{\ell+1}$ -atom α we have that $\xi \models \alpha$ implies $\xi \models \alpha(x_1 \cdots x_{\ell+1-s} x_{\ell-s} \cdots x_{\ell+1-2s})$. If an $(\ell+1)$ -tuple has a proper suffix that is a non-trivial palindrome of length $2s + 1$, then its adjacent type in any structure is s -hooked. Observe the strict inequality $2s + 1 < \ell + 1$ governing the parameter s in this last definition: if ξ is palindromic (and $\ell + 1$ is odd), we do not say that ξ is $(\ell/2)$ -hooked.

Now let \bar{a} be an $(\ell-1)$ -tuple over A and $b \in A$, and suppose $\mathfrak{A}^+ \models \mathbf{r}_{\sigma,c,\xi}[\bar{a}b]$. By the construction of \mathfrak{A}^+ , there exists $a \in A$ such that b is the ξ -witness with respect to $a\bar{a}$, and, moreover, $\text{col}^{\mathfrak{A}^+}[a\bar{a}] = c$. If ξ is palindromic, blunt or s -hooked for some s ($2 < 2s + 1 < \ell + 1$), then the $(\ell+1)$ -tuple $a\bar{a}b$ exhibits certain properties, which we proceed to describe. Consider first the case where $\ell + 1$ is odd and ξ is not palindromic, and suppose in addition that $\mathfrak{A}^+ \models d_{\ell-1}[\bar{a}]$. By the construction of \mathfrak{A}^+ again, \bar{a} is a palindrome, i.e. $\bar{a} = \tilde{a}$. It follows that $a \neq b$, since otherwise, $a\bar{a}b$ is a palindrome, contradicting the assumption that $\xi = \text{atp}^{\mathfrak{A}^+}[a\bar{a}b]$ is not palindromic. And since b is a witness with respect to $a\bar{a}$ with $a \neq b$, the ordered pair (a, b) is an edge of the directed graph $G_{\bar{a}} = G_{\tilde{a}}$, so that by the construction of \mathfrak{A}^+ , we have $c = \text{col}^{\mathfrak{A}^+}[a\bar{a}] \neq \text{col}^{\mathfrak{A}^+}[b\bar{a}] = \text{col}^{\mathfrak{A}^+}[b\tilde{a}]$. Thus we have shown that $\mathfrak{A}^+ \models \psi_{5,1}$, where, for $(\ell+1)$ odd, $\psi_{5,1}$ is

$$\bigwedge_{\sigma \in \text{STR}} \bigwedge_{c \in C} \bigwedge_{\substack{\xi \text{ not palindromic} \\ \xi \in \text{dom}(\sigma)}} \forall \mathbf{x}_\ell \left(\mathbf{r}_{\sigma,c,\xi}(\mathbf{x}_\ell) \wedge \mathbf{d}_{\ell-1}(\mathbf{x}_{\ell-1}) \rightarrow \neg \mathbf{c}(x_\ell \cdots x_1) \right), \quad (\psi_{5,1})$$

and for $(\ell+1)$ even, $\psi_{5,1} := \top$. We remark that even-length palindromes have immediately repeated letters in the middle, which obviates—as we shall see later—the need for an analogue of the odd-length case.

Second, consider the case where ξ is not blunt. Here, we do not need to add any conjuncts to ψ , since the consistency of ξ requires that it contains the inequality literal $x_\ell \neq x_{\ell+1}$. Clearly, $\text{atp}^{\mathfrak{A}^+}[\bar{a}b]$ contains the inequality literal $x_{\ell-1} \neq x_\ell$, which is all the information we shall require.

Third, consider the case where, for any s in the range $2 < 2s + 1 < \ell + 1$, ξ is not s -hooked. Since $\xi = \text{atp}^{\mathfrak{A}^+}[a\bar{a}b]$, it follows that $a\bar{a}b$ has no proper suffix of length $2s + 1$ that is a palindrome, and therefore, $\bar{a}b$ has no suffix of the same length that is a palindrome. Thus we have shown that $\mathfrak{A}^+ \models \psi_{5,2}$, where

$$\bigwedge_{\sigma \in \text{STR}} \bigwedge_{c \in C} \bigwedge_{s=1}^{s \leq (\ell-1)/2} \bigwedge_{\substack{\xi \text{ not } s\text{-hooked} \\ \xi \in \text{dom}(\sigma)}} \forall \mathbf{x}_\ell \left(\mathbf{r}_{\sigma,c,\xi}(\mathbf{x}_\ell) \rightarrow \neg \mathbf{d}_{2s+1}(x_{\ell-2s} \cdots x_\ell) \right). \quad (\psi_{5,2})$$

Writing ψ_5 for the following conjunction of formulas

$$\psi_{5,1} \wedge \psi_{5,2} \tag{\psi_5}$$

we have shown that $\mathfrak{A}^+ \models \psi_5$. This completes the definition of the formula ψ .

Summarizing the above discussion, and recalling that, by Lemma 5.2.5, $\varphi \models \text{acl}(\varphi)$, we have:

Lemma 5.5.2. *Suppose $\mathfrak{A} \models \varphi$. Then we can expand \mathfrak{A} to a model $\mathfrak{A}^+ \models \psi$.*

Having defined ψ and established Lemma 5.5.2, we establish a converse in the form of the following lemma.

Lemma 5.5.3. *Suppose $\mathfrak{A} \models \psi$. Then we can construct a model $\mathfrak{A}' \models \varphi$ over the same domain.*

Proof. Since $\psi \in \mathcal{AF}^\ell$, by Lemma 5.3.7 we may take \mathfrak{A} to be a structure of primitive height ℓ . That is, we assume that adjacent types of primitive $(\ell+1)$ -tuples are left undefined. Setting \mathfrak{A}^- to be the Σ -reduct of \mathfrak{A} , we proceed by expanding \mathfrak{A}^- into a model of φ by setting the interpretations of the predicates in Σ with respect to the primitive $(\ell+1)$ -tuples over A .

Let a be any element of A and \bar{a} any $(\ell-1)$ -tuple over A . By ψ_0 , $a\bar{a}$ has a unique colour, say $c = \text{col}^{\mathfrak{A}}[a\bar{a}]$, and by $\psi_{2,1}$, there is a unique $\sigma \in \text{STR}$ such that $\mathfrak{A} \models \mathbf{s}_\sigma[a\bar{a}]$. Let $\zeta = \text{atp}(\sigma)$ be the underlying ℓ -type of σ , and let us fix, for the moment, any $\xi \in \text{dom}(\sigma)$. Thus, $\zeta := \xi|_{[1,\ell]}$. By $\psi_{2,2}$, we have that $\text{atp}^{\mathfrak{A}^-}[a\bar{a}] = \zeta$. Moreover, by $\psi_{2,3}$, $\mathfrak{A} \models \mathbf{q}_{\sigma,c}[\bar{a}]$, and hence, by the fact that $\mathfrak{A} \models \psi_{2,4}$, there exists $b \in A$ such that $\mathfrak{A} \models \mathbf{r}_{\sigma,c,\xi}[\bar{a}b]$. Let $\eta = \text{tl}(\xi)$; it follows from $\psi_{2,5}$ that $\text{atp}^{\mathfrak{A}^-}[\bar{a}b] = \eta$ and $\text{col}^{\mathfrak{A}}[\bar{a}b] = \sigma(\xi)$. The intention is that we should set the interpretations of the predicates in the structure \mathfrak{A}' in such a way that the $(\ell+1)$ -tuple $a\bar{a}b$ is assigned the adjacent type ξ . The various other conjuncts of ψ will ensure that this can be done consistently.

Suppose on the one hand that the $(\ell+1)$ -tuple $a\bar{a}b$ is primitive. Thus, no prefix or suffix of $a\bar{a}b$ is a non-trivial palindrome, and $a\bar{a}b$ contains no immediately repeated letters. Hence, any defect $\langle i, j \rangle$ of $a\bar{a}b$ satisfies $2 \leq i \leq j \leq \ell$, with $j-i+1$ odd. Letting D now be the set of defects $\langle i, j \rangle$ of the $(\ell-1)$ -tuple \bar{a} (so that $D \subseteq \mathbf{D}_{\ell-1}^\circ$), we see that the set of defects of $a\bar{a}b$ is given by $D^+ = \{\langle i+1, j+1 \rangle \mid \langle i, j \rangle \in D\}$. By ψ_1 , we have that $\mathfrak{A} \models \delta_D[\bar{a}]$; and by ψ_3 , bearing in mind that $\mathfrak{A} \models \mathbf{r}_{\sigma,c,\xi}[\bar{a}b]$, we have that $\partial\xi$ is D^+ -compatible. Lemma 5.4.1 thus ensures

that it is meaningful to set $\text{itp}_{\ell+1}^{\mathfrak{A}'}(a\bar{a}b) := \partial\xi$. Moreover, since $\text{atp}^{\mathfrak{A}^-}[a\bar{a}] = \zeta$, $\text{atp}^{\mathfrak{A}^-}[\bar{a}b] = \eta$, and $\xi = \zeta \cup \eta^+ \cup \partial\xi$, we have $\text{atp}_{\ell+1}^{\mathfrak{A}'}[a\bar{a}b] = \xi$. We remark that the adjacent type of the primitive $(\ell+1)$ -tuple $b\bar{a}a$ is also set in this process, but no other primitive $(\ell+1)$ -tuples have their adjacent types set. Observe that, since $\xi \models \widehat{\beta}$, the newly-set adjacent types of $a\bar{a}b$ and $b\bar{a}a$ will not violate β .

Suppose, on the other hand, that $a\bar{a}b$ is not primitive. We show that, in this case, \mathfrak{A} already provides a ξ -witnesses with respect to $a\bar{a}$, so that there is nothing to do. Here, we make use of the fact that $\mathfrak{A} \models \text{acl}(\varphi) \wedge \psi_5$. By Lemma 5.5.1, either b is the last element of $a\bar{a}$, or $a\bar{a}b$ has a suffix that is a non-trivial odd-length palindrome, or $a\bar{a}$ is terminal. The case where $a\bar{a}$ is terminal is easily dealt with. There exists a k -tuple \bar{d} and a final walk $f \in [1, \ell] \rightarrow [1, k]$ such that $a\bar{a} = \bar{d}^f$ for some k ($2 \leq k < \ell$). Defining, as before, $f^+ = f \cup \langle k+1, \ell+1 \rangle$, by $\text{acl}(\varphi)$ there exists, for each $t \in T$, some $b' \in A$ such that $\mathfrak{A} \models \gamma_t[(\bar{d}b')^{f^+}]$, i.e. $\mathfrak{A} \models \gamma_t[a\bar{a}b']$. In effect, we are discarding our originally chosen element b , since the required witness is already present.

Consider next the case where b is the last element of \bar{a} . We have already argued that $\text{atp}^{\mathfrak{A}^-}[\bar{a}b] = \eta = \text{tl}(\xi)$, whence $\text{tl}(\xi) \models x_{\ell-1} = x_\ell$, whence $\xi \models x_\ell = x_{\ell+1}$. Thus, ξ is blunt, by consistency of ξ . But in that case, writing $\bar{a} = a_1 \cdots a_{\ell-1}$, we have $\mathfrak{A}^- \models \zeta[aa_1 \cdots a_{\ell-1}]$ implies $\mathfrak{A}^- \models \xi[aa_1 \cdots a_{\ell-1}a_{\ell-1}]$, that is, $\mathfrak{A}^- \models \xi[a\bar{a}b]$. Thus, our chosen element b is already a ξ -witness with respect to $a\bar{a}$, without our having to do anything.

Consider finally the case where $a\bar{a}b$ is not terminal and has a suffix that is an odd-length, non-trivial palindrome. Suppose, on the one hand, that the suffix in question is the *whole* of $a\bar{a}b$ —that is, $\ell+1$ is odd, and $a\bar{a}b$ has the form $aa_1 \cdots a_{\ell/2}a_{\ell/2-1} \cdots a_1a$. Thus, $a = b$ and $\bar{a} = \tilde{a}$. Hence, $a\bar{a} = b\tilde{a}$, whence $\text{col}^{\mathfrak{A}^-}[b\tilde{a}] = \text{col}^{\mathfrak{A}^-}[a\tilde{a}] = c$. It follows by ψ_1 that $\mathfrak{A} \models \mathbf{d}_{\ell-1}[\bar{a}]$, and, by $\psi_{5,1}$, that ξ is palindromic. Thus, $\text{atp}^{\mathfrak{A}^-}[aa_1 \cdots a_{\ell/2}] = \zeta|_{[1, \ell/2+1]} = \xi|_{[1, \ell/2+1]}$, whence $\text{atp}^{\mathfrak{A}^-}[aa_1 \cdots a_{\ell/2}a_{\ell/2-1} \cdots a_1a] = \xi$. Again, our chosen element $b = a$ is already a ξ -witness with respect to $a\bar{a}$.

On the other hand, suppose that $a\bar{a}b$ has a *proper* suffix that is a non-trivial, odd-length palindrome. Thus, $a\bar{a}b$ is of the form $aa_1 \cdots a_{\ell-s}a_{\ell-(s+1)} \cdots a_{\ell-2s}$ (for some s with $2 < 2s+1 \leq \ell$). But by ψ_1 , $\mathfrak{A} \models \mathbf{d}_s[a_{\ell-2s} \cdots a_{\ell-s}a_{\ell-(s+1)} \cdots a_{\ell-2s}]$. Hence, by $\psi_{5,2}$, we have that ξ is s -hooked. As in the previous case, it follows that $b = a_{(\ell-2s)}$ is already a ξ -witness with respect to $a\bar{a}$. This completes the process of finding a ξ -witness with respect to $a\bar{a}$.

Keeping a and \bar{a} fixed for the moment, carry out the above assignments for all $\xi \in \text{dom}(\sigma)$, choosing, for each such ξ , a ξ -witness b_ξ with respect to $a\bar{a}$. It follows from $\psi_{2,7}$ that the various elements b_ξ must be distinct, so no clashes can arise. This completes the process of finding all the required witnesses with respect to $a\bar{a}$. From the properties of the elements of **STR**, we see that, however \mathfrak{A}' is completed, the ℓ -tuple $a\bar{a}$ will satisfy $\exists x_{\ell+1} \gamma_i$ in \mathfrak{A}' for every $t \in T$. Moreover, each of the newly-fixed primitive $(\ell+1)$ -tuples (and their reversals) satisfies β .

Now vary a and \bar{a} freely. We must show that, for a' and \bar{a}' with $a\bar{a} \neq a'\bar{a}'$, the chosen witnesses b and b' can never lead to a clash. Suppose then that some m -tuple \bar{d} is assigned (or not) to the extension of an m -ary predicate p when defining the adjacent $(\ell+1)$ -types of both of the primitive tuples $a\bar{a}b$ and $a'\bar{a}'b'$. Then \bar{d} is generated by both $a\bar{a}b$ and $a'\bar{a}'b'$, and there exist adjacent $(\ell+1)$ -types ξ and ξ' such that b is selected as the ξ -witness with respect to $a\bar{a}$ and b' is selected as the ξ' -witness with respect to $a'\bar{a}'$. Since, by assumption, $a\bar{a} \neq a'\bar{a}'$, Theorem 5.2.1 implies $a\bar{a}b$ is the reversal of $a'\bar{a}'b'$, so that $a' = b$, $b' = a$ and $\bar{a}' = \bar{a}$; henceforth, we shall write $b\bar{a}$ in preference to $a'\bar{a}'$. It suffices to show that $\xi' = \xi^{-1}$, as then there can be no clash in the assignment of \bar{d} to the extension of the predicate p .

Suppose, for contradiction that $\xi' \neq \xi^{-1}$. By $\psi_{2,1}$, let σ and σ' be the unique elements of **STR** such that $\mathfrak{A} \models \mathbf{s}_\sigma[a\bar{a}]$ and $\mathfrak{A} \models \mathbf{s}_{\sigma'}[b\bar{a}]$; and by ψ_0 , let c and c' be the unique elements of C such that $\mathfrak{A} \models \mathbf{c}[a\bar{a}]$ and $\mathfrak{A} \models \mathbf{c}'[b\bar{a}]$. Since b was chosen as a ξ -witness for $a\bar{a}$, we have $\mathfrak{A} \models \mathbf{r}_{\sigma,c,\xi}[\bar{a}b]$, and similarly $\mathfrak{A} \models \mathbf{r}_{\sigma',c',\xi'}[\bar{a}a]$. Recalling that $\mathfrak{A} \models \mathbf{r}_{\sigma,c,\xi}[\bar{a}b]$ and $\xi' \neq \xi^{-1}$, we have by $\psi_{2,6}$ that $\sigma'(\xi') \neq c$. Moreover, since $\mathfrak{A} \models \mathbf{r}_{\sigma',c',\xi'}[\bar{a}a]$, by $\psi_{2,5}$ we have $\mathfrak{A} \models (\sigma'(\xi'))[a\bar{a}]$, contradicting $\mathfrak{A} \models \mathbf{c}[a\bar{a}]$. Thus, $\xi' = \xi^{-1}$ as required.

This completes the process of finding witnesses for all ℓ -tuples from A , in the course of which we have partially defined \mathfrak{A}' , such that, however it is completed, $\mathfrak{A}' \models \forall \mathbf{x}_\ell \exists x_{\ell+1} \gamma_t$ for every $t \in T$. Moreover, each of the newly-fixed primitive $(\ell+1)$ -tuples (and their reversals) satisfies β .

All that remains is to complete the construction of \mathfrak{A}' by assigning adjacent types to all primitive $(\ell+1)$ -tuples whose adjacent types have not yet been fixed, without violating the condition $\forall \mathbf{x}_{\ell+1} \beta$. Suppose, then $\bar{c} = a\bar{a}b$ is such a primitive $(\ell+1)$ -tuple. By primitiveness, any defect $\langle i, j \rangle$ of \bar{c} is such that $2 < i \leq j \leq \ell$ and $j-i+1$ is odd. Thus, writing D for the set of defects of \bar{a} , we have that $D \subseteq \mathbf{D}_{\ell-1}^\circ$, and D^+ is exactly the defect set of \bar{c} . By ψ_1 we have $\mathfrak{A} \models \delta_D[\bar{a}]$ and, writing $\zeta = \text{atp}^{\mathfrak{A}^-}[a\bar{a}]$, we have that $\mathfrak{A} \models \mathbf{p}_\zeta[\bar{a}]$ by $\psi_{2,0}$. Writing $\eta =$

$\text{atp}^{\mathfrak{A}^-}[\bar{a}b]$, by ψ_4 there exists an adjacent $(\ell+1)$ -type ξ such that $\xi \models \zeta \wedge \widehat{\beta} \wedge \eta^+$ and $\partial\xi$ is D^+ -compatible. By the second statement of Lemma 5.4.1, then, we can consistently assign $\text{itp}^{\mathfrak{A}'}[\bar{c}] = \partial\xi$. Since $\zeta = \text{atp}^{\mathfrak{A}^-}[a\bar{a}]$, $\eta = \text{atp}^{\mathfrak{A}^-}[\bar{a}b]$, and $\xi = \zeta \cup \eta^+ \cup \partial\xi$, we have $\text{atp}^{\mathfrak{A}'}[\bar{c}] = \xi$. Additionally, since $\xi \models \widehat{\beta}$, both \bar{c} and \tilde{c} satisfy the universal requirements β in \mathfrak{A}' . Repeated applications of this procedure result in all primitive $(\ell+1)$ -tuples having their adjacent types defined in such a way that $\mathfrak{A}' \models \forall \mathbf{x}_{\ell+1} \beta$. Hence $\mathfrak{A}' \models \varphi$. \square

Taken together, Lemmas 5.5.2 and 5.5.3 reduce the satisfiability problem for $\mathcal{AF}^{\ell+1}$ to that for \mathcal{AF}^ℓ ($\ell \geq 2$), though with exponential blow-up. We thus obtain the decidability of satisfiability for the whole of \mathcal{AF} . More precisely:

Theorem 5.5.4. *If φ is a satisfiable \mathcal{AF}^ℓ -formula, with $\ell \geq 2$, then φ is satisfied in a structure of size at most $t(\ell-1, \|\varphi\|^{O(1)})$. Hence the satisfiability problem for \mathcal{AF}^ℓ is in $(\ell-1)$ -NEXPTIME, and the adjacent fragment is TOWER-complete.*

Proof. Suppose that φ_ℓ a \mathcal{AF}^ℓ -sentence. By Lemma 5.1.1, we may assume that it is in normal-form. Above we provided a procedure that, when given a sentence such as φ_ℓ , produces in exponential time an $\mathcal{AF}^{\ell-1}$ -sentence $\varphi_{\ell-1}$. By Lemma 5.5.2, every model of φ_ℓ can be expanded into a model of $\varphi_{\ell-1}$. Lemma 5.5.3 shows the converse; i.e. each model \mathfrak{A} of $\varphi_{\ell-1}$ gives rise to a model \mathfrak{A}' over the same domain such that $\mathfrak{A}' \models \varphi_\ell$. It is easy to verify that $\|\psi_1\|$ is a constant whilst $\|\psi_0\|$ and $\|\text{acl}(\varphi_\ell)\|$ are bounded by $O(\|\varphi_\ell\|)$. Referencing Lemma 5.4.6, we again have that atoms that are not part of $\{\alpha^f \mid \alpha \text{ is an atom of } \varphi_\ell\}$ for some walk $f : [1, \ell] \rightarrow [1, \ell]$ do not play any role in entailing quantifier-free subformulas of φ_ℓ and thus can safely be discarded. Writing Σ_ℓ for the signature of φ_ℓ , we have that $|\text{ATP}_\ell^{\Sigma_\ell}|$ and $|\text{ATP}_{\ell-1}^{\Sigma_\ell}|$ are then bounded by $2^{O(\|\varphi_\ell\|)}$. We now claim that $|\text{STR}|$ is bounded by $2^{\|\varphi_\ell\|^{O(1)}}$. To see this, recall that each star is a function from a subset $\text{ATP}_\ell^{\Sigma_\ell}$ of size at most $|T|$ to the set C , which is of cardinality at most $2(|T|^2 + |T|) + 1$. Clearly, the total number of stars does not exceed $\sum_{k=1}^{|T|} (2^{O(\|\varphi_\ell\|)} \cdot |C|)^k$ is bounded by $2^{\|\varphi_\ell\|^{O(1)}}$ as required. With the above it is easy to verify that the sizes of ψ_2 , ψ_3 , ψ_4 and ψ_5 are bounded by a polynomial function on the cardinalities of $\text{ATP}_{\ell-1}^{\Sigma_\ell}$ and STR ; i.e. $2^{\|\varphi_\ell\|^{O(1)}}$.

Noting that $\varphi_{\ell-1}$ is a normal-form $\mathcal{AF}^{\ell-1}$ -sentence, we may apply the reduction above and obtain a sequence of sentences $\varphi_{\ell-2} \dots \varphi_2$ of formulas each

exponentially larger than the last. Recalling the function $\mathfrak{t} : \mathbb{N} \times \mathbb{N} \rightarrow \mathbb{N}$:

$$\begin{aligned}\mathfrak{t}(0, n) &:= n, \\ \mathfrak{t}(m+1, n) &:= 2^{\mathfrak{t}(m, n)}.\end{aligned}$$

We see, by an easy induction on the size of the formulas, that $\mathfrak{t}(\ell-2, \|\varphi_\ell\|^{O(1)})$ is a bound on $\|\varphi_2\|$. Since $\varphi_2 \in \mathcal{FO}^2$, Theorem 2.2.2 allows us to check the satisfiability status of φ_2 in time bounded by $2^{\|\varphi_2\|^{O(1)}}$. By Lemmas 5.5.2 and 5.5.3, this also checks the satisfiability prospects of φ_ℓ . Thus, the total time required is $\mathfrak{t}(\ell-1, \|\varphi_\ell\|^{O(1)})$. Additionally, by Lemma 2.2.1, if φ_2 is satisfiable, then it has a model of size $2^{\|\varphi_2\|^{O(1)}} = \mathfrak{t}(\ell-1, \|\varphi_\ell\|^{O(1)})$. The same then holds for φ_ℓ as it is satisfiable over the same domains as φ_2 .

We see that when the number of variables is not fixed, the value ℓ is bounded by $\|\varphi\|$. Thus, satisfiability for sentences φ of the *full* adjacent fragment can be checked in time $\mathfrak{t}(\|\varphi\|, \|\varphi\|^{O(1)})$. Thus, $\text{Sat}(\mathcal{AF})$ resides in the class TOWER. TOWER-hardness is inherited from $\text{Sat}(fl)$ [Pratt-Hartmann et al., 2019]. \square

It may not have escaped the reader's attention that the equality predicate barely features in the proof of Theorem 5.5.4. This is because the reduction from the satisfiability problem for $\mathcal{AF}^{\ell+1}$ to that for \mathcal{AF}^ℓ effected in Lemmas 5.5.2 and 5.5.3 does not interfere with predicates of arity less than 3, whilst also retaining the domain of the original structure. The special status of the equality predicate is thus, in a sense, pushed back into the two-variable fragment.

This observation suggests a generalization similar to that found in Corollary 4.4.1. Recent decades have witnessed numerous attempts to investigate the decidability of the satisfiability problem for \mathcal{FO}^2 augmented with distinguished symbols having fixed semantics. There include but are not limited to linear orders [Kieroński and Tendera, 2009, Schwentick and Zeume, 2012], trees [Charatonik et al., 2014, Bednarczyk et al., 2017], equivalences [Kieroński and Tendera, 2009, Kieroński and Otto, 2012]. If such an extension of \mathcal{FO}^2 is decidable for (finite) satisfiability, then, by utilising the reduction outlined in this section, the decidability status carries over to the adjacent fragment; it being understood that the relations in question are of arity at most 2. On the other hand, if the extension of \mathcal{FO}^2 is undecidable for (finite) satisfiability, then, by Theorem 5.3.3, this undecidability result immediately transfers to \mathcal{AF} as well.

Corollary 5.5.5. *Let Σ^* be some set of symbols of arity at most 2 with a fixed semantic interpretation. Then, the (finite) satisfiability problem for \mathcal{AF} is decidable over signatures incorporating the predefined symbols Σ^* if and only if the (finite) satisfiability problem for \mathcal{FO}^2 is decidable over signatures incorporating the predefined symbols Σ^* .*

5.6 The Guarded Subfragment

We next shift our attention to the *guarded adjacent fragment*, denoted \mathcal{GA} , defined as the intersection of \mathcal{AF} with the guarded fragment, \mathcal{GF} [Andréka et al., 1998, Sec. 4.1]. In \mathcal{GF} , quantification is relativized by atoms, i.e. all quantification takes the form $\forall \bar{x}(\alpha \rightarrow \psi)$ and $\exists \bar{x}(\alpha \wedge \psi)$, where α (a *guard*) is an atom featuring all the variables in \bar{x} and all the free variables of ψ . The satisfiability problem for \mathcal{GF} is 2EXPTIME-complete [Grädel, 1999, Thm. 4.4]. We show that the satisfiability problem for \mathcal{GA} remains 2EXPTIME-complete, and thus is as hard as for full \mathcal{GF} . This contrasts with the EXPTIME-completeness of the k -variable guarded fragment \mathcal{GF}^k [Grädel, 1999, Cor. 4.6] and of the guarded forward fragment [Bednarczyk, 2021, Thm. 4]. Our proof follows the same reduction strategy as the 2EXPTIME-hardness proof for \mathcal{GF} by E. Grädel [Grädel, 1999, Thm. 4.4]. Because we are working in the guarded *adjacent* fragment, however, Grädel's reduction is not directly available, and some new techniques are required. To aid readability, we employ the variable names w, x, y, z, \dots in adjacent formulas, rather than the official $x_1, x_2, x_3, x_4, \dots$. The reader may easily check that all formulas in question are indeed in \mathcal{AF} modulo variable renaming.

Generating Words.

We start with a combinatorial exercise concerning the generation of words by the recurrent application of certain walks. Let $m \in \mathbb{N}$ and consider the walks $\lambda_1, \lambda_2, \lambda_3: [1, m+2] \rightarrow [1, m+1]$ defined by the following courses of values

$(\lambda_i(1) \ \cdots \ \lambda_i(m+2))$:

$$\begin{aligned}\lambda_1 &:= (1 \ 2 \ 2 \ 3 \ 4 \ \dots \ m+1), \\ \lambda_2 &:= (1 \ 2 \ 1 \ 2 \ 3 \ \dots \ m), \\ \lambda_3 &:= (1 \ 2 \ 3 \ 3 \ 4 \ \dots \ m+1).\end{aligned}$$

Intuitively, we picture the functions as being walks on a word of length $(m+1)$, in each case starting at the left-most position, and proceeding generally rightwards: λ_1 pauses at the second time step before continuing; λ_2 returns to the beginning after the second time step, but then resumes its rightward journey (without quite reaching the end); and λ_3 pauses at the third time step before continuing.

We show that repeated application of these functions to the bit-string 011^m yields the whole of the (exponentially large) language $01\{0,1\}^m$.

Lemma 5.6.1. *Let $W_0 \subseteq \{0,1\}^*$ contain 011^m and $W_i := W_{i-1} \cup \{\bar{w}^{\lambda_1}, \bar{w}^{\lambda_2}, \bar{w}^{\lambda_3} \mid \bar{w} \in W_{i-1}\}$. Setting $W := \bigcup_{i \geq 0} W_i$, we have $01\{0,1\}^m \subseteq W$.*

Proof. We establish by induction on $i \in [0, m]$ that, for any word $\bar{c} \in \{0,1\}^i$, the word $01\bar{c}1^{m-i}$ is in W ; the case $i = m$ then yields the statement of the lemma. The base case, $i = 0$ follows from the assumption that W_0 contains 011^m . Suppose now $i > 0$. We show that, for any $x \in \{0,1\}$, the word $\bar{u} = 01x\bar{c}1^{m-i-1}$ is in W . Letting c_1 be the first character of \bar{c} , and writing $\bar{c} = c_1\bar{d}$, it follows by inductive hypothesis that the words $\bar{v} = 01\bar{c}1^{m-i-1}1$ and $\bar{w} = 01\bar{d}1^{m-i-1}11$ are in W . Consider cases: (i) if $x = 1$ then $\bar{u} = \bar{v}^{\lambda_1}$, (ii) if both $x = 0$ and $c_1 = 0$ then $\bar{u} = \bar{v}^{\lambda_3}$, and otherwise, (iii) $x = 0$, $c_1 = 1$ and $\bar{u} = \bar{w}^{\lambda_2}$. Thus $\bar{u} \in W$. \square

We now apply Lemma 5.6.1 to \mathcal{GA} . For $m \geq 0$, let \mathbf{G}_m be an $(m+2)$ -ary predicate. We proceed to write, for any binary predicate \mathbf{p} , a \mathcal{GA} -sentence $\overrightarrow{\text{allSeq}}_m(\mathbf{p})$ ensuring that, if \mathbf{p} is satisfied by a pair of objects, say ab , then \mathbf{G}_m is satisfied by any $(m+2)$ -tuple $ab\bar{w}$ for $\bar{w} \in \{a,b\}^m$. By Lemma 5.6.1, it suffices so take $\overrightarrow{\text{allSeq}}_m(\mathbf{p})$ to be

$$\forall xy \left(\mathbf{p}(xy) \rightarrow \mathbf{G}_m(xy \underbrace{y \cdots y}_m) \right) \wedge \bigwedge_{i \in [1,3]} \forall \mathbf{u}_{m+2} \left(\mathbf{G}_m(\mathbf{u}_{m+2}) \rightarrow \mathbf{G}_m(\mathbf{u}_{m+2}^{\lambda_i}) \right).$$

Thus, if $\mathfrak{A} \models \overrightarrow{\text{allSeq}}_m(\mathbf{p})$, and $\mathfrak{A} \models \mathbf{p}[ab]$, we may freely quantify over words of length $m+2$ over the alphabet $\{a,b\}$ as long as they have the prefix ab , since \mathbf{G}_m

can always be used as a guard.

We shall require a ‘mirrored’ version of the above device, this time involving a *pair* of 2-element alphabets and a *pair* of words of length m over these alphabets. For $m \geq 0$, let F_m be a $(2m+4)$ -ary predicate. We proceed to write, for any quaternary predicate \mathbf{r} , a \mathcal{GA} -sentence $\overleftrightarrow{\text{allSeq}}_m(\mathbf{r})$ ensuring that, if \mathbf{r} is satisfied by a quartet of objects $abcd$, then, for any m -tuple \bar{u} over the alphabet $\{a, b\}$, and any m -tuple \bar{v} over the alphabet $\{c, d\}$, the predicate F_m is satisfied by $\bar{u}abcd\bar{v}$. Let λ_0 denote the identity function on $[1, m+2]$, and take $\lambda_1, \lambda_2, \lambda_3$ as defined above. In addition, for any word \bar{w} , we write \bar{w}^{-1} for the reversal of \bar{w} . (Thus, $\bar{w}^{-1} = \widetilde{\bar{w}}$; but the new notation is more readable in what follows.) By two applications of Lemma 5.6.1, it suffices so take $\overleftrightarrow{\text{allSeq}}_m(\mathbf{r})$ to be:

$$\forall yxzt \left(\mathbf{r}(yxzt) \rightarrow F_m(\underbrace{y \dots y}_m yxzt \underbrace{t \dots t}_m) \right) \wedge \\ \bigwedge_{i,j \in [0,3]} \forall \mathbf{u}_{m+2}^{-1} \mathbf{v}_{m+2} \left(F_m(\mathbf{u}_{m+2}^{-1} \mathbf{v}_{m+2}) \rightarrow F_m((\mathbf{u}_{m+2}^{\lambda_i})^{-1} \mathbf{v}_{m+2}^{\lambda_j}) \right).$$

Again, if $\mathfrak{A} \models \overleftrightarrow{\text{allSeq}}_m(\mathbf{r})$, and $\mathfrak{A} \models \mathbf{r}[abcd]$, we may freely quantify over words of the language $\{a, b\}^m abcd \{c, d\}^m$ in \mathcal{GA} , since F_m can always be used as a guard. Note that the variables of $\mathbf{u}_{m+2} = u_1 \cdots u_{m+2}$ are quantified above in ‘reverse order’. This ensures, after renaming, the adjacency of the formula $\overleftrightarrow{\text{allSeq}}_m(\mathbf{r})$.

Alternating Turing Machines.

An *alternating Turing machine* [Chandra et al., 1981] is a tuple $\mathcal{M} := \langle Q, \Gamma, q_0, \Delta \rangle$, where Q is a non-empty finite set (the *states* of \mathcal{M}), Γ a non-empty finite alphabet, q_0 an element of Q (the *initial state*), and Δ a set *transitions* δ , defined presently. We imagine \mathcal{M} to operate on an 1-way infinite tape by means of a read-write head as usual, but we take Q to be partitioned into the sets Q_{\exists} (*existential states*), Q_{\forall} (*universal states*) and $\{q_a, q_r\}$ (the *accepting* and *rejecting state*, respectively). Writing Γ' for the alphabet Γ augmented with the blank cell symbol ‘ \sqcup ’, we define a *transition* $\delta \in \Delta$ to be a relation

$$\delta \in Q \times \Gamma' \times Q \times \Gamma' \times \{-1, 0, 1\}.$$

The transition $\delta = \langle q, s, q', s', k' \rangle$ is enabled when the machine is in state q and the read-write head is positioned over a tape square containing the letter s . On

execution of δ , the current tape square is overwritten with s' , the head is moved by k' squares, and the current state is updated to q' . We denote the set of $\delta \in \Delta$ enabled by state q and letter s by $\Delta(q, s)$.

A *configuration* \mathcal{C} of \mathcal{M} is a triple $\langle q, \bar{w}, h \rangle$, where $q \in Q$, \bar{w} is an infinite word over Γ' and h a non-negative integer. We read the triple \mathcal{C} as stating that the machine is in state q , the tape contents are given by \bar{w} , and the head is situated over the h th tape square (counting from 0). If \bar{w}_0 is a finite word over Γ representing the input to the machine, the *initial configuration* is $\langle q_0, \bar{w}_0 \sqcup^*, 0 \rangle$, where \sqcup^* represents an infinite series of blanks. The *successors* of \mathcal{C} are defined in the usual way via transitions in Δ . A *halting* configuration is one which is in state q_a or q_r . We assume that halting configurations have no enabled transitions, and non-halting configurations always have at least one enabled transition.

We shall be interested in the case where every computation of \mathcal{M} (understood as a sequence of enabled transitions starting in the initial configuration) is of finite length, so that there is a function f such that, when \mathcal{M} runs on input w_0 , the read-write head never reaches positions beyond $f(|w_0|) - 1$. In that case, we may as well take a configuration to have the form $\langle q_0, \bar{w}, h \rangle$ where \bar{w} is a (finite) word over Γ' of length $f(|w_0|) - 1$ and h is an integer in the range $[0, f(|w_0|) - 1]$. The notions of acceptance and rejection may then be defined as follows: a halting configuration is *accepting* if it is in state q_a ; an existential configuration is *accepting* if it has an accepting successor; a universal configuration is *accepting* if all its successors are accepting; a configuration which is not accepting is *rejecting*. We take \mathcal{M} to accept \bar{w}_0 if the initial configuration is accepting.

We witness acceptance of an input \bar{w}_0 by \mathcal{M} using an *acceptance tree* \mathcal{T} . This is a finite tree with vertices labelled by (accepting) configurations. The root is labelled by an *initial configuration*; and for any vertex labelled with a particular configuration, its children are labelled with the results of executing enabled transitions in that configuration. Vertices labelled with existential configurations have *at least one* child corresponding to an enabled transition; those labelled with universal configurations have a child corresponding to *every* enabled transition; the leaves of the tree are labelled with accepting configurations.

We are interested in the case where the function f bounding the space required by \mathcal{M} is of the form $f(n) = 2^n$. Thus, \mathcal{M} accesses at most $2^{|\bar{w}_0|}$ tape squares in the course of any computation on input \bar{w}_0 . We now fix such a machine \mathcal{M} , and show how, for a given input \bar{w}_0 , we can manufacture a \mathcal{GA} -formula

$\psi_{\mathcal{M}, \bar{w}_0}$ satisfiable if and only if \mathcal{M} accepts \bar{w}_0 . The computation of $\psi_{\mathcal{M}, \bar{w}_0}$ runs in polynomial time (in fact, in logarithmic space) as a function of $|\bar{w}_0|$. Since there are problems in $\text{ASPACE}(2^n)$ that are complete for AEXPSPACE , we can thereby reduce any problem in AEXPSPACE to the satisfiability problem for \mathcal{GA} . Hence the latter problem is AEXPSPACE -hard. Using the well-known equation $\text{AEXPSPACE} = 2\text{EXPTIME}$, this achieves our goal.

Encoding numbers.

In the sequel, we will consider structures \mathfrak{A} interpreting a unary predicate $\mathbf{0}$. Whenever $\mathfrak{A} \models \alpha[ab]$, where $\alpha(xy)$ is the formula $\neg \mathbf{0}(x) \wedge \mathbf{0}(y)$, we say that a and b *act as zero and unit bits* ($\mathbf{0}$ for “One”), and for any word \bar{u} over $\{a, b\}$, we write $\text{VAL}^{\mathfrak{A}}(\bar{u})$ to denote the integer represented by \bar{u} , considered as a bit-string, with a standing for 0 and b for 1 (most significant bit first). Notice that there may be other elements, say c and d , such that $\mathfrak{A} \models \alpha[cd]$, in which case may write $\text{VAL}^{\mathfrak{A}}(\bar{w})$ for the integer represented by any word \bar{w} over $\{c, d\}$. Clearly, the \mathcal{GA} -formula $\text{EQ}(\mathbf{u}_n, \mathbf{v}_n) := \bigwedge_{i=1}^n \mathbf{0}(u_i) \leftrightarrow \mathbf{0}(v_i)$ satisfies $\mathfrak{A} \models \text{EQ}[\bar{c}, \bar{d}]$ if and only if $\text{VAL}^{\mathfrak{A}}(\bar{c}) = \text{VAL}^{\mathfrak{A}}(\bar{d})$. Other arithmetical properties can be expressed similarly. In particular we may write

$$\text{EQ}(\mathbf{u}_n, \mathbf{v}_n + 1) := \bigwedge_{i=1}^n \left((\mathbf{0}(u_i) \leftrightarrow \mathbf{0}(v_i)) \leftrightarrow \bigvee_{j=1}^{i-1} \mathbf{0}(v_j) \right),$$

with the formulas $\text{EQ}(\mathbf{u}_n, \mathbf{v}_n - 1)$ and $\text{EQ}(\mathbf{u}_n, \mathbf{v}_n + 0)$ being defined as $\text{EQ}(\mathbf{v}_n, \mathbf{u}_n + 1)$ and $\text{EQ}(\mathbf{u}_n, \mathbf{v}_n)$ respectively to axiomatise, for each $k \in \{-1, +1\}$,

$$\mathfrak{A} \models \text{EQ}[\bar{c}, \bar{d} + k] \text{ iff } \text{VAL}^{\mathfrak{A}}(\bar{c}) = \text{VAL}^{\mathfrak{A}}(\bar{d}) + k.$$

Observe that the formula $\text{EQ}(\bar{x}, \bar{y} + 1)$ is adjacent. Indeed, all the predicates appearing in this formula are unary; thus, this formula may appear in adjacent formulas whatever the order of quantification among the variables \bar{x} and \bar{y} .

Encoding configurations.

We proceed to describe a method of encoding, within a certain class of structures, configurations of an Alternating Turing Machine $\mathcal{M} = \langle Q, \Gamma, q_0, \Delta \rangle$ that never accesses more than $f(n) = 2^n$ tape squares on an input of size n . Recall that we

agreed to regard configurations in this case as triples $\langle q, \bar{w}, h \rangle$, where $q \in Q$, \bar{w} is a word over Γ' of length $N = 2^n$, and $0 \leq h < N$. In doing so, we indentify various states $q \in Q$ as (typographically decorated) *binary* predicates \mathbf{q} : a configuration is then represented by an ordered pair of (distinct) elements ab satisfying any of these predicates. If $\mathfrak{A} \models \mathbf{q}[ab]$, we take it that the represented configuration has state q . Since we shall want words over the alphabet $\{a, b\}$ to represent integers, we shall insist that, in this case, ab satisfies the formula α defined in above. This we ensure by writing the \mathcal{GA} -sentence:

$$\bigwedge_{q \in Q} \forall xy (\mathbf{q}(xy) \rightarrow \neg \mathbf{0}(x) \wedge \mathbf{0}(y)). \quad (\psi_\alpha)$$

(Thus, we think of a as a 0 and b as a 1). To represent further aspects of the configuration ab , we identify, for each symbol $s \in \Gamma'$, the (typographically decorated) letter \mathbf{s} as an n -ary predicate, and additionally employ an n -ary predicate \mathbf{H} . Specifically, for any $\bar{w} \in \{a, b\}^n$, we read $\mathfrak{A} \models \mathbf{H}[\bar{w}]$ as “the head of the configuration represented by ab is at position $\text{VAL}^{\mathfrak{A}}(\bar{w})$ ”, and we read $\mathfrak{A} \models \mathbf{s}[\bar{w}]$ as “the tape square $\text{VAL}^{\mathfrak{A}}(\bar{w})$ of the configuration represented by ab contains the symbol s ”. Of course, these interpretations are only meaningful if:

1. there is at most one string in $\{a, b\}^n$ satisfying \mathbf{H} and thus encoding the head position;
2. each bit-string over $\{a, b\}^n$ satisfies at most one predicate \mathbf{s} for $s \in \Gamma'$, thus ensuring that a tape cell contains at most one symbol; and
3. ab satisfies at most one \mathbf{q} for $q \in Q$, thus ensuring that the configuration is in at most one state.

For any state $q \in Q$ we construct a \mathcal{GA} -formula ψ_q :

$$\psi_{q,1} \wedge \psi_{q,2} \wedge \psi_{q,3}, \quad (\psi_q)$$

enforcing these conditions for any configuration whose state is q . The first conjunct, $\psi_{q,1}$, may be given as follows:

$$\overrightarrow{\text{allSeq}}_{2n}(\mathbf{q}) \wedge \forall xy \mathbf{u}_n \mathbf{v}_n \left(\mathbf{G}_{2n}(xy \mathbf{u}_n \mathbf{v}_n) \rightarrow \left((\mathbf{H}(\mathbf{u}_n) \wedge \mathbf{H}(\mathbf{v}_n)) \rightarrow \text{EQ}(\mathbf{u}_n, \mathbf{v}_n) \right) \right). \quad (\psi_{q,1})$$

Note that this sentence is (adjacent and) guarded, with the atom $\mathbf{G}_{2n}(xy\mathbf{u}_n\mathbf{v}_n)$ acting as a guard. Of course, thanks to the conjunct $\overrightarrow{\text{allSeq}_{2n}(\mathbf{q})}$, this guard is, as it were, *semantically inert*, since, assuming that \mathbf{q} is satisfied by ab in some structure, say, \mathfrak{A} , \mathbf{G}_{2n} is satisfied by *every* word of $ab\{a, b\}^{2n}$. It is then easy to see that there is at most one string $\bar{w} \in \{a, b\}^n$ such that $\mathfrak{A} \models \mathbf{H}[\bar{w}]$.

Using similar arguments, it is easy to verify that the second requirement is represented by the following \mathcal{GA} -sentences:

$$\overrightarrow{\text{allSeq}_n(\mathbf{q})} \wedge \forall xy\mathbf{u}_n \left(\mathbf{G}_n(xy\mathbf{u}_n) \rightarrow \bigwedge_{s, s' \in \Gamma'}^{s \neq s'} (\neg \mathbf{s}(\mathbf{u}_n) \vee \neg \mathbf{s}'(\mathbf{u}_n)) \right), \quad (\psi_{q,2})$$

Indeed, the conjunct $\overrightarrow{\text{allSeq}_n(\mathbf{q})}$ guards the sequences $ab\{a, b\}^n$, so long as ab represent a configuration in state q . Then, $\mathfrak{A} \models \psi_{q,2}$ ensures that no bit-string over $\{a, b\}^n$ satisfies more than one \mathbf{s} for $s \in \Gamma$.

The final condition is ensured trivially:

$$\forall xy \left(\mathbf{q}(xy) \rightarrow \bigwedge_{q' \in Q}^{q \neq q'} \neg \mathbf{q}'(xy) \right). \quad (\psi_{q,3})$$

Encoding instances of transitions

For each transition $\delta \in \Delta$ we employ a quaternary predicate \mathbf{E}_δ and read $\mathfrak{A} \models \mathbf{E}_\delta[abcd]$ as “the configuration encoded by ab enables a transition δ thus producing the configuration encoded by cd ”. (Do note the reversal of ab in $\mathbf{E}_\delta[abcd]$). We call $abcd$ a δ -*transition instance* in \mathfrak{A} with ab corresponding to the predecessor configuration and cd to the successor configuration. To make sure that each δ -transition instance is indeed a result of transitioning via $\delta = \langle q, s, q', s', k' \rangle$ we write the following in \mathcal{GA} :

1. the successor configuration is in state q' ;
2. the head of the successor configuration is moved by k' relative to the predecessor configurations head;
3. the h -th tape cell on the successor configuration is occupied by s' , where h is the position of the predecessor’s head;
4. all tape cells that the predecessor’s head does not point to are inherited by the successor.

For any transition $\delta \in \Delta$, we construct a \mathcal{GA} -formula ψ_δ :

$$\psi_{\delta,1} \wedge \cdots \wedge \psi_{\delta,4} \quad (\psi_\delta)$$

enforcing these conditions. We begin axiomatisation in reverse order as $\psi_{\delta,4}$ is the most intricate of the four:

$$\begin{aligned} & \overleftarrow{\text{allSeq}_n(\mathbf{E}_\delta)} \wedge \forall \mathbf{u}_n yxzt \mathbf{v}_n \left(\mathbf{F}_n(\mathbf{u}_n yxzt \mathbf{v}_n) \rightarrow \right. \\ & \quad \left. \left((\mathbf{E}_\delta(yxzt) \wedge \neg \mathbf{H}(\mathbf{u}_n) \wedge \mathbf{EQ}(\mathbf{u}_n, \mathbf{v}_n)) \rightarrow \bigwedge_{s \in \Gamma} (\mathbf{s}(\mathbf{u}_n) \leftrightarrow \mathbf{s}(\mathbf{v}_n)) \right) \right). \quad (\psi_{\delta,4}) \end{aligned}$$

Suppose now that $\mathfrak{A} \models \psi_{\delta,4}$, and moreover, that $\mathfrak{A} \models \mathbf{E}_\delta[\overleftarrow{bacd}]$ for some $bacd \in A^4$, where ab and cd both encode configurations. By $\overleftarrow{\text{allSeq}_n(\mathbf{E}_\delta)}$, we have that $\mathfrak{A} \models \mathbf{F}_n[\bar{u}bacd\bar{v}]$ for all $\bar{u} \in \{a, b\}^n$ and $\bar{v} \in \{c, d\}^n$. By picking any $0 \leq i < 2^n$ that is not the head position of the configuration encoded by ab , and $\bar{u} \in \{a, b\}^n$, $\bar{v} \in \{c, d\}^n$ with $\text{VAL}^{\mathfrak{A}}(\bar{u}) = \text{VAL}^{\mathfrak{A}}(\bar{v}) = i$, we are guaranteed that, for each symbol $s \in \Gamma$, $\mathfrak{A} \models \mathbf{s}[\bar{u}]$ if and only if $\mathfrak{A} \models \mathbf{s}[\bar{v}]$.

Recalling that $\delta = \langle q, s, q', s', k' \rangle$ we have the following satisfy condition 3:

$$\begin{aligned} & \overleftarrow{\text{allSeq}_n(\mathbf{E}_\delta)} \wedge \forall \mathbf{u}_n yxzt \mathbf{v}_n \left(\mathbf{F}_n(\mathbf{u}_n yxzt \mathbf{v}_n) \rightarrow \right. \\ & \quad \left. \left((\mathbf{E}_\delta(yxzt) \wedge \mathbf{H}(\mathbf{u}_n) \wedge \mathbf{EQ}(\mathbf{u}_n, \mathbf{v}_n)) \rightarrow \mathbf{s}'(\mathbf{v}_n) \right) \right). \quad (\psi_{\delta,3}) \end{aligned}$$

Indeed, now taking $\mathfrak{A} \models \psi_{\delta,3}$ and $\mathfrak{A} \models \mathbf{E}_\delta[\overleftarrow{bacd}]$ by $\overleftarrow{\text{allSeq}_n(\mathbf{E}_\delta)}$ we again have that $\mathfrak{A} \models \mathbf{F}_n[\bar{u}bacd\bar{v}]$ for all $\bar{u} \in \{a, b\}^n$ and $\bar{v} \in \{c, d\}^n$. By taking i to be the head position of the configuration encoded by ab , we have $\mathfrak{A} \models \mathbf{s}'[\bar{v}]$ when $\text{VAL}^{\mathfrak{A}}(\bar{v}) = i$.

The second condition is then formalised as follows:

$$\begin{aligned} & \overleftarrow{\text{allSeq}_n(\mathbf{E}_\delta)} \wedge \forall \mathbf{u}_n yxzt \mathbf{v}_n \left(\mathbf{F}_n(\mathbf{u}_n yxzt \mathbf{v}_n) \rightarrow \right. \\ & \quad \left. \left((\mathbf{E}_\delta(yxzt) \wedge \mathbf{H}(\mathbf{u}_n) \wedge \mathbf{EQ}(\mathbf{v}_n, \mathbf{u}_n + k')) \rightarrow \mathbf{H}(\mathbf{v}_n) \right) \right). \quad (\psi_{\delta,2}) \end{aligned}$$

Again, taking $\mathfrak{A} \models \psi_{\delta,2}$ and a δ -transition instance $bacd \in A^4$, suppose that i is the head position of the configuration encoded by ab . Similar reasoning as before allows us to conclude that $\mathfrak{A} \models \mathbf{H}[\bar{v}]$, where $\bar{v} \in \{c, d\}^n$ is such that

$\text{VAL}^{\mathfrak{A}}(\bar{v}) = i + k'$, with k' taken from δ .

The first conjunct is trivial:

$$\forall yxzt \left(\mathbf{E}_\delta(yxzt) \rightarrow \mathbf{q}'(zy) \right). \quad (\psi_{\delta,1})$$

Encoding acceptance trees

The last step in our reduction is to write a formula whose models contain configurations arranged as an acceptance tree witnessing the fact that \mathcal{M} accepts some input $w_0 = s_1 \cdots s_n$. Recall that, the root of this tree is labelled with the initial configuration, namely $\mathcal{C} = \langle q_0, \bar{w}_0 \sqcup^\ell, 0 \rangle$, where $\ell = 2^n - |\bar{w}_0|$. We ensure the existence of such a root configuration with the following \mathcal{GA} -sentence:

$$\begin{aligned} \overrightarrow{\text{allSeq}}_n(\mathbf{q}_0) \wedge \exists xy \left(\mathbf{q}_0(xy) \wedge \mathbf{H}("0") \wedge \bigwedge_{i=0}^{i < |\bar{w}_0|} \mathbf{s}_{i+1}("i") \wedge \right. \\ \left. \forall \mathbf{u}_n \left(\mathbf{G}(xy\mathbf{u}_n) \rightarrow \left(\bigwedge_{i=0}^{i < |\bar{w}_0|} \neg \mathbf{EQ}("i", \mathbf{u}_n) \rightarrow \sqcup(\mathbf{u}_n) \right) \right) \right), \quad (\psi_{\mathcal{C}}) \end{aligned}$$

where “ i ” is the binary encoding of i using x as a zero bit and y as a unit bit. We ensure the existence of successor configurations required by existential configurations as follows. Suppose that q is an existential state, and s a symbol. The following sentence ensures that any configuration in state $q \in Q_\exists$ with the head reading symbol $s \in \Gamma'$ has a child in the acceptance tree:

$$\begin{aligned} \overrightarrow{\text{allSeq}}_n(\mathbf{q}) \wedge \forall \mathbf{u}_n^{-1}yx \left(\mathbf{G}_n(xy\mathbf{u}_n) \rightarrow \right. \\ \left. \left((\mathbf{q}(xy) \wedge \mathbf{H}(\mathbf{u}_n) \wedge \mathbf{s}(\mathbf{u}_n)) \rightarrow \bigvee_{\delta \in \Delta(q,s)} \exists zt \mathbf{E}_\delta(yxzt) \right) \right). \quad (\psi_\exists) \end{aligned}$$

In case q is universal, the disjunction over $\Delta(q, s)$ is replaced by a conjunction resulting in the sentence ψ_\forall . Lastly, taking q_r to be the unique rejecting state of \mathcal{M} , we write the sentence $\psi_r := \neg \exists xy \mathbf{q}_r(xy)$ to ensure that a rejecting configuration is never encoded by any pair of elements in any structure.

For any input string w_0 , let $\psi_{\mathcal{M}, \bar{w}_0}$ be the following:

$$\psi_\alpha \wedge \bigwedge_{q \in Q} \psi_q \wedge \bigwedge_{\delta \in \Delta} \psi_\delta \wedge \psi_{\mathcal{C}} \wedge \psi_\exists \wedge \psi_\forall \wedge \psi_r. \quad (\psi_{\mathcal{M}, \bar{w}_0})$$

We remark that $\psi_{\mathcal{M}, \bar{w}_0}$ does not feature the equality predicate. Bearing in mind that \mathcal{M} is guaranteed to terminate on \bar{w}_0 in a finite number of steps, accessing no more than $2^{|\bar{w}_0|}$ tape squares, we see that any model of $\psi_{\mathcal{M}, \bar{w}_0}$ embeds an acceptance tree for \mathcal{M} on input \bar{w}_0 . Conversely, any acceptance trees for \mathcal{M} on input \bar{w}_0 can be expanded to a model of $\psi_{\mathcal{M}, \bar{w}_0}$ by interpreting the relevant predicates as suggested above. We conclude:

Theorem 5.6.2. *The finite and general satisfiability problems for the (equality-free) guarded adjacent fragment is 2EXPTIME-hard. Thus, the finite and general satisfiability problems for the guarded adjacent fragment is 2EXPTIME-complete.*

5.7 Extending the Adjacent Fragment

The adjacent fragment \mathcal{AF} is defined as the union of the formulas sets $\mathcal{AF}^{[k]}$, each of which restricts the allowed argument sequences appearing in atomic formulas to what we call adjacent words over the alphabet \mathbf{x}_k . The question arises as to whether these restrictions might be further relaxed without compromising the decidability of satisfiability. Under reasonable assumptions about the fragment in question, the answer must be no, as we formally show next. Take any function $f : [1, m] \rightarrow [1, k]$ that is not a walk (i.e. there is some position $1 \leq j < m$ for which $|f(j+1) - f(j)| \geq 2$). Then, \mathcal{AF}^f denotes the extension of \mathcal{AF} obtained by allowing atoms of the form $\mathbf{p}(\mathbf{x}_k^f)$. That is to say, f is a “honorary adjacent walk” that is considered in Item (1) of the definition of $\mathcal{AF}^{[k]}$ (page 54). As it turns out, \mathcal{AF}^f is powerful enough to express that a given relation is transitive.

Lemma 5.7.1. *Let \mathbf{T} be a binary predicate, and \mathbf{Q} an m -ary predicate. There exists a formula $\varphi_{\mathbf{T}, \mathbf{Q}}$ of \mathcal{AF}^f , such that: (i) if $\mathfrak{A} \models \varphi_{\mathbf{T}, \mathbf{Q}}$, then $\mathbf{T}^{\mathfrak{A}}$ is transitive, and (ii) any structure \mathfrak{A} interpreting \mathbf{T} as a transitive relation, but not interpreting \mathbf{Q} , can be expanded to a model of $\varphi_{\mathbf{T}, \mathbf{Q}}$.*

Proof. Since f is not a walk, let us fix an index $j \in [1, m-1]$ for which $|f(j+1) - f(j)| \geq 2$. We may assume, without loss of generality, that $f(j+1) > f(j)$, as the proof for the other case is obtained by swapping all occurrences of j and $j+1$. Notice that, in this case, we must have $f(j) < f(j+1) - 1$ and hence

$f(j) < k-1$. Define $\varphi_{\mathbb{T},\mathbb{Q}} := \forall \mathbf{x}_m \varphi_{\mathbb{T},\mathbb{Q}}^1 \wedge \forall \mathbf{x}_k \varphi_{\mathbb{T},\mathbb{Q}}^2$, where

$$\begin{aligned} \varphi_{\mathbb{T},\mathbb{Q}}^1 &:= \mathbb{Q}(\mathbf{x}_m) \rightarrow \mathbb{T}(x_j x_{j+1}), \\ \varphi_{\mathbb{T},\mathbb{Q}}^2 &:= \left(\mathbb{T}(x_{f(j)} x_{f(j)+1}) \wedge \mathbb{T}(x_{f(j)+1} x_{f(j)+2}) \wedge \bigwedge_{i=f(j)+2}^{f(j+1)-1} x_i = x_{i+1} \right) \rightarrow \mathbb{Q}(\mathbf{x}_k^f). \end{aligned}$$

To establish (i), we take any $\mathfrak{A} \models \varphi_{\mathbb{T},\mathbb{Q}}$ and any elements a, b, c of \mathfrak{A} with $ab \in \mathbb{T}^{\mathfrak{A}}$ and $bc \in \mathbb{T}^{\mathfrak{A}}$. We must show that $ac \in \mathbb{T}^{\mathfrak{A}}$. Let $\bar{d} = d_1 \cdots d_k$ be a tuple with $d_{f(j)} = a$, $d_{f(j)+1} = b$, and $d_i = c$ for all i except $f(j)$ and $f(j)+1$. Thus, $\mathfrak{A} \models \mathbb{T}[d_{f(j)} d_{f(j)+1}]$, $\mathfrak{A} \models \mathbb{T}[d_{f(j)+1} d_{f(j)+2}]$, and $d_{f(j)+2} = \cdots = d_{f(j+1)}$. Since $\mathfrak{A} \models \varphi_{\mathbb{T}}^2[\bar{d}]$, we have $\mathfrak{A} \models \mathbb{Q}[\bar{d}^f]$. But the j th position of \bar{d}^f is occupied by $d_{f(j)} = a$, and the $(j+1)$ th position is occupied by $d_{f(j+1)} = c$; and since $\mathfrak{A} \models \varphi_{\mathbb{T}}^1[\bar{d}^f]$, we have $ac \in \mathbb{T}^{\mathfrak{A}}$, as required.

To establish (ii), suppose \mathfrak{A} interprets \mathbb{T} as a transitive relation. We expand \mathfrak{A} to \mathfrak{A}^+ by fixing the interpretation of \mathbb{Q} to be the set of all m -tuples \bar{a} such that $a_j a_{j+1} \in \mathbb{T}^{\mathfrak{A}}$. It is immediate that $\mathfrak{A}^+ \models \forall \mathbf{x}_m \varphi_{\mathbb{T},\mathbb{Q}}^1$. Now take any k -tuple \bar{c} satisfying the antecedent of $\varphi_{\mathbb{T},\mathbb{Q}}^2$ in \mathfrak{A}^+ . Thus, $c_{f(j)} c_{f(j)+1}$ and $c_{f(j)+1} c_{f(j)+2}$ are both in the relation $\mathbb{T}^{\mathfrak{A}}$, and, moreover, $c_{f(j)+2} = \cdots = c_{f(j+1)}$. By transitivity, $c_{f(j)} c_{f(j)+2} \in \mathbb{T}^{\mathfrak{A}}$, whence $c_{f(j)} c_{f(j+1)} \in \mathbb{T}^{\mathfrak{A}}$. But $c_{f(j)}$ and $c_{f(j+1)}$ are, respectively, the j th and $(j+1)$ th element of \bar{c}^f , and so, by construction, $\mathfrak{A}^+ \models \mathbb{Q}[\bar{c}^f]$. Thus, $\mathfrak{A}^+ \models \forall \mathbf{x}_k \varphi_{\mathbb{T},\mathbb{Q}}^2$. \square

Noting that \mathcal{FO}^2 with 2 transitive relations has undecidable finite and general satisfiability problems [Kieronski, 2005], the same is almost immediate for \mathcal{AF}^f given the lemma above. Indeed, by taking any \mathcal{FO}^2 -sentence φ we have, by Theorem 5.3.3, that it is equivalent to an \mathcal{AF} -sentence, say, φ' . Combining this observation with Lemma 5.7.1 we have that φ is satisfiable in a (finite) model in which binary relations \mathbb{T} and \mathbb{T}' are transitive if and only if $\varphi' \wedge \varphi_{\mathbb{T},\mathbb{Q}} \wedge \varphi_{\mathbb{T}',\mathbb{Q}'}$ is (finitely) satisfiable; it being understood that \mathbb{Q}, \mathbb{Q}' are not in the signature of φ . We thus conclude:

Theorem 5.7.2. *Take $f: [1, m] \rightarrow [1, k]$ to be a function that is not a walk. Then, the finite and general satisfiability problems for \mathcal{AF}^f are undecidable.*

Having ruled out syntactic relaxations as a way of generalising the adjacent fragment, we now turn to extensions such as threshold, periodic quantification.

Having positive cases (in regard to decidability of satisfiability) for the fluted fragment such as in [Pratt-Hartmann, 2021] and Theorem 4.3.5, it is natural to conjecture that the same holds for the adjacent fragment with the appropriate extensions, especially since the base cases, i.e. \mathcal{C}^2 and $\mathcal{FO}_{\text{Pres}}^2$, have already been shown to have a decidable finite and general satisfiability problem in [Pratt-Hartmann, 2010] and [Benedikt et al., 2024]. Moreover, having already established a variable reduction procedure that preserves domain elements in Section 5.5 it should be all but certain that the adjacent fragment with counting extensions retains decidability. This will be the subject of the following two chapters.

Chapter 6

The Adjacent Fragment with Counting is Undecidable

The following chapter is an expanded version of Section 5 of the conference paper [Kojelis, 2025]. All results presented are due to the Ph. d. candidate.

In the following two chapters we consider the adjacent fragment augmented with various kinds of counting quantifiers. Informally, in the *adjacent fragment with (periodic) counting*, we allow $\exists_{[\geq n]}$ (resp. $\exists_{[n+p]}$) in place of \exists . More formally, we define the formulas $\mathcal{AFC}^{[\ell]}$ and $\mathcal{AFPC}^{[\ell]}$ by simultaneous structural induction for $\ell \geq 0$ as follows:

1. every adjacent ℓ -atom is in $\mathcal{AFC}^{[\ell]}$ and $\mathcal{AFPC}^{[\ell]}$;
2. $\mathcal{AFC}^{[\ell]}$ and $\mathcal{AFPC}^{[\ell]}$ are closed under Boolean combinations;
3. if φ is in $\mathcal{AFC}^{[\ell+1]}$, then $\exists_{[\geq n]}x \varphi$ is in $\mathcal{AFC}^{[k]}$ for all $k \geq \ell$ and $n \in \mathbb{N}$.
4. if φ is in $\mathcal{AFPC}^{[\ell+1]}$, then $\exists_{[n+p]}x \varphi$ is in $\mathcal{AFPC}^{[k]}$ for all $k \geq \ell$ and $n, p \in \mathbb{N}$.

We then define the adjacent fragment with counting to be $\mathcal{AFC} := \bigcup_{\ell \geq 0} \mathcal{AFC}^{[\ell]}$, and the adjacent fragment with periodic counting to be $\mathcal{AFPC} := \bigcup_{\ell \geq 0} \mathcal{AFPC}^{[\ell]}$. The ℓ -variable variants of the languages are denoted as $\mathcal{AFC}^\ell := \mathcal{AFC} \cap \mathcal{FO}^\ell$ and $\mathcal{AFPC}^\ell := \mathcal{AFPC} \cap \mathcal{FO}^\ell$ respectively. To avoid notational clutter in \mathcal{AFC} , we allow formulas of the form $\exists_{[\leq n]}x \theta$ and $\exists_{[=n]}x \theta$ in place of their more formal counterparts $\neg \exists_{[\geq n+1]}x \theta$ and $\exists_{[\leq n]}x \theta \wedge \exists_{[\geq n]}x \theta$. Additionally we, use $\exists x \theta$ instead of $\exists_{[\geq 1]}x \theta$ and $\forall x \theta$ in place of $\exists_{[=0]}x \neg \theta$. Since (standard) counting quantifiers can be simulated using periodic counting, we assume that all formulas permitted in \mathcal{AFC} are also legal in \mathcal{AFPC} .

The adjacent fragment with (periodic) counting quantifiers extends the fluted fragment with (periodic) counting. It is thus immediate that \mathcal{AFC} is a multi-variable extension of description logics such as \mathcal{ALCHQ} and, since relations of the form $\mathbf{r}(x_\ell \cdots x_1)$ are allowed, \mathcal{ALCHIQ} . Since the guarded fragment is undecidable for both finite and general satisfiability in the presence of counting quantifiers [Grädel, 1999], we are motivated to ask if \mathcal{AFC} is a decidable (in terms of finite and general satisfiability) multi-variable extension of \mathcal{ALCHIQ} .

In this chapter we show that all hope is lost for decidability of the finite and (for the most part) general satisfiability problems for the adjacent fragment with counting quantifiers. In Section 6.1 we construct a \mathcal{AFC} -sentence φ with as few as 3 variables that is finitely satisfiable if and only if a given Diophantine equation \mathcal{E} has a solution over \mathbb{N} . Finding solutions to such equations is known as *Hilbert's 10th problem* and is Σ_1^0 -complete when considered over \mathbb{N} . Since finite models can be enumerated, our reduction will render $\text{FinSat}(\mathcal{AFC}^3)$ (and thus $\text{FinSat}(\mathcal{AFC})$) complete for Σ_1^0 . When periodic counting quantifiers are allowed, the Σ_1^0 -hardness result transfers to the general satisfiability problem. Assuming the former result, the latter is immediate as we have:

$$(\varphi \wedge \exists_{[0+1]x} \top) \in \text{Sat}(\mathcal{AFP}^3) \iff \varphi \in \text{FinSat}(\mathcal{AFC}^3).$$

We note that, in our construction, φ will be satisfiable in any (i.e. possibly infinite) model if and only if \mathcal{E} has a solution over $\mathbb{N}^* = \mathbb{N} \cup \{\aleph_0\}$. This, however, does not imply undecidability of general satisfiability for \mathcal{AFC}^3 , as solutions to Hilbert's 10th problem over \mathbb{N}^* can be found in NPTIME [Jeřábek, 2016].

To show undecidability of general satisfiability, we thus opt for a different approach. Consider the tuple $\Phi := \langle \mathcal{T}, \mathcal{H}, \mathcal{V} \rangle$, where $\mathcal{H}, \mathcal{V} \subseteq \mathcal{T} \times \mathcal{T}$. We regard \mathcal{T} as being a set of *tiles* whilst \mathcal{H} and \mathcal{V} are respectively called *horizontal* and *vertical tiling constraints*. We say that Φ tiles the infinite $(\mathbb{N} \times \mathbb{N})$ -plane just in case there is a function $f : \mathbb{N} \times \mathbb{N} \rightarrow \mathcal{T}$ such that $\langle f(i, j), f(i, j+1) \rangle \in \mathcal{H}$ and $\langle f(i, j), f(i+1, j) \rangle \in \mathcal{V}$ for each $i, j \in \mathbb{N}$. Determining the existence of f is Π_1^0 -complete [Berger, 1966]. We show undecidability of $\text{Sat}(\mathcal{AFC}^4)$ by producing an \mathcal{AFC}^4 -sentence that is satisfiable if and only if Φ tiles the infinite $(\mathbb{N} \times \mathbb{N})$ -plane. We further show that, if periodic counting is permitted, then it is possible to request, via an \mathcal{AFP}^2 -sentence, that a designated tile appears infinitely often in the first column of the tiling. Noting that the requirement of infinite recurrence renders the tiling problem Σ_1^1 -complete [Harel, 1984], we will be able to classify

$\text{Sat}(\mathcal{AFPC}^4)$ as Σ_1^1 -hard.

The reader might have noticed that there is an odd gap in terms of the undecidability results mentioned. Whilst both finite and general satisfiability for \mathcal{AFPC} is undecidable with as little as 3 variables, the same cannot be said for \mathcal{AFC} . Indeed, for \mathcal{AFC} the minimum number of variables required to establish undecidability of general satisfiability is 4. For finite satisfiability the number is 3. Anticipating Chapter 7, we mention that this is not a shortcoming of the methodology used here but rather an idiosyncrasy of adjacent fragment with counting.

Note that whilst the target language for the undecidability results is \mathcal{AF} with the appropriate counting extensions, the formulas in the reduction do not use all features of adjacency, nor do they feature the equality symbol. We mention now that all formulas defined are actually in the fluted fragment with appropriate counting extensions and *reversed variable sequences*. In this language, if $\mathbf{r}(x_i \cdots x_\ell)$ is an ℓ -atom, then so is $\mathbf{r}(x_\ell \cdots x_i)$.

For readability, we default to variables x, y, z, w when constructing formulas. We use (sub)sequences of $xyzw$ and $wzyx$ in atoms and quantification. By renaming $x \mapsto x_1, y \mapsto x_2, z \mapsto x_3$ and $w \mapsto x_4$ when the quantification order is $xyzw$, (and symmetrically when it is $wzyx$) it is easy to verify that every formula is in \mathcal{AF} with the appropriate counting extensions.

6.1 Undecidability of Finite Satisfiability

We proceed by reducing Hilbert's 10th problem to the finite satisfiability problem of \mathcal{AFC}^3 . Let \mathcal{E} be a system of Diophantine equations. We assume that each equation $e \in \mathcal{E}$ is of one of the following (simple) forms:

- (i) $u = 1$,
- (ii) $u + v = w$, or
- (iii) $u \cdot v = w$,

where u, v, w are mutually disjoint variables. Clearly, no loss of generality is incurred as any Diophantine equation can be rewritten into the simpler form by introducing new variables. For each $e \in \mathcal{E}$ we will define a formula φ_e depending

on the form that e takes. Then, φ , defined as (n.b. ψ is given later):

$$\bigwedge_{e \in \mathcal{E}} \varphi_e \wedge \psi, \quad (\varphi)$$

will be the advertised formula that is finitely satisfiable if and only if \mathcal{E} has a solution over \mathbb{N} . We specify that the signature of φ includes

- (1) unary predicates \mathbf{A}_u for each variable of u in \mathcal{E} ,
- (2) binary predicates \mathbf{R}_e for each $e \in \mathcal{E}$ of the form (ii), and
- (3) ternary predicates \mathbf{P}_e for each $e \in \mathcal{E}$ of the form (iii).

In the sequel we will argue that

- if $\mathfrak{A} \models \varphi$, then \mathcal{E} has a solution with $\pi^{\mathfrak{A}}(u) := |\mathbf{A}_u^{\mathfrak{A}}|$ for each variable u , and
- if \mathcal{E} has a satisfying assignment π , then there is a structure $\mathfrak{A} \models \varphi$ satisfying $|\mathbf{A}_u^{\mathfrak{A}}| = \pi(u)$ for each variable u .

We proceed with the definition of φ . For technical reasons we wish for the sets $\mathbf{A}_u^{\mathfrak{A}}$ and $\mathbf{A}_v^{\mathfrak{A}}$ with $u \neq v$ to be disjoint in any structure \mathfrak{A} . Denoting $\text{vars}(\mathcal{E})$ for the set of variables in \mathcal{E} , we first define ψ to be

$$\bigwedge_{\substack{u \neq v \\ u, v \in \text{vars}(\mathcal{E})}} \forall x (\neg \mathbf{A}_u(x) \vee \neg \mathbf{A}_v(x)), \quad (\psi)$$

which clearly has the required effect. We proceed by taking $e \in \mathcal{E}$ in turn.

Suppose first that e is of the form (i) $u = 1$. We ensure that every model \mathfrak{A} of φ will have $|\mathbf{A}_u^{\mathfrak{A}}| = 1$ by defining φ_e to be

$$\exists_{[=1]} x \mathbf{A}_u(x).$$

Now, supposing that e is of the form (ii) $u + v = w$, we define φ_e with the intent that models \mathfrak{A} of φ will have $\mathbf{R}_e^{\mathfrak{A}}$ be a bijection between $\mathbf{A}_u^{\mathfrak{A}} \cup \mathbf{A}_v^{\mathfrak{A}}$ and $\mathbf{A}_w^{\mathfrak{A}}$ (which, together with our requirement that $u \neq v$, implies that $|\mathbf{A}_u^{\mathfrak{A}}| + |\mathbf{A}_v^{\mathfrak{A}}| = |\mathbf{A}_w^{\mathfrak{A}}|$):

$$\begin{aligned} & \forall x \left((\mathbf{A}_u(x) \vee \mathbf{A}_v(x)) \rightarrow \exists_{[=1]} y (\mathbf{A}_w(y) \wedge \mathbf{R}_e(xy)) \right) \wedge \\ & \forall y \left(\mathbf{A}_w(y) \rightarrow \exists_{[=1]} x ((\mathbf{A}_u(x) \vee \mathbf{A}_v(x)) \wedge \mathbf{R}_e(xy)) \right). \end{aligned}$$

Lastly, if e is of the form (iii) $u \cdot v = w$, then with φ_e we will have that $\mathfrak{A} \models \varphi$ entails $P_e^{\mathfrak{A}}$ being a bijection between $\mathbf{A}_u^{\mathfrak{A}} \times \mathbf{A}_v^{\mathfrak{A}}$ and $\mathbf{A}_w^{\mathfrak{A}}$ (thus making $|\mathbf{A}_u^{\mathfrak{A}}| \cdot |\mathbf{A}_v^{\mathfrak{A}}| = |\mathbf{A}_w^{\mathfrak{A}}|$):

$$\begin{aligned} & \forall x \left(\mathbf{A}_u(x) \rightarrow \forall y \left(\mathbf{A}_v(y) \rightarrow \exists_{[=1]} z \left(\mathbf{A}_w(z) \wedge P_e(xyz) \right) \right) \right) \wedge \\ & \forall z \left(\mathbf{A}_w(z) \rightarrow \exists_{[=1]} y \left(\mathbf{A}_v(y) \wedge \exists x \left(\mathbf{A}_u(x) \wedge P_e(xyz) \right) \right) \right) \wedge \\ & \forall z \left(\mathbf{A}_w(z) \rightarrow \exists y \left(\mathbf{A}_v(y) \wedge \exists_{[=1]} x \left(\mathbf{A}_u(x) \wedge P_e(xyz) \right) \right) \right). \end{aligned}$$

We argue that the construction above satisfies the advertised properties by showing the following:

Lemma 6.1.1. *Suppose \mathcal{E} is an instance of Hilbert's 10th problem and φ is computed from \mathcal{E} as described above. Then, \mathcal{E} has a solution over \mathbb{N} if and only if φ is finitely satisfiable.*

Proof. For the “if” direction suppose $\mathfrak{A} \models \varphi$. We claim that $\pi^{\mathfrak{A}} := \{u \mapsto |\mathbf{A}_u^{\mathfrak{A}}| \mid u \in \text{vars}(\mathcal{E})\}$ is a satisfying assignment for \mathcal{E} . Thus, again taking $e \in \mathcal{E}$ in turn, we have that if e is of the form (i) $u = 1$, then $\mathfrak{A} \models \varphi_e \implies |\mathbf{A}_u^{\mathfrak{A}}| = 1$. If e takes the form (ii) $u + v = w$, we then claim that $R_e^{\mathfrak{A}}$ is a bijection between $\mathbf{A}_u^{\mathfrak{A}} \cup \mathbf{A}_v^{\mathfrak{A}}$ and $\mathbf{A}_w^{\mathfrak{A}}$. Indeed, by the first conjunct of φ_e we have that each element in $\mathbf{A}_u^{\mathfrak{A}} \cup \mathbf{A}_v^{\mathfrak{A}}$ is paired with a single element in $\mathbf{A}_w^{\mathfrak{A}}$; the converse is established by the second conjunct. Thus, $|\mathbf{A}_u^{\mathfrak{A}} \cup \mathbf{A}_v^{\mathfrak{A}}| = |\mathbf{A}_w^{\mathfrak{A}}|$. Additionally, by the requirement that $u \neq v$ and since $\mathfrak{A} \models \psi$, we have that $\mathbf{A}_u^{\mathfrak{A}} \cap \mathbf{A}_v^{\mathfrak{A}} = \emptyset$. This gives us $|\mathbf{A}_u^{\mathfrak{A}} \cup \mathbf{A}_v^{\mathfrak{A}}| = |\mathbf{A}_u^{\mathfrak{A}}| + |\mathbf{A}_v^{\mathfrak{A}}|$ as required. Lastly, suppose e takes the form (iii) $u \cdot v = w$. By the first conjunct of φ_e we have that for each $ab \in \mathbf{A}_u^{\mathfrak{A}} \times \mathbf{A}_v^{\mathfrak{A}}$ there is a single $c \in \mathbf{A}_w^{\mathfrak{A}}$ such that $abc \in P_e^{\mathfrak{A}}$. Hence, $P_e^{\mathfrak{A}}$ gives rise to a function

$$f := \{ab \mapsto c \mid abc \in P_e^{\mathfrak{A}} \cap (\mathbf{A}_u^{\mathfrak{A}} \times \mathbf{A}_v^{\mathfrak{A}} \times \mathbf{A}_w^{\mathfrak{A}})\}.$$

Writing $f(xy) = z$ in place of the atom $P_e(xyz)$, we claim that f is a bijection between $\mathbf{A}_u^{\mathfrak{A}} \times \mathbf{A}_v^{\mathfrak{A}}$ and $\mathbf{A}_w^{\mathfrak{A}}$. It is easily seen that f is surjective from the second conjunct of φ_e . To establish injectivity suppose $f(ab) = f(a'b') = c$ for some $a, b, a', b' \in A$. But, by the second conjunct of φ_e , we have that:

$$\mathfrak{A}, c \models \exists_{[=1]} y \left(\mathbf{A}_v(y) \wedge \exists x \left(\mathbf{A}_u(x) \wedge f(xy) = z \right) \right).$$

Thus, $b = b'$. Notice that this establishes that b is the only element in $\mathbf{A}_v^{\mathfrak{A}}$ for

which, under the assignment $c \mapsto z, b \mapsto y$, we have:

$$\mathfrak{A}, cb \models \exists x (\mathbf{A}_u(x) \wedge f(xy) = z).$$

Combining this fact with the third conjunct of φ_e we have, again under the assignment $c \mapsto z, b \mapsto y$, that:

$$\mathfrak{A}, cb \models \exists_{[=1]} x (\mathbf{A}_u(x) \wedge f(xy) = z).$$

Clearly, $a = a'$ and thus $|\mathbf{A}_u^{\mathfrak{A}}| \cdot |\mathbf{A}_v^{\mathfrak{A}}| = |\mathbf{A}_u^{\mathfrak{A}} \times \mathbf{A}_v^{\mathfrak{A}}| = |\mathbf{A}_w^{\mathfrak{A}}|$ as required.

For the “only-if” direction, suppose \mathcal{E} has a solution. Let $\pi : \text{vars}(\mathcal{E}) \rightarrow \mathbb{N}$ be the satisfying assignment for \mathcal{E} and construct a model \mathfrak{A} of φ as follows. For each variable $u \in \text{vars}(\mathcal{E})$ define $\mathbf{A}_u^{\mathfrak{A}}$ to be a set of $\pi(u)$ distinct elements and set the domain A of \mathfrak{A} to be the disjoint union of $\mathbf{A}_u^{\mathfrak{A}}$, where $u \in \text{vars}(\mathcal{E})$. Clearly, $\mathfrak{A} \models \psi$. If φ_e was constructed from $e \in \mathcal{E}$ of the form (i) $u = 1$, then $|\mathbf{A}_u^{\mathfrak{A}}| = \pi(u) = 1$ as required by φ_e . On the other hand, if e is of the form (ii) $u + v = w$, we have that $\mathfrak{A} \models \varphi_e$ by setting $\mathbf{R}_e^{\mathfrak{A}}$ to be a bijection between $\mathbf{A}_u^{\mathfrak{A}} \cup \mathbf{A}_v^{\mathfrak{A}}$ and $\mathbf{A}_w^{\mathfrak{A}}$ (this can be done as $\pi(u) + \pi(v) = \pi(w)$ and, by initial assumption, $u \neq v$). Lastly, if e is (iii) $u \cdot v = w$, then $\pi(u) \cdot \pi(v) = \pi(w)$. Thus, index elements of $\mathbf{A}_u^{\mathfrak{A}}$ as $a_1 \dots a_{\pi(u)}$, elements of $\mathbf{A}_v^{\mathfrak{A}}$ as $b_1 \dots b_{\pi(v)}$ and elements of $\mathbf{A}_w^{\mathfrak{A}}$ as $(c_{i,j})_{\substack{1 \leq i \leq \pi(u) \\ 1 \leq j \leq \pi(v)}}$. Clearly, by setting $\mathbf{P}_e^{\mathfrak{A}} := \{a_i b_j c_{i,j} \mid 1 \leq i \leq \pi(u), 1 \leq j \leq \pi(v)\}$ we have that $\mathfrak{A} \models \varphi_e$ thus concluding the proof. \square

Noting again that $(\varphi \wedge \exists_{[0+1]} x. \top) \in \text{Sat}(\mathcal{AFPC}^3) \iff \varphi \in \text{FinSat}(\mathcal{AFC}^3)$ we have proved the following theorem:

Theorem 6.1.2. *The finite satisfiability problem for \mathcal{AFC}^3 is Σ_1^0 -complete. The finite and general satisfiability problems for \mathcal{AFPC}^3 are Σ_1^0 -hard.*

6.2 Undecidability of General Satisfiability

As mentioned before, the reduction in the previous section does not imply undecidability of general satisfiability for \mathcal{AFC} . Thus, to show such undecidability results we resort to tiling problems. Given an instance $\Phi = \langle \mathcal{T}, \mathcal{H}, \mathcal{V} \rangle$ we produce an \mathcal{AFC} -sentence φ in 4 variables that is satisfiable if and only if Φ tiles the infinite $(\mathbb{N} \times \mathbb{N})$ -plane. We then argue that one can append additional \mathcal{AFPC}^2 conjuncts to φ and thus obtain a reduction from the *recurring tiling problem*. That is, the

problem of determining if a given instance Φ tiles the $(\mathbb{N} \times \mathbb{N})$ -plane with a designated tile appearing infinitely often on the first column. Such reductions classify $\text{Sat}(\mathcal{AFC}^4)$ as Π_1^0 -hard, whilst guaranteeing Σ_1^1 -hardness for $\text{Sat}(\mathcal{AFPC}^4)$.

Take \mathbf{G} to be a unary predicate and \mathbf{H} and \mathbf{V} to be binary predicates. We define the *canonical* $(\mathbb{N} \times \mathbb{N})$ -grid to be a $\{\mathbf{G}, \mathbf{H}, \mathbf{V}\}$ -structure \mathfrak{G} over the domain $\mathbb{N} \times \mathbb{N}$ with the following extensions:

- $\mathbf{G}^{\mathfrak{A}} := \mathbb{N} \times \mathbb{N}$,
- $\mathbf{H}^{\mathfrak{A}} := \{\langle (i, j), (i, j+1) \rangle \mid i, j \in \mathbb{N}\}$, and
- $\mathbf{V}^{\mathfrak{A}} := \{\langle (i, j), (i+1, j) \rangle \mid i, j \in \mathbb{N}\}$.

We say that a structure \mathfrak{A} is a $(\mathbb{N} \times \mathbb{N})$ -grid if \mathfrak{A} restricted to elements $\mathbf{G}^{\mathfrak{A}}$ and the signature $\{\mathbf{G}, \mathbf{H}, \mathbf{V}\}$ is isomorphic to the canonical $(\mathbb{N} \times \mathbb{N})$ -grid. More leniently, \mathfrak{A} is *grid-like* if it contains a homomorphic embedding of \mathfrak{G} . It is well known that the satisfiability problem posed over grid-like structures is undecidable even for less expressive logics such as \mathcal{FL}^2 . Delaying the axiomatisation of grid-like structures by just a bit, we first show the following:

Proposition 6.2.1. *The satisfiability problem for \mathcal{AFC}^4 over grid-like structures is Π_1^0 -hard. The satisfiability problem for \mathcal{AFPC}^4 over $(\mathbb{N} \times \mathbb{N})$ -grids is Σ_1^1 -hard.*

Proof. Let us fix $\Phi = \langle \mathcal{T}, \mathcal{H}, \mathcal{V} \rangle$ to be some instance of the tiling problem with \mathcal{T} being a set of tiles and $\mathcal{H}, \mathcal{V} \subseteq \mathcal{T} \times \mathcal{T}$ being a set of constraints. For the first statement we need only secure that positions are tiled by a single tile and that adjacent tiles respect the rules imposed by \mathcal{H} and \mathcal{V} . For this purpose we identify tiles $t \in \mathcal{T}$ as unary predicates \mathfrak{t} and write:

$$\forall x \left(\bigvee_{t \in \mathcal{T}} \mathfrak{t}(x) \wedge \bigwedge_{\substack{t \neq t' \\ t, t' \in \mathcal{T}}} (\neg \mathfrak{t}(x) \vee \neg \mathfrak{t}'(x)) \right), \quad (\psi_1)$$

$$\forall xy \left((\mathbf{G}(x) \wedge \mathbf{G}(y) \wedge \mathbf{H}(xy)) \rightarrow \bigvee_{(t, t') \in \mathcal{H}} (\mathfrak{t}(x) \wedge \mathfrak{t}'(y)) \right), \quad (\psi_2)$$

$$\forall xy \left((\mathbf{G}(x) \wedge \mathbf{G}(y) \wedge \mathbf{V}(xy)) \rightarrow \bigvee_{(t, t') \in \mathcal{V}} (\mathfrak{t}(x) \wedge \mathfrak{t}'(y)) \right). \quad (\psi_3)$$

Writing $\psi := \psi_1 \wedge \psi_2 \wedge \psi_3$ we claim that there is a grid-like model of ψ if and only if Φ is a positive instance of the tiling problem. On the one hand, if Φ is a positive instance of the problem, then, by taking $f : \mathbb{N} \times \mathbb{N} \rightarrow \mathcal{T}$ to be the

function witnessing the tiling, we have that any $(\mathbb{N} \times \mathbb{N})$ -grid \mathfrak{A} can be expanded into a model \mathfrak{A}^+ of ψ . Supposing $h : \mathbb{N} \times \mathbb{N} \rightarrow \mathbf{G}^{\mathfrak{A}}$ is the promised isomorphism from \mathfrak{G} to \mathfrak{A} (restricted to elements in $\mathbf{G}^{\mathfrak{A}}$ and the signature $\{\mathbf{G}, \mathbf{H}, \mathbf{V}\}$), we obtain the required result by setting $h(i, j) \in \mathfrak{t}^{\mathfrak{A}^+}$ if and only if $f(i, j) = t$ for each $i, j \in \mathbb{N}$. Referencing the definitions of $\mathbf{H}^{\mathfrak{G}}$ and $\mathbf{V}^{\mathfrak{G}}$ it is immediate that $\mathfrak{A}^+ \models \psi$.

Now, on the other hand, if \mathfrak{A} is a grid-like model of ψ , then a tiling function $f : \mathbb{N} \times \mathbb{N} \rightarrow \mathcal{T}$ can be constructed as follows. Let $h : \mathbb{N} \times \mathbb{N} \hookrightarrow \mathbf{G}^{\mathfrak{A}}$ be the function embedding the canonical $(\mathbb{N} \times \mathbb{N})$ -grid into \mathfrak{A} . By ψ_1 we have that each $h(i, j)$ is in exactly one $\mathfrak{t}^{\mathfrak{A}}$, where $t \in \mathcal{T}$. Thus, the following assignment is well-defined: $f(i, j) := t$ if and only if $h(i, j) \in \mathfrak{t}^{\mathfrak{A}}$ for each $i, j \in \mathbb{N}$ and $t \in \mathcal{T}$. To see that f witnesses a tiling for Φ take any $i, j \in \mathbb{N}$ and suppose $f(i, j) = t$, whilst $f(i, j+1) = t'$. Since $\langle h(i, j), h(i, j+1) \rangle \in \mathbf{H}^{\mathfrak{A}}$, we must have, by ψ_2 , that $(t, t') \in \mathcal{H}$ as required. Utilising ψ_3 the same can be shown for positions (i, j) , $(i+1, j)$ and the constraint \mathcal{V} .

Considering the second statement of the lemma, let $t^\times \in \mathcal{T}$ be some designated tile. When the satisfiability problem is posed over $(\mathbb{N} \times \mathbb{N})$ -grids, we can force, via a \mathcal{AFPC}^2 sentence θ , the infinite recurrence of t^\times on the first column as follows:

$$\neg \exists_{[0+1]y} (\mathfrak{t}^\times(y) \wedge \mathbf{G}(y) \wedge \forall x \neg \mathbf{H}(xy)). \quad (\theta)$$

It is easy to see (using the same argument as given above) that $(\mathbb{N} \times \mathbb{N})$ -grids can be expanded to a model of $\psi \wedge \theta$ when given a tiling of Φ with t^\times appearing infinitely often on the first column. The converse also holds: if \mathfrak{A} is a $(\mathbb{N} \times \mathbb{N})$ -grid structure satisfying $\psi \wedge \theta$, then, utilising the same construction as above, we will have a tiling f of Φ . By θ , there is an infinite number of elements in $\mathbf{G}^{\mathfrak{A}}$ that have no \mathbf{H} -predecessor and satisfy \mathfrak{t}^\times . By the definition of $(\mathbb{N} \times \mathbb{N})$ -grids, the structure \mathfrak{A} restricted to $\mathbf{G}^{\mathfrak{A}}$ and the signature $\{\mathbf{G}, \mathbf{H}, \mathbf{V}\}$ is isomorphic to the canonical $(\mathbb{N} \times \mathbb{N})$ -grid. Thus, by the construction above, t^\times appears infinitely often on the first column in the tiling. \square

The lemma above is, of course, the “easy” part of a much larger reduction. The axiomatisation of grid-like structures and $(\mathbb{N} \times \mathbb{N})$ -grids is where the expressive power of \mathcal{AFC}^4 and \mathcal{AFPC}^4 is needed. Before writing the advertised formulas, we build the motivating structure we will be looking for in three steps. Suppose \mathfrak{G} is the canonical $(\mathbb{N} \times \mathbb{N})$ -grid. Letting $\mathbf{E}_\mathbf{H}$, $\mathbf{E}_\mathbf{V}$ be binary and $\mathbf{0}$ be unary predicates, we define the *graphed expansion* of \mathfrak{G} to be the structure \mathfrak{G}^+ over the domain

$(\mathbb{N} \times \mathbb{N}) \cup \mathbb{N}$ with the following extensions:

- $\mathfrak{G}^+|_{\mathbb{N} \times \mathbb{N}} := \mathfrak{G}$,
- $k \notin \mathfrak{G}^+$ for each $k \in \mathbb{N}$,
- $k \in \mathfrak{O}^+$ if and only if $k = 0$,
- $\mathbf{E}_H^+ := \bigcup_{i,j \in \mathbb{N}} \{\langle (i, j), k \rangle \mid 1 \leq k \leq j\}$, and
- $\mathbf{E}_V^+ := \bigcup_{i,j \in \mathbb{N}} \{\langle (i, j), k \rangle \mid 1 \leq k \leq i\}$.

Intuitively, \mathfrak{G}^+ , when restricted to $\mathbb{N} \times \mathbb{N} = \mathfrak{G}^+$, is the canonical $(\mathbb{N} \times \mathbb{N})$ -grid. Notice that each $(i, j) \in \mathbb{N} \times \mathbb{N}$ has j elements in \mathbb{N} that are \mathbf{E}_H -successors and i elements that are \mathbf{E}_V -successors. In other words, the coordinates of (i, j) are explicitly encoded in \mathfrak{G}^+ as the out-degrees of \mathbf{E}_H and \mathbf{E}_V respectively. (In the future, we will simply speak of \mathbf{E}_H - and \mathbf{E}_V -degree with “out” being left implicit). We invite the reader to regard \mathbb{N} as the set of extra elements which help encode positions of grid elements. Notice that the singleton $0 \in \mathfrak{O}^+$ is not featured in any binary relations (most notably, \mathbf{E}_H and \mathbf{E}_V). This is deliberate, as it will act as a *spare part* in the constructions to come.

We now define *the mapped expansion* \mathfrak{G}^* of \mathfrak{G}^+ , where \mathfrak{G}^+ itself is the graphed expansion of \mathfrak{G} . For this, we introduce quaternary predicates $\mathbf{R}_H, \mathbf{R}_V, \mathbf{S}_H, \mathbf{S}_V$ and ternary predicates $\mathbf{C}_H, \mathbf{C}_V$, whilst setting the following extensions:

- $\mathbf{R}_H^{\mathfrak{G}^*} := \bigcup_{i,j,i',j' \in \mathbb{N}}^{j \leq j'} \{\langle k, (i, j), (i', j'), k \rangle \mid 1 \leq k \leq j\}$,
- $\mathbf{R}_V^{\mathfrak{G}^*} := \bigcup_{i,j,i',j' \in \mathbb{N}}^{i \leq i'} \{\langle k, (i, j), (i', j'), k \rangle \mid 1 \leq k \leq i\}$,
- $\mathbf{S}_H^{\mathfrak{G}^*} := \bigcup_{i,j,i' \in \mathbb{N}} \{\langle k-1, (i, j), (i', j+1), k \rangle \mid 1 \leq k \leq j+1\}$,
- $\mathbf{S}_V^{\mathfrak{G}^*} := \bigcup_{i,j,j' \in \mathbb{N}} \{\langle k-1, (i, j), (i+1, j'), k \rangle \mid 1 \leq k \leq i+1\}$,
- $\mathbf{C}_H^{\mathfrak{G}^*} := \bigcup_{i,j,i',j' \in \mathbb{N}}^{j \leq j'} \{\langle k, (i', j'), (i, j) \rangle \mid 1 \leq k \leq j\}$, and

$$\bullet \mathfrak{C}_V^{\mathfrak{G}^*} := \bigcup_{i,j,i',j' \in \mathbb{N}}^{i \leq i'} \{ \langle k, (i', j'), (i, j) \rangle \mid 1 \leq k \leq i \}.$$

To understand the assignments above we first introduce some terminology. Given a relation \mathfrak{p} denote its reversal by \mathfrak{p}^{-1} . Now, take \mathfrak{A} to be any structure interpreting a binary predicate \mathfrak{p} . Taking $a \in A$ we define a \mathfrak{p} -edge originating from a in \mathfrak{A} to be some tuple ab , where $\mathfrak{A} \models \mathfrak{p}[ab]$. Similarly, a \mathfrak{p}^{-1} -edge terminating at a in \mathfrak{A} is some tuple ba , where $\mathfrak{A} \models \mathfrak{p}[ab]$. References to \mathfrak{A} are suppressed when the structure is clear from context.

Returning to \mathfrak{G}^* fix some $(i, j), (i', j') \in \mathbb{N} \times \mathbb{N}$ with $j \leq j'$. We defined $\mathfrak{R}_H^{\mathfrak{G}^*}$ in a way that injectively maps \mathfrak{E}_H^{-1} -edges terminating at (i, j) to \mathfrak{E}_H -edges originating from (i', j') . On the other hand, $\mathfrak{S}_H^{\mathfrak{G}^*}$ is defined as a bijection between the \mathfrak{E}_H^{-1} -edges terminating at (i, j) together with the spare part 0 and \mathfrak{E}_H -edges originating from $(i', j+1)$. Lastly, $\mathfrak{C}_H^{\mathfrak{G}^*}$ remembers which \mathfrak{E}_H^{-1} -edges terminating at (i', j') are mapped to \mathfrak{E}_H -edge originating from (i, j) via $(\mathfrak{R}_H^{\mathfrak{G}^*})^{-1}$. (Note the reversal of \mathfrak{R}_H and order of elements). Relations $\mathfrak{R}_V^{\mathfrak{G}^*}$, $\mathfrak{S}_V^{\mathfrak{G}^*}$ and $\mathfrak{C}_V^{\mathfrak{G}^*}$ act similarly.

Lastly, we say that $\mathfrak{G}^\#$ is the ordered expansion of \mathfrak{G}^* , where \mathfrak{G}^* itself is the graphed and mapped expansion of \mathfrak{G} , if the signature contains two additional binary relations \preceq_H and \preceq_V which we will define to be total orders over $\mathbb{N} \times \mathbb{N}$. For aid motivation, we forget that grid elements a and b are pairs of natural numbers and instead focus on the \mathfrak{E}_H - and \mathfrak{E}_V -degrees of the elements. In $\mathfrak{G}^\#$ we have:

- $a \preceq_H^{\mathfrak{G}^\#} b$ if and only if the \mathfrak{E}_H -degree of a is no more than that of b , and
- $a \preceq_V^{\mathfrak{G}^\#} b$ if and only if the \mathfrak{E}_V -degree of a is no more than that of b .

We now define the sentence φ one conjunct at a time:

$$\varphi_1 \wedge \cdots \wedge \varphi_{13}. \tag{\varphi}$$

At a high level, the conjuncts simply state facts about the graphed mapped and ordered expansion $\mathfrak{G}^\#$ of \mathfrak{G} . In the sequel we will argue that the satisfaction of φ by some \mathfrak{A} is sufficient to deduce that \mathfrak{A} is grid-like.

Keeping $\mathfrak{G}^\#$ fixed recall that there is a unique spare part element $0 \in \mathfrak{G}^\#$. Since this element is not part of the grid ($0 \notin \mathfrak{G}^{\mathfrak{G}^\#}$) we have that $\mathfrak{G}^\#$ models:

$$\exists_{[=1]}x \mathfrak{O}(x) \wedge \forall x (\mathfrak{O}(x) \rightarrow \neg \mathfrak{G}(x)). \tag{\varphi_1}$$

Additionally, recall that the spare part 0 has no incoming edges E_H - or E_V -edges in $\mathfrak{G}^\#$. Thus, $\mathfrak{G}^\#$ also models:

$$\forall x \left(\mathbf{0}(x) \rightarrow \forall y (\neg E_H(yx) \wedge \neg E_V(yx)) \right). \quad (\varphi_2)$$

Moving to grid elements, we see that there is a single element in $\mathfrak{G}^{\mathfrak{A}\#}$ with no E_H - or E_V -degree. (This is the starting coordinate $(0, 0)$). Thus, $\mathfrak{G}^\#$ is a model of:

$$\exists_{[=1]} x \left(\mathbf{G}(x) \wedge \forall y (\neg E_H(xy) \wedge \neg E_V(xy)) \right). \quad (\varphi_3)$$

Note that elements in $\mathbb{N} \times \mathbb{N}$ have a single H- and V-successor. Thus, $\mathfrak{G}^\#$ models:

$$\forall x \left(\mathbf{G}(x) \rightarrow \left(\exists_{[=1]} y (\mathbf{H}(xy) \wedge \mathbf{G}(y)) \wedge \exists_{[=1]} y (\mathbf{V}(xy) \wedge \mathbf{G}(y)) \right) \right). \quad (\varphi_4)$$

For the remaining conjuncts let us fix $(i, j) \in \mathbb{N} \times \mathbb{N}$. Notice that the H-successor $(i, j+1)$ has an E_H -degree that is larger by 1 when compared to its predecessor (i, j) . Thus, there is a bijection between the set of E_H^{-1} -edges terminating at (i, j) together with the spare part 0 and the set of E_H -edges originating from $(i, j+1)$. One such bijection is the relation $\mathbf{S}_H^{\mathfrak{G}^\#}$. Noting that V-successors have analogous properties we conclude that $\mathfrak{G}^\#$ models the following sentences:

$$\bigwedge_{x \in \{H, V\}} \forall xyz \left(((E_x(yx) \vee \mathbf{0}(x)) \wedge \mathbf{X}(yz)) \rightarrow \exists_{[=1]} w (E_x(zw) \wedge \mathbf{S}_x(xyzw)) \right), \quad (\varphi_5)$$

$$\bigwedge_{x \in \{H, V\}} \forall wzy \left((E_x(zw) \wedge \mathbf{X}(yz)) \rightarrow \exists_{[=1]} x ((E_x(yx) \vee \mathbf{0}(x)) \wedge \mathbf{S}_x(xyzw)) \right). \quad (\varphi_6)$$

Recall that, by our construction, grid elements $(i, j), (i', j') \in \mathbb{N} \times \mathbb{N}$ satisfy $\mathfrak{G}^\# \models (i, j) \preceq_V (i', j')$ if and only if the E_V -degree of (i, j) is no more than that of (i', j') . Taking (i, j) and its H-successor $(i, j+1)$ we see that $(i, j) \preceq_V^{\mathfrak{G}^\#} (i, j+1)$ and $(i, j+1) \preceq_V^{\mathfrak{G}^\#} (i, j)$. That is, (i, j) and its H-successor have the same E_V -degrees. Analogously, we have that $(i, j) \preceq_H^{\mathfrak{G}^\#} (i+1, j)$ and $(i+1, j) \preceq_H^{\mathfrak{G}^\#} (i, j)$. Thus, $\mathfrak{G}^\#$ is a model of:

$$\bigwedge_{\substack{x \neq y \\ x, y \in \{H, V\}}} \forall xy \left(\mathbf{X}(xy) \rightarrow (x \preceq_y y \wedge y \preceq_y x) \right). \quad (\varphi_7)$$

Since \preceq_X is total on $G^{\mathfrak{G}^\#}$ for both $X \in \{H, V\}$, we have that $\mathfrak{G}^\#$ models:

$$\bigwedge_{X \in \{H, V\}} \forall xy \left((\mathbf{G}(x) \wedge \mathbf{G}(y)) \rightarrow (x \preceq_X y \vee y \preceq_X x) \right). \quad (\varphi_8)$$

Still keeping $(i, j), (i', j') \in \mathbb{N} \times \mathbb{N}$ suppose $(i, j) \preceq_H^{\mathfrak{G}^\#} (i', j')$. By our construction, the E_H^{-1} -edges terminating at (i, j) are injectively mapped to E_H -edges originating from (i', j') . We write down this implication in the form of φ_9 and φ_{10} .

First, we capture functionality of R_H ; i.e. we will require E_H^{-1} -edges terminating at (i, j) to be mapped (via R_H) to a single E_H -edge originating from (i', j') :

$$\bigwedge_{X \in \{H, V\}} \forall xyz \left((y \preceq_X z \wedge E_X(yx)) \rightarrow \exists_{[=1]} w (E_X(zw) \wedge R_X(xyzw)) \right). \quad (\varphi_9)$$

With the next sentence we secure injectivity. I.e. we write that E_H -edges originating from (i', j') are mapped (via R_H) to at most one E_H^{-1} -edge terminating at (i, j) :

$$\bigwedge_{X \in \{H, V\}} \forall wzy \left((y \preceq_X z \wedge E_X(zw)) \rightarrow \exists_{[\leq 1]} x (E_X(yx) \wedge R_X(xyzw)) \right). \quad (\varphi_{10})$$

It is easy to verify that this is indeed how $R_H^{\mathfrak{G}^\#}$ is set up. Noting that \preceq_V and R_V behave symmetrically we conclude $\mathfrak{G}^\# \models \varphi_9 \wedge \varphi_{10}$.

Let us continue with the assumption $(i, j) \preceq_H^{\mathfrak{G}^\#} (i', j')$. Notice that (i, j) and (i', j') are equal in regards to \preceq_H in case R_H is a bijection between E_H^{-1} -edges terminating at (i, j) and E_H -edges originating from (i', j') . We write this implication in the form of sentences φ_{11} and φ_{12} .

By the assumption $(i, j) \preceq_H^{\mathfrak{G}^\#} (i', j')$ we will already have that R_H injectively maps E_H^{-1} -edges terminating at (i, j) to E_H -edges originating from (i', j') . Recall that, by our construction, $\langle k, (i', j'), (i, j) \rangle \in C_H^{\mathfrak{G}^\#}$ if and only if there is some $k' \in \mathbb{N}$ for which $\langle k', (i, j), (i', j'), k \rangle \in R_H^{\mathfrak{G}^\#}$. In other words, C_H remembers which E_H^{-1} -edges terminating at (i', j') are featured in a mapping (by R_H) with E_H -edges originating from (i, j) . We axiomatise this relationship as follows:

$$\bigwedge_{X \in \{H, V\}} \forall wzy \left(C_X(wzy) \leftrightarrow \exists_{[=1]} x (E_X(yx) \wedge R_X(xyzw)) \right). \quad (\varphi_{11})$$

Utilising C_H we can then test if all E_H^{-1} -edges terminating at (i', j') are mapped to some E_H -edge originating from (i, j) via R_E^{-1} . Clearly, if the property holds, then

$(i', j') \preceq_{\mathfrak{H}}^{\mathfrak{G}^{\#}} (i, j)$. We write this implication as follows:

$$\bigwedge_{x=H,V} \forall yz \left((G(y) \wedge G(z) \wedge \forall w (\mathbf{E}_x(zw) \rightarrow \mathbf{C}_x(wzy))) \rightarrow z \preceq_x y \right). \quad (\varphi_{12})$$

Noting that \preceq_V , R_V and \mathbf{C}_V behave similarly, we have that $\mathfrak{G}^{\#} \models \varphi_{11} \wedge \varphi_{12}$.

Lastly, notice that there are no two grid elements that have the same \mathbf{E}_H - and \mathbf{E}_V -degrees. Thus, $\mathfrak{G}^{\#}$ models the uniqueness requirement as given by the following:

$$\forall y \left(\mathbf{G}(y) \rightarrow \exists_{[=1]} z \left(\bigwedge_{x=H,V} (y \preceq_x z \wedge z \preceq_x y) \right) \right). \quad (\varphi_{13})$$

Let us now briefly step inside the realm of periodic counting. We may now capture the fact that, in $\mathfrak{G}^{\#}$, there are no transfinite positions by defining the sentence χ limiting the \mathbf{E}_H - and \mathbf{E}_V -degrees of elements to finite values:

$$\bigwedge_{x=H,V} \forall x \exists_{[0+1]} y \mathbf{E}_x(xy). \quad (\chi)$$

Recalling that $\varphi = \varphi_1 \wedge \dots \wedge \varphi_{13}$ we have showed the following:

Lemma 6.2.2. *The graphed, mapped and ordered expansion of the canonical $(\mathbb{N} \times \mathbb{N})$ -grid is a model of $\varphi \wedge \chi$.*

Lemma 6.2.3. *Suppose $\mathfrak{A} \models \varphi$. Then \mathfrak{A} is a grid-like structure. In addition, if $\mathfrak{A} \models \chi$, then \mathfrak{A} is an $(\mathbb{N} \times \mathbb{N})$ -grid.*

Proof. We start by showing the first statement. Notice that by φ_1 there is exactly one element that satisfies $\mathbf{0}$ in \mathfrak{A} and, by φ_2 , has no incoming \mathbf{E}_H - and \mathbf{E}_V -edges. This will be our spare part element in the argument to come. Now, take any element $a_{0,0} \in A$ such that $a_{0,0} \in \mathbf{G}^{\mathfrak{A}}$ (i.e. $a_{0,0}$ is a grid element) with finite \mathbf{E}_H - and \mathbf{E}_V -degree. Such an element is guaranteed to exist by φ_3 . Then, φ_4 tells us that $a_{0,0}$ has an H-successor $a_{0,1}$ and a V-successor $a_{1,0}$. Notice that, by φ_5 , each \mathbf{E}_H^{-1} -edge terminating at $a_{0,0}$ along with the spare part is paired with exactly one \mathbf{E}_H -edge originating from $a_{0,1}$ in $\mathbf{S}_H^{\mathfrak{A}}$. That is to say, writing $U = \{b \in A \mid a_{0,0}b \in \mathbf{E}_H^{\mathfrak{A}} \text{ or } b \in \mathbf{0}^{\mathfrak{A}}\}$ and $U' = \{c \in A \mid a_{0,1}c \in \mathbf{E}_H^{\mathfrak{A}}\}$, we have that for each $b \in U$ there is exactly one $c \in U'$ such that $ba_{0,0}a_{0,1}c \in \mathbf{S}_H^{\mathfrak{A}}$. The reverse is established by φ_6 . Clearly, there is a bijection between U and U' thus making the \mathbf{E}_H -degree of $a_{0,1}$ one greater than that of $a_{0,0}$. By φ_7 we have that $a_{0,0} \preceq_V^{\mathfrak{A}} a_{0,1}$ and $a_{0,1} \preceq_V^{\mathfrak{A}} a_{0,0}$. We first fixate on the fact that $a_{0,0} \preceq_V^{\mathfrak{A}} a_{0,1}$. Writing

$U = \{b \in A \mid a_{0,0}b \in E_V^{\mathfrak{A}}\}$ and $U' = \{c \in A \mid a_{0,1}c \in E_V^{\mathfrak{A}}\}$ we have, by φ_9 , that for each $b \in U$ there is exactly one $c \in U'$ such that $ba_{0,0}a_{0,1}c \in R_V^{\mathfrak{A}}$. By φ_{10} , for each $c \in U'$ there is at most a single $b \in U$ such that $ba_{0,0}a_{0,1}c \in R_V^{\mathfrak{A}}$. We may thus regard $R_V^{\mathfrak{A}}$ as an injective map between E_V^{-1} -edges terminating at $a_{0,0}$ and E_V -edges originating from $a_{1,0}$. Then, again by φ_9 , φ_{10} and the fact that $a_{0,1} \preceq_V^{\mathfrak{A}} a_{0,0}$, we have that $R_V^{\mathfrak{A}}$ injectively maps the E_V^{-1} -edges terminating at $a_{0,1}$ to E_V -edges originating from $a_{0,0}$. By the Schröder-Bernstein theorem, there is a bijection between U and U' thus ensuring that the E_V -degrees of the elements in question coincide. A symmetric argument holds for the E_V - and E_H -degree of $a_{1,0}$.

Now, let $a_{1,1}$ and $a'_{1,1}$ be, respectively, the V -successor of $a_{0,1}$ and the H -successor of $a_{1,0}$ promised by φ_4 . Using the same arguments as in the paragraph above, it is easy to see that the E_H -degrees of $a_{1,1}$ and $a'_{1,1}$ coincide; as do E_V -degrees. We claim that $a_{1,1} = a'_{1,1}$. By φ_{13} we need only show that $a_{1,1} \preceq_X^{\mathfrak{A}} a'_{1,1}$ and $a'_{1,1} \preceq_X^{\mathfrak{A}} a_{1,1}$ for both $X \in \{H, V\}$. Fixating on E_H -edges first, we have, by φ_8 , that $a_{1,1}$ and $a'_{1,1}$ are comparable by $\preceq_H^{\mathfrak{A}}$ in some way. Suppose, without loss of generality, that $a_{1,1} \preceq_H^{\mathfrak{A}} a'_{1,1}$. Writing $U = \{b \in A \mid a_{1,1}b \in E_H^{\mathfrak{A}}\}$ and $U' = \{c \in A \mid a'_{1,1}c \in E_H^{\mathfrak{A}}\}$ we have, by φ_9 , that for each $b \in U$ there is exactly one $c \in U'$ such that $ba_{1,1}a'_{1,1}c \in R_H^{\mathfrak{A}}$, and, by φ_{10} , for each $c \in U'$ there is at most one $b \in U$ such that the same holds. That is to say, $R_H^{\mathfrak{A}}$ injectively maps E_H^{-1} -edges terminating at $a_{1,1}$ to E_H -edges originating from $a'_{1,1}$. Since $a_{1,1}$ and $a'_{1,1}$ both have an equal and finite E_H -degree, we conclude that $R_H^{\mathfrak{A}}$ is a bijection between the edges. Hence, we have, by φ_{11} , that $ca'_{1,1}a_{1,1} \in C_H^{\mathfrak{A}}$ for each $c \in U'$. Clearly, the antecedents of φ_{12} are met and thus $a'_{1,1} \preceq_H^{\mathfrak{A}} a_{1,1}$ as required. Repeating the argument for $\preceq_V^{\mathfrak{A}}$ we conclude that $a_{1,1} \preceq_X^{\mathfrak{A}} a'_{1,1}$ and $a'_{1,1} \preceq_X^{\mathfrak{A}} a_{1,1}$ for both $X \in \{H, V\}$ thus closing the grid.

By repeating the argument above on every element in $G^{\mathfrak{A}}$ with finite E_H - and E_V -degree we conclude that \mathfrak{A} contains a homomorphic embedding of the canonical $(\mathbb{N} \times \mathbb{N})$ -grid thus making it grid-like.

Suppose now, in addition, that $\mathfrak{A} \models \chi$. We have that each element in $G^{\mathfrak{A}}$ has a finite E_H - and E_V -degree. We may thus unambiguously identify these elements as the pair of their E_H -degree $j \in \mathbb{N}$ and E_V -degree $i \in \mathbb{N}$. Hence, the structure \mathfrak{A} restricted to elements in $G^{\mathfrak{A}}$ and signature $\{G, H, V\}$ is isomorphic to the canonical $(\mathbb{N} \times \mathbb{N})$ -grid thus making \mathfrak{A} an $(\mathbb{N} \times \mathbb{N})$ -grid as required. \square

Taking Lemmas 6.2.2 and 6.2.3 we see that φ is an \mathcal{AFC}^4 -sentence satisfied exclusively by grid-like structures, whilst the \mathcal{AFPC}^4 -sentence $\varphi \wedge \chi$ only has $(\mathbb{N} \times \mathbb{N})$ -grids as models. The subclasses of grid-like structures and $(\mathbb{N} \times \mathbb{N})$ -grids

the above sentences define are sufficient for entailing undecidability as in Proposition 6.2.1. Now, recall that \mathcal{AFC} is a fragment of \mathcal{FO} . By completeness of first-order logic we have that satisfiability of \mathcal{AFC} resides in Π_1^0 . On the other hand we have that each \mathcal{AFPC} -sentence is logically equivalent to a computable formula of $\mathcal{L}_{\omega_1, \omega}$. By *Barwise* completeness [Barwise, 1967]¹, we have that the satisfiability problem of \mathcal{AFPC} is in Σ_1^1 . Hence, we proved:

Theorem 6.2.4. *The satisfiability problem for \mathcal{AFC}^4 is Π_1^0 -complete. The same problem for \mathcal{AFPC}^4 is Σ_1^1 -complete.*

¹We recommend the survey paper [Keisler and Knight, 2004] as a starting point.

Chapter 7

The 3-variable Adjacent Fragment with Counting

The following chapter is unpublished work of the Ph. d. candidate.

In this section we show that the status of the satisfiability problem for the adjacent fragment with counting is not as bleak as inferred previously. To do so, let us confine ourselves to the \mathcal{AFC}^3 ; i.e. the three-variable adjacent fragment with counting. Note that, whilst we have established Σ_1^0 -completeness for $\text{FinSat}(\mathcal{AFC}^3)$ in Theorem 6.1.2, the subject of general satisfiability for \mathcal{AFC}^3 has been left untouched in the previous section. Indeed, the construction leading up to Theorem 6.2.4 operates on 4 variables. As will become apparent in the sequel, this is not a weakness of our undecidability reduction but rather a consequence of $\text{Sat}(\mathcal{AFC}^3)$ being decidable.

We will briefly illustrate the capabilities of \mathcal{AFC}^3 . To this end we introduce the notion of *Härtig quantification*. Let \mathfrak{A} be any structure, $\bar{a} \in A^\ell$ and $\gamma(\mathbf{x}_\ell x)$, $\delta(\mathbf{x}_\ell y)$ any first-order formulas. The non-first-order Härtig quantifier $I(x, y)(\gamma, \delta)$ then behaves as follows:

$$\mathfrak{A}, \bar{a} \models I(x, y)(\gamma, \delta) \text{ if and only if } |\{b \in A \mid \mathfrak{A} \models \gamma[\bar{a}b]\}| = |\{c \in A \mid \mathfrak{A} \models \delta[\bar{a}c]\}|.$$

In the language of graphs $\{\mathbf{E}\}$, we may, using Härtig quantification and at most two-variables, describe the class of directed graphs that have nodes with equal in- and out-degrees as follows:

$$\forall x I(y, y)(\mathbf{E}(yx), \mathbf{E}(xy)). \quad (\varphi)$$

Clearly, such expressivity is outside the realm of \mathcal{C}^2 . It is, however, possible to mimic this behaviour in \mathcal{AFC}^3 by using an auxiliary ternary predicate \mathbf{R} :

$$\begin{aligned} & \forall x_1 x_2 \left(\mathbf{E}(x_1 x_2) \rightarrow \exists_{[=1]} x_3 \left(\mathbf{E}(x_2 x_3) \wedge \mathbf{R}(x_1 x_2 x_3) \right) \right) \wedge \\ & \forall x_3 x_2 \left(\mathbf{E}(x_2 x_3) \rightarrow \exists_{[=1]} x_1 \left(\mathbf{E}(x_1 x_2) \wedge \mathbf{R}(x_1 x_2 x_3) \right) \right). \end{aligned} \quad (\psi)$$

It should be clear that $\psi \models \varphi$ as, when given $\mathfrak{A} \models \psi$, we have that $\mathbf{R}^{\mathfrak{A}}$ is a bijection between incoming and outgoing edges for each given node in the domain. Note, however, that there are properties expressible in two-variable logic with Härtig quantification that are impossible to express (or even simulate with the help of additional predicates) using \mathcal{AFC}^3 . Indeed, it is easy to see that the formula $\neg \mathbf{I}(y, y)(x \neq y, \top)$ is only satisfied by finite structures and thus (by an easy compactness argument) cannot be simulated by any fragment of \mathcal{FO} .

Nonetheless, we have identified the two-variable fragment with Härtig quantification as being “close enough” in expressive power to \mathcal{AFC}^3 to be considered as a target of satisfiability-preserving reductions.

7.1 Limiting Härtig Quantification

The two-variable fragment with Härtig quantifiers, at its full expressivity, has a Σ_1^1 -hard satisfiability problem [Grädel et al., 1999]. The crux of the above result is that one can, using a Härtig quantifier, quantify the variables x and y at the same time. Thus, the following 2-variable formula (in the language of graphs) $\psi(xy) := \mathbf{I}(y, x)(\mathbf{E}(xy), \mathbf{E}(yx))$ allows us to equate out-degrees of any two nodes; similarly as was done above Theorem 6.2.4. Note that this method is not the only way of showing undecidability. Let us assume that Härtig quantifiers can only quantify the variable y (thus preventing us from writing ψ). In the sequel, when given an \mathcal{AFC}^3 -sentence such as φ , we recursively compute a sentence ψ in two-variable logic with Härtig quantifiers that is satisfiable over the same domains as φ . But then, by Theorem 6.1.2, we will have that the satisfiability problem for the two-variable fragment with Härtig quantifiers is Σ_1^0 -hard as:

$$\varphi \in \text{FinSat}(\mathcal{AFC}^3) \Leftrightarrow \psi \wedge \exists x \neg \mathbf{I}(y, y)(x \neq y, \top) \text{ is satisfiable.}$$

To evade undecidability we confine ourselves to a limited form of Härtig quantification we call *uniform Härtig assertions*, which are formulas of the form:

$$\forall \mathbf{x}_n I(y, y)(\gamma, \delta) \text{ with } \gamma, \delta \text{ quantifier-free.}$$

Notice that in the above only the variable y is allowed to be quantified via I . Denote the set of uniform Härtig assertions as \mathcal{UHA} , and let \mathcal{UHA}^2 be the subset of \mathcal{UHA} that features sentences only in the variables x and y . Without loss of generality, we may assume that each formula in \mathcal{UHA}^2 takes the form $\forall x I(y, y)(\gamma, \delta)$ with γ, δ being some quantifier-free \mathcal{FO}^2 -formulas. Now, take any \mathcal{C}^2 -sentence φ and a finite set of \mathcal{UHA}^2 -formulas Φ . In the sequel we show that checking the satisfiability status of formulas such as $\varphi \wedge \bigwedge \Phi$ is decidable.

Still keeping $\varphi \in \mathcal{C}^2$ and $\Phi \subseteq \mathcal{UHA}^2$ finite, we will say that sentences of the form $\varphi \wedge \bigwedge \Phi$ are $\mathcal{C}_{\mathcal{UHA}}^2$ -sentences. A $\mathcal{C}_{\mathcal{UHA}}^2$ -sentence is in normal-form if it takes the following shape:

$$\forall xy \alpha \wedge \bigwedge_{s \in S} \forall x \exists_{[=M_s]} y \beta_s \wedge \bigwedge_{t \in T} \forall x I(y, y)(\gamma_t, \delta_t), \quad (\mathcal{C}_{\mathcal{UHA}}^2\text{-nmf})$$

where M_s is a positive integer, and $\alpha, \beta_s, \gamma_t, \delta_t$ are quantifier-free \mathcal{FO}^2 -formulas indexed by finite sets S and T . The following is immediate using standard rewriting techniques for \mathcal{C}^2 -formulas (see [Pratt-Hartmann, 2023, Lemma 8.3] for \mathcal{C}^2 and Lemma 7.3.2 for a generalisation to \mathcal{AFC}^3):

Lemma 7.1.1. *Suppose that φ is a $\mathcal{C}_{\mathcal{UHA}}^2$ -sentence. Then, we may compute, in polynomial time, a $\mathcal{C}_{\mathcal{UHA}}^2$ -sentence ψ in normal-form such that φ and ψ are satisfiable over the same domains of size at least $M_{\max}+1$, where M_{\max} is the maximum numeric subscript (of a counting quantifier) occurring in φ .*

We note that, whilst addition of uniform Härtig assertions substantially increases the expressive power of \mathcal{C}^2 , the properties definable in $\mathcal{C}_{\mathcal{UHA}}^2$ can still be simulated in \mathcal{FO} and, in fact, \mathcal{AFC}^3 :

Lemma 7.1.2. *Suppose φ is a normal-form $\mathcal{C}_{\mathcal{UHA}}^2$ -sentence. Then, we may compute, in polynomial time, an \mathcal{AFC}^3 -sentence ψ that is satisfiable over the same domains as φ .*

Proof. Recall that φ is of the form ($\mathcal{C}_{\mathcal{UHA}}^2$ -nmf) recapitulated below:

$$\forall xy \alpha \wedge \bigwedge_{s \in S} \forall x \exists_{[=M_s]} y \beta_s \wedge \bigwedge_{t \in T} \forall x I(y, y)(\gamma_t, \delta_t).$$

The promised sentence ψ will be a conjunction of two formulas, the first being:

$$\forall x_1 x_2 \alpha(x_1 x_2) \wedge \bigwedge_{s \in S} \forall x_1 \exists_{[=M_s]} x_2 \beta(x_1 x_2). \quad (\psi_1)$$

We define the second formula, ψ_2 , as follows. Fix some $t \in T$ and let \mathbf{R}_t be a new ternary predicate. We write $\psi_{2,t}$ for the conjunction of the following:

$$\begin{aligned} & \forall x_1 x_2 \left(\gamma_t(x_2 x_1) \rightarrow \exists_{[=1]} x_3 (\delta_t(x_2 x_3) \wedge \mathbf{R}_t(x_1 x_2 x_3)) \right), \\ & \forall x_3 x_2 \left(\delta_t(x_2 x_3) \rightarrow \exists_{[=1]} x_1 (\gamma_t(x_2 x_1) \wedge \mathbf{R}_t(x_1 x_2 x_3)) \right). \end{aligned}$$

After renaming the variable sequences in the second conjunct of $\psi_{2,t}$, and, repeating the computation for each $t \in T$, we will have $\psi_2 := \bigwedge_{t \in T} \psi_{2,t}$.

To verify that φ and ψ are equisatisfiable, proceed as follows. For the easy direction, suppose $\mathfrak{A} \models \psi$. Clearly, the universal and existential requirements of φ are met as $\mathfrak{A} \models \psi_1$. We need only show that $\mathfrak{A} \models \forall x I(y, y)(\gamma_t, \delta_t)$ for each $t \in T$. To this end, fix any $t \in T$ and take some element $b \in A$. Let us write $U_\gamma := \{a \in A \mid \mathfrak{A} \models \gamma_t[ba]\}$ and $V_\delta := \{c \in A \mid \mathfrak{A} \models \delta_t[bc]\}$. Since $\mathfrak{A} \models \psi_{2,t}$ we have that for each $a \in U_\gamma$ there is exactly one $c \in V_\delta$ such that $\mathfrak{A} \models \mathbf{R}_t[abc]$. Conversely, for each $c \in V_\delta$ there is exactly one $a \in U_\gamma$ such that $\mathfrak{A} \models \mathbf{R}_t[abc]$. Clearly, the relation $\{\langle a, c \rangle \in U_\gamma \times V_\delta \mid \mathfrak{A} \models \mathbf{R}_t[abc]\}$ is a bijection between γ_t - and δ_t -witnesses of b . Thus, $|U_\gamma| = |V_\delta|$ and $\mathfrak{A}, b \models I(y, y)(\gamma_t, \delta_t)$ as required.

Conversely, suppose $\mathfrak{A} \models \varphi$. We will expand \mathfrak{A} into a model $\mathfrak{B} \models \psi$ by providing extensions to predicates \mathbf{R}_t for each $t \in T$. Firstly, set $\mathfrak{B} := \mathfrak{A}$. Thus, $\mathfrak{B} \models \psi_1$. To enforce $\mathfrak{B} \models \psi_2$ fix some $t \in T$ and $b \in A$. Define $U_\gamma := \{a \in A \mid \mathfrak{A} \models \gamma_t[ba]\}$ and $V_\delta := \{c \in A \mid \mathfrak{A} \models \delta_t[bc]\}$. By the assertion $\mathfrak{A}, b \models I(y, y)(\gamma_t, \delta_t)$, we have that the sets are of the same cardinality. We may thus define $f : U_\gamma \rightarrow V_\delta$ to be some bijection and set $\mathfrak{B} \models \mathbf{R}_t[abf(a)]$ for all $a \in U_b$. This clearly has the required effect. Repeating the procedure for all $t \in T$ and $b \in A$ we will have the required result. \square

Let Σ be any constant- and function-free signature. Recall that an atomic 1- and 2-type over Σ is a maximal consistent set of literals formed from $\Sigma \cup \{=\}$ in,

respectively, 1- and 2-variables. Given a structure \mathfrak{A} and $i \in \{1, 2\}$, each $\bar{a} \in A^i$ realises a unique atomic i -type which we denote by $\text{tp}^{\mathfrak{A}}[\bar{a}]$. We will write ATP_i^Σ for the set of all atomic i -types over Σ . In the 2-variable setting, a *profile* is a function mapping atomic 2-types to cardinal numbers. Given an element $a \in A$ we define the profile ρ of a in \mathfrak{A} to be defined on atomic 2-types ξ as follows:

$$\rho(\xi) := |\{b \in A \mid \mathfrak{A} \models \xi[ab]\}|.$$

For brevity, we write $\text{pr}^{\mathfrak{A}}[a]$ for the profile of a in \mathfrak{A} . Given quantifier-free $\mathcal{C}_{\mathcal{UHA}}^2$ -formulas ψ, θ over the signature Σ , a positive integer M , and a profile ρ , we write:

1. $\rho \models \forall y \psi$ whenever $\sum_{\xi \in \text{ATP}_2^\Sigma}^{\xi \not\models \psi} \rho(\xi) = 0$,
2. $\rho \models \exists_{[=M]} y \theta$ whenever $\sum_{\xi \in \text{ATP}_2^\Sigma}^{\xi \models \theta} \rho(\xi) = M$, and
3. $\rho \models \text{I}(y, y)(\psi, \delta)$ whenever $\sum_{\xi \in \text{ATP}_2^\Sigma}^{\xi \models \psi} \rho(\xi) = \sum_{\mu \in \text{ATP}_2^\Sigma}^{\mu \models \delta} \rho(\mu)$.

Taking any structure \mathfrak{A} and $a \in A$ it should be clear that, for formulas χ taking any of the forms $\forall y \psi$, $\exists_{[=M]} y \theta$, or $\text{I}(y, y)(\psi, \delta)$, we have $\mathfrak{A} \models \chi[a]$ if and only if $\text{pr}^{\mathfrak{A}}[a] \models \chi$. Supposing that φ is a normal-form $\mathcal{C}_{\mathcal{UHA}}^2$ -sentence, we say that a profile ρ is φ -compliant just in case the following hold:

1. $\rho \models \forall y \alpha$,
2. $\rho \models \exists_{[=M_s]} y \beta_s$ for each $s \in S$,
3. $\rho \models \text{I}(y, y)(\gamma_t, \delta_t)$ for each $t \in T$,
4. $\rho(\xi) = 1$ for exactly one 2-type ξ featuring the literal $x = y$, and
5. $\rho(\xi) \leq 1$ for all 2-type ξ featuring the literal $x = y$.

Say that a profile ρ is realised in a structure \mathfrak{A} just in case there is some $a \in A$ having $\text{pr}^{\mathfrak{A}}[a] = \rho$. Referencing φ -compliance, the following is immediate:

Lemma 7.1.3. *Take a normal-form $\mathcal{C}_{\mathcal{UHA}}^2$ -sentence φ and a structure \mathfrak{A} . Then, $\mathfrak{A} \models \varphi$ if and only if each profile realised in \mathfrak{A} is φ -compliant.*

7.2 Deciding $\text{Sat}(\mathcal{C}_{\mathcal{U}\mathcal{H}\mathcal{A}}^2)$

Let us fix a satisfiable normal-form $\mathcal{C}_{\mathcal{U}\mathcal{H}\mathcal{A}}^2$ -sentence φ recapitulated here:

$$\forall xy \alpha \wedge \bigwedge_{s \in S} \forall x \exists_{[=M_s]} y \beta_s \wedge \bigwedge_{t \in T} \forall x I(y, y)(\gamma_t, \delta_t).$$

Now, suppose that \mathfrak{A} is a countable model of φ . This comes with no loss of generality as, by Lemma 7.1.2, there is an \mathcal{AFC}^3 -formula ψ satisfiable over the same domains as φ . Since ψ is in \mathcal{FO} , by the downward Löwenheim-Skolem theorem it has a countable model; and so does φ .

Now, take some element $a \in A$ and let ρ be its profile in \mathfrak{A} . We define the *character of ρ* , denoted $\text{chr}(\rho)$, to be a profile-like function χ mapping 2-types to the 3-element set¹ $\{0, 1^{+1}, \aleph_0\}$. Formally, if ξ is an 2-type, then

$$\chi(\xi) := \begin{cases} 0 & \text{if } \rho(\xi) = 0, \\ 1^{+1} & \text{if } \rho(\xi) \in \mathbb{N} \setminus \{0\}, \\ \aleph_0 & \text{otherwise.} \end{cases}$$

For convenience, we define the *character of a in \mathfrak{A}* , denoted as $\text{chr}^{\mathfrak{A}}[a]$, to simply be the character of $\text{pr}^{\mathfrak{A}}[a]$. We say that a character χ is realised in \mathfrak{A} if there is some $a \in A$ for which $\text{chr}^{\mathfrak{A}}[a] = \chi$.

The ensuing decidability argument depends on the observation that there is a model of φ in which elements of the same character realise a finite number of different profiles. Let us take some character χ realised in \mathfrak{A} and a 2-type ξ . Additionally suppose that $\chi(\xi) = 1^{+1}$. We say that ξ is *converging* for χ in \mathfrak{A} if there is some $n \in \mathbb{N}$ such that $\rho(\xi) \leq n$ for each profile ρ of character χ realised in \mathfrak{A} . If no such $n \in \mathbb{N}$ exists, we say that ξ is *diverging* for χ in \mathfrak{A} . Alternatively, divergence of ξ for χ can be viewed as the absence of an upper-bound for the following subset of natural numbers: $\{\rho(\xi) \mid \rho \text{ is realised in } \mathfrak{A} \text{ and } \text{chr}(\rho) = \chi\}$. Convergence and divergence is left undefined when $\chi(\xi) = 0$ or $\chi(\xi) = \aleph_0$. Note the following implication of divergence:

Lemma 7.2.1. *Pick some uniform Härtig assertion $t \in T$ in φ and some character χ realised in \mathfrak{A} . If some 2-type $\xi \models \gamma_t$ is diverging for χ in \mathfrak{A} , then there is some 2-type $\mu \models \delta_t$ that is either diverging for χ in \mathfrak{A} or such that $\chi(\mu) = \aleph_0$.*

¹Here 1^{+1} is a single element representing cardinalities that are positive but not infinite.

Proof. Suppose that $\mathfrak{A} \models \varphi$ realises some character χ for which the 2-type $\xi \models \gamma_t$ is diverging. Let S be the set of all 2-types μ such that $\mu \models \delta_t$ and $\chi(\mu) \neq 0$. For contradiction suppose that every 2-type μ in V is converging for χ in \mathfrak{A} . (By definition of convergence, $\chi(\mu) \neq \aleph_0$). Thus, by convergence, let $n \in \mathbb{N}$ be the smallest integer such that $\rho(\mu) \leq n$ for each ρ in \mathfrak{A} of character χ and all $\mu \in V$. On the other hand, since ξ is diverging for χ , we may find an element a with $\rho := \text{pr}^{\mathfrak{A}}[a]$ such that $\rho(\xi) \geq n|V| + 1$. But $\sum_{\mu \in S} \rho(\mu) \leq n|S|$ thus contradicting our initial assumption that $\mathfrak{A}, a \models I(y, y)(\gamma_t, \delta_t)$. \square

Note that the lemma above holds with γ_t and δ_t transposed. Keeping the character χ fixed, we define $\lceil \chi \rceil^{\mathfrak{A}}$ to be a character defined over 2-types ξ as follows:

$$\lceil \chi \rceil^{\mathfrak{A}}(\xi) := \begin{cases} \aleph_0 & \text{if } \xi \text{ is diverging for } \chi \text{ in } \mathfrak{A}, \\ \chi(\xi) & \text{otherwise.} \end{cases}$$

That is, we view $\lceil \chi \rceil^{\mathfrak{A}}$ as being the profile χ but with the cardinalities of diverging 2-types being ‘‘rounded up’’ to \aleph_0 . Similarly, when given a profile ρ we write $\lceil \rho \rceil^{\mathfrak{A}}$ for the function defined on 2-types ξ as follows:

$$\lceil \rho \rceil^{\mathfrak{A}}(\xi) := \begin{cases} \aleph_0 & \text{if } \xi \text{ is diverging for } \text{chr}(\rho) \text{ in } \mathfrak{A}, \\ \rho(\xi) & \text{otherwise.} \end{cases}$$

Given Lemma 7.2.1, the following is almost immediate:

Lemma 7.2.2. *Take some profile ρ realised in \mathfrak{A} . Then, $\lceil \rho \rceil^{\mathfrak{A}}$ is φ -compliant.*

Proof. Since $\rho(\xi) = 0$ implies $\lceil \rho \rceil^{\mathfrak{A}}(\xi) = 0$, we have that the universal conjuncts of φ are not violated by any 2-type ξ satisfying $\lceil \rho \rceil^{\mathfrak{A}}(\xi) \geq 1$. Thus, $\lceil \rho \rceil^{\mathfrak{A}} \models \forall y \alpha(xy)$.

Now, write $\chi = \text{chr}(\rho)$ and take any $s \in S$. We have that no 2-type $\xi \models \beta_s$ can be diverging for χ as otherwise there is some element $a \in A$ with profile ρ^* of character χ that has $\rho^*(\xi) \geq M_s + 1$ thus contradicting our assumption that $\mathfrak{A}, a \models \exists_{[=M_s]} y \beta_s$. We must thus have $\lceil \rho \rceil^{\mathfrak{A}}(\xi) = \rho(\xi)$ for each 2-type $\xi \models \beta_s$ allowing us to conclude $\lceil \rho \rceil^{\mathfrak{A}} \models \exists_{[=M_s]} y \beta_s(xy)$.

Lastly, take some $t \in T$ and let U be the set of all 2-type $\xi \models \gamma_t$, whilst making V the set of all 2-types $\mu \models \delta_t$. If no 2-type amongst $U \cup V$ is diverging for χ , then there is nothing to show as $\lceil \rho \rceil^{\mathfrak{A}}(\xi) = \rho(\xi)$. If there is some diverging ξ for χ in say U , then, by Lemma 7.2.1, there is some $\mu \in V$ that is either also diverging

for χ or having $\chi(\mu) = \rho(\mu) = \aleph_0$. Either way we have that $\lceil \bar{\rho} \rceil(\xi) = \lceil \bar{\rho} \rceil(\mu) = \aleph_0$ thus securing $\lceil \bar{\rho} \rceil \models I(y, y)(\gamma_t, \delta_t)$. \square

Utilising the results above, we show that if a normal-form sentence φ is satisfiable, then it is satisfiable in a structure \mathfrak{B} without diverging characters; that is to say, $\chi = \lceil \bar{\chi} \rceil$ for each character χ realised in \mathfrak{B} . The following technical lemma will be useful in the main part of our argument.

Lemma 7.2.3. *Suppose ξ is diverging for χ and write $B := \{b \in A \mid \text{chr}^{\mathfrak{A}}[b] = \chi\}$. Then, there is an infinite set $C \subseteq A$ and an injection $f : C \rightarrow B$ satisfying:*

- $\text{tp}^{\mathfrak{A}}[f(c)c] = \xi$ for each $c \in C$,
- f is asymmetric (i.e. $f(c) = d \Rightarrow f(d) \neq c$ for each $c, d \in C$).

Proof. Let $D := \{d \in A \mid \text{atp}^{\mathfrak{A}}[bd] = \xi \text{ for some } b \in B\}$. That is, $d \in A$ is included in D just in case it is a ξ -witness for some element with the character χ . By divergence of ξ for χ , we have that $|D| = \aleph_0$. We proceed by inductively defining a series of sets $B_i \subsetneq B$, $C_i \subsetneq D$ and functions $f_i : C_i \rightarrow B_i$ for $i < \omega$. The limit set $C := \bigcup_{i < \omega} C_i$ and function $f := \bigcup_{i < \omega} f_i$ will then be as required by the lemma. We maintain the following for $i < \omega$:

0. $B_j \subsetneq B_i$ for each $j < i$ with $|B_i| = i$,
1. $C_j \subsetneq C_i$ for each $j < i$ with $|C_i| = i$,
2. $f_j \subsetneq f_i$ for each $j < i$, with $f_i : C_i \rightarrow B_i$ injective,
3. $\text{tp}^{\mathfrak{A}}[f_i(c)c] = \xi$ for each $c \in C_i$, and
4. f_i is asymmetric.

We proceed by induction on the elements of B . Intuitively, B_i is simply the set of processed elements.

For the base case we set $B_0 = C_0 = f_0 = \emptyset$. Clearly, properties 0–4 are met by this assignment. At step $1 \leq i < \omega$ we take C_j and f_j to be defined compatibly with 0–4 for each $j < i$. Moving to the construction, let $B' := B \setminus B_{i-1}$. Since B_{i-1} is of cardinality $i-1$, we have that B' is infinite. Notice that, by divergence of ξ for χ in \mathfrak{A} , there is some element $b \in B'$ realising a profile ρ , such that $\rho(\xi) \geq |C_{i-1}| + 2$. Thus, by setting $D_b := \{d \in D \mid \mathfrak{A} \models \xi[bd]\}$ we can find at

least two distinct elements $d_1, d_2 \in D_b$ such that $d_1, d_2 \notin C_{i-1}$. It should be clear that $b \notin D_b$ as otherwise $x = y$ is a literal of ξ thus making $|D_b| = 1$. Since $B \cap C_{i-1}$ is not necessarily empty, we might have that b is a member of C_{i-1} and even that $f_{i-1}(b) = d_k$ for some $k \in \{1, 2\}$. If this is indeed the case, set $d := d_{3-k}$. Otherwise, pick d to be d_1 or d_2 arbitrarily. Setting $B_i := B_{i-1} \cup \{b\}$, and $C_i := C_{i-1} \cup \{d\}$, and $f_i := f_{i-1} \cup \{d \mapsto b\}$ we will have maintained the required properties 0–4.

It is then easy to verify that $C = \bigcup_{i < \omega} C_i$ and $f = \bigcup_{i < \omega} f_i$ are as required. \square

With Lemma 7.2.3 we are now ready to construct the advertised model.

Lemma 7.2.4. *Take some character χ realised in \mathfrak{A} . There is a structure $\mathfrak{A}^* \models \varphi$ over the same domain as \mathfrak{A} such that:*

- $\text{pr}^{\mathfrak{A}^*}[a] = \text{pr}^{\mathfrak{A}}[a]$ for each $a \in A$ satisfying $\text{chr}^{\mathfrak{A}}[a] \neq \chi$, and
- $\text{pr}^{\mathfrak{A}^*}[a] = \lceil \rho \rceil$ for each $a \in A$ having $\text{chr}^{\mathfrak{A}}[a] = \chi$ and where $\rho := \text{pr}^{\mathfrak{A}}[a]$.

Proof. Let us fix ξ to be a 2-type that is diverging for χ and write B for the set $\{b \in A \mid \text{chr}^{\mathfrak{A}}[b] = \chi\}$. Now, take $C \subseteq A$ and $f : C \rightarrow B$ to be as described in Lemma 7.2.3 and recall that $\text{tp}^{\mathfrak{A}}[f(c)c] = \xi$ for each $c \in C$. By divergence of ξ for χ , we must have that B and C are of cardinality \aleph_0 . Note that $B \cap C$ is not necessarily empty. This causes complications that are remedied by the fact that $f : C \rightarrow B$ is asymmetric.

Let us enumerate elements of B as b_1, b_2, \dots and partition the set C into pairwise disjoint infinite sets C_1, C_2, \dots . We present a high-level overview of the construction. Intuitively, C_i is treated as a set of potential ξ -witnesses for b_i . Recall that each element $c \in C_i$ is already a ξ -witness for $f(c)$ in \mathfrak{A} . The trick in our construction is simple: repurpose some infinite subset of the elements $c \in C_i$ to be ξ -witnesses for b_i by swapping the 2-types of $f(c)c$ and $b_i c$. Clearly, such an operation does not alter the profile of c . Yet caution is still needed as no guarantees are made for the profiles of $f(c)$ and b_i (in regards to the number of 2-types $\text{tp}^{\mathfrak{A}}[b_i c]$ emitted). However, we are getting ahead of ourselves. We provide a safe way of picking $c \in C_i$ in the following paragraph. For now, we simply that each element $b \in B$ will be provided with \aleph_0 witnesses for ξ .

Keeping b_1, b_2, \dots and C_1, C_2, \dots as before we set \mathfrak{A}' to be \mathfrak{A} but with 2-types of $B \times C$ left undefined. Now, for each $i < \omega$ we make the following observation about b_i and C_i . Since the number of different 2-types is finite and C_i is infinite,

we have, by the infinite pigeonhole principle, that there is some 2-type realised between b_i and C_i infinitely often in \mathfrak{A} . Denoting this 2-type by μ_i we note that it need not be the case that $\mu_i = \mu_j$ for each $j < \omega$. On the other hand, since elements b_j realise the character χ and $\chi(\mu_i) = \aleph_0$, we have that each b_j has infinitely many witnesses for μ_i in \mathfrak{A} .

Having picked μ_i for each $i < \omega$ let us now define the set D_i as follows:

$$\{c \in C_i \mid \text{tp}^{\mathfrak{A}}[b_i c] = \mu_i \text{ and } c \neq f(b_i), \text{ and } c \neq b_j \text{ for all } j \leq i\}. \quad (D_i)$$

Alternatively, D_i can be viewed as the set of μ_i -witnesses of b_i in \mathfrak{A} that do not cause unwanted dependencies. Clearly D_i is infinite as at most $i+1$ μ_i -witnesses of b_i are barred from being in the set. Thus, we may partition D_i into two disjoint infinite sets D_i^ξ and $D_i^{\mu_i}$ with the intent that elements of D_i^ξ will serve as ξ -witnesses, whilst elements of $D_i^{\mu_i}$ will be *retained* μ_i -witnesses.

We are now in a position to perform the following action for each $i < \omega$:

- set $\text{tp}^{\mathfrak{A}'}[b_i d] := \xi$ and $\text{tp}^{\mathfrak{A}'}[f(d)d] := \mu_i$ for each $d \in D_i^\xi$,
- set $\text{tp}^{\mathfrak{A}'}[b_i d] := \mu_i$ and $\text{tp}^{\mathfrak{A}'}[f(d)d] := \xi$ for each $d \in D_i^{\mu_i}$.

We claim that no clashes arise in the above as no pair of elements is ever assigned a 2-type twice. To see this take any two elements $a_1, a_2 \in C$. We have that $a_1 a_2$ is considered for a 2-type when at least one of the following conditions is met:

- (i) for some $k < \omega$ we have $a_1 = b_k$ and $a_2 \in D_k$;
- (ii) for some $\ell < \omega$ we have $a_2 = b_\ell$ and $a_1 \in D_\ell$;
- (iii) $f(a_2) = a_1$;
- (iv) $f(a_1) = a_2$;

It is easy to verify that the above list is exhaustive. We now claim that the events (i)–(iv) are disjoint. For suppose conditions (i) and (ii) are met. Without loss of generality, say that $k \leq \ell$. Then, by (ii), $a_2 = b_\ell$ and $a_1 \in D_\ell$. By definition, $b_k \notin D_\ell$ as $k \leq \ell$. But $a_1 = b_k$ by (i) thus contradicting our initial assumption. On the other hand, let us suppose that (i) and (iii) both hold. By (i) we have that $a_2 \in D_k$, thus ensuring $\text{tp}^{\mathfrak{A}'}[a_1 a_2] = \mu_k$. By (iii) we have $f(a_2) = a_1$ thus implying $\text{tp}^{\mathfrak{A}'}[a_1 a_2] = \xi$. But if $\xi = \mu_k$, then $\chi(\xi) = \chi(\mu_k) = \aleph_0$ contradicts divergence of ξ for χ . When considering (i) and (iv) we have $f(b_k) = f(a_1) = a_2$ and $a_2 \in D_k$.

But, by definition, we cannot have $f(b_k) \in D_k$. Moving on, we see that the case of both (ii) and (iii) holding is symmetric to (i) and (iv) holding, whilst (ii) and (iv) is symmetric to (i) and (iii). We thus need only consider (iii) and (iv) being true at the same time. But this is immediate by the fact that f is defined to be asymmetric in Lemma 7.2.3. With this we have safely secured an infinite number of ξ -witnesses for each element of B .

Ending the construction phase of our lemma we set $\text{tp}^{\mathfrak{A}'}[bc] := \text{tp}^{\mathfrak{A}}[bc]$ for each pair $bc \in B \times C$ that was not given a 2-type in the above.

We claim that $\text{pr}^{\mathfrak{A}'}[a] = \text{pr}^{\mathfrak{A}}[a]$ for all $a \in A \setminus B$ and, for $b \in B$, that $(\text{pr}^{\mathfrak{A}'}[b])(\xi) = \aleph_0$ while $(\text{pr}^{\mathfrak{A}'}[b])(\mu) = (\text{pr}^{\mathfrak{A}}[b])(\mu)$ for all 2-types $\mu \neq \xi$. We start by showing the former. The statement is trivially true if $a \notin C$ as then $\text{tp}^{\mathfrak{A}'}[aa'] = \text{tp}^{\mathfrak{A}}[aa']$ for each $a' \in A$. Otherwise, suppose that $a \in C$. If $i < \omega$ is such that $a \in C_i$, we have $\text{tp}^{\mathfrak{A}'}[aa'] = \text{tp}^{\mathfrak{A}}[aa']$ for each $a' \in A \setminus \{b_i, f(a)\}$. Then, if $a \in D_i^\xi$, we have $\text{tp}^{\mathfrak{A}'}[b_i a] = \text{tp}^{\mathfrak{A}}[f(a)a]$ and $\text{tp}^{\mathfrak{A}'}[f(a)a] = \text{tp}^{\mathfrak{A}}[b_i a]$. Otherwise, $\text{tp}^{\mathfrak{A}'}[b_i a] = \text{tp}^{\mathfrak{A}}[b_i a]$ and $\text{tp}^{\mathfrak{A}'}[f(a)a] = \text{tp}^{\mathfrak{A}}[f(a)a]$. Clearly, in both cases, $\text{pr}^{\mathfrak{A}'}[a] = \text{pr}^{\mathfrak{A}}[a]$ as required.

We now show the later claim. Take $b_i \in B$ and write $\rho := \text{pr}^{\mathfrak{A}}[b_i]$, and $\rho' := \text{pr}^{\mathfrak{A}'}[b_i]$. By our assignments concerning b_i and sets $D_i^\xi, D_i^{\mu_i}$, we have $\rho'(\xi) = \aleph_0$ and $\rho'(\mu_i) = \rho(\mu_i) = \aleph_0$. We need only show that $\rho'(\mu) = \rho(\mu)$ for 2-types $\mu \notin \{\xi, \mu_i\}$. Let us suppose that $b_i \in D_k$ for some $k < i$. (The case where this does not hold is similar but easier). It is easy to verify that, for each $a \in A$, we have $\text{tp}^{\mathfrak{A}'}[b_i a] = \text{tp}^{\mathfrak{A}}[b_i a]$ as long as $a \notin D_i^\xi$, $a \neq b_k$, $a \neq f(b_i)$, and $f(a) \neq b_i$. By our construction, $\text{tp}^{\mathfrak{A}'}[b_i a] = \xi$ for all $a \in D_i^\xi$, whilst, in the original model we had $\text{tp}^{\mathfrak{A}}[b_i a] = \mu_i$. Thus, in \mathfrak{A}' , we need only account for the 2-types of $b_k b_i$, $f(b_i) b_i$ and $b_i a$ when $f(a) = b_i$. Since $b_i \in D_k$, we have $\text{tp}^{\mathfrak{A}'}[b_k b_i] = \text{tp}^{\mathfrak{A}}[f(b_i) b_i]$ and $\text{tp}^{\mathfrak{A}'}[f(b_i) b_i] = \text{tp}^{\mathfrak{A}}[b_k b_i]$, or $\text{tp}^{\mathfrak{A}'}[b_k b_i] = \text{tp}^{\mathfrak{A}}[b_k b_i]$ and $\text{tp}^{\mathfrak{A}'}[f(b_i) b_i] = \text{tp}^{\mathfrak{A}}[f(b_i) b_i]$. Thus, from the perspective of b_i , the count of 2-types was not altered by this (possible) switch. Turning to the tuple $b_i a$ notice that, since f is injective, there is at most one element $a \in A$ for which $f(a) = b_i$. If such an element indeed exists, then $\text{tp}^{\mathfrak{A}'}[b_i a] = \xi$. Since we already argued that $\rho'(\xi) = \aleph_0$, the element a can be repurposed. This is exactly what happens when $a \in D_\ell^\xi$ for some $\ell < \omega$. Then $\text{tp}^{\mathfrak{A}'}[b_i a] = \text{tp}^{\mathfrak{A}}[b_\ell a] = \mu_\ell$ thus making $\rho'(\mu_\ell) = \rho(\mu_\ell) + 1$. But, μ_ℓ was picked in such a way that $\chi(\mu_\ell) = \aleph_0 (= \rho(\mu_\ell))$. If no such ℓ exists, then $\text{tp}^{\mathfrak{A}'}[b_i a] = \text{tp}^{\mathfrak{A}}[b_i a] = \xi$. In either case, a makes no meaningful contribution to the profile of b_i in \mathfrak{A}' .

Notice that elements of B now realise a profile χ' for which ξ is no longer diverging. Additionally, no new diverging 2-types are created in the process. By repeating this procedure on \mathfrak{A}' and χ' (and subsequent structures and characters) we reach a model \mathfrak{A}^* in which elements that were of character χ in \mathfrak{A} now realise $\lceil \chi \rceil$ in \mathfrak{A}^* . In fact, for each b or character χ in \mathfrak{A} we have $\text{pr}^{\mathfrak{A}^*}[b] = \lceil \rho \rceil$, where $\rho := \text{pr}^{\mathfrak{A}}[b]$. By Lemma 7.2.2, $\lceil \rho \rceil$ is φ -compliant. Moreover, elements $a \in A \setminus B$ retain the φ -compliant profile $\text{pr}^{\mathfrak{A}}[a]$ in \mathfrak{A}^* . Thus, $\mathfrak{A}^* \models \varphi$ by Lemma 7.1.3. \square

Suppose that χ_1, \dots, χ_n is the list of diverging characters in \mathfrak{A} . Writing $\mathfrak{A}_0 := \mathfrak{A}$ let $\mathfrak{A}_1, \dots, \mathfrak{A}_n$ be such that, for $i \in [1, n]$, \mathfrak{A}_i is the structure obtained by applying Lemma 7.2.4 on \mathfrak{A}_{i-1} and the character χ_i . Clearly, $\mathfrak{A}_i \models \varphi$. Additionally, when compared to \mathfrak{A}_{i-1} , no other profiles other than those of character χ_i are altered in \mathfrak{A}_i . Thus, \mathfrak{A}_i has strictly fewer diverging characters than \mathfrak{A}_{i-1} . In fact, the final structure, \mathfrak{A}_n , has none at all.

When considering satisfiable normal-form sentences of $\mathcal{C}_{\mathcal{U}\mathcal{H}\mathcal{A}}^2$ we can now, due to Lemma 7.2.4, confine attention to models in which no diverging characters are realised. In other words, we will only consider structures in which the number of 2-types ξ emitted from any given element is either bounded by some $n \in \mathbb{N}$ (i.e. converging) or is exactly \aleph_0 . Using this fact, we show the following:

Lemma 7.2.5. *The satisfiability problem for $\mathcal{C}_{\mathcal{U}\mathcal{H}\mathcal{A}}^2$ is RECURSIVELY ENUMERABLE.*

Proof. Suppose φ is a $\mathcal{C}_{\mathcal{U}\mathcal{H}\mathcal{A}}^2$ -sentence over Σ given as input. By Lemma 7.1.1, we assume that φ is in normal-form recapitulated below for convenience:

$$\forall xy \alpha \wedge \bigwedge_{s \in S} \forall x \exists_{[=M_s]} y \beta_s \wedge \bigwedge_{t \in T} \forall x \text{I}(y, y)(\gamma_t, \delta_t).$$

In the forthcoming procedure we will not deal with uniform Härtig assertions directly. Instead, we compute an infinite series ψ_1, ψ_2, \dots of $\mathcal{FO}_{\text{Pres}}^2$ -sentences such that φ is satisfiable if and only if there is some $k \in \mathbb{N}$ for which ψ_k is satisfiable. Informally, in ψ_k we “guess” that the number of elements considered by any given uniform Härtig assertion is at most k or \aleph_0 .

We construct the new sentences as follows. Let us take $k \in \mathbb{N}$. For each $t \in T$ define a series of fresh unary predicates $\mathbf{p}_0^t, \dots, \mathbf{p}_k^t$ alongside $\mathbf{p}_{\aleph_0}^t$. We define ψ_k to be a conjunction of three $\mathcal{FO}_{\text{Pres}}^2$ -sentences θ_1, θ_2 and θ_3 defined as follows. The

first sentence is simply the universal and existential conjuncts copied from φ :

$$\forall xy \alpha \wedge \bigwedge_{s \in S} \forall x \exists_{=[M_s]} y \beta. \quad (\theta_1)$$

We now turn to θ_2 which is a technical requirement regarding our new predicates. For each $t \in T$, we require that each element satisfies exactly one of the unary predicates $\mathfrak{p}_0^t, \dots, \mathfrak{p}_k^t, \mathfrak{p}_{\aleph_0}^t$:

$$\bigwedge_{t \in T} \forall x \left(\bigvee_{i \in [0, k] \cup \{\aleph_0\}} \mathfrak{p}_i^t(x) \wedge \bigwedge_{\substack{i \neq j \\ i, j \in [0, k] \cup \{\aleph_0\}}} (\neg \mathfrak{p}_i^t(x) \vee \neg \mathfrak{p}_j^t(x)) \right). \quad (\theta_2)$$

The satisfaction of \mathfrak{p}_i^t by a in some model \mathfrak{A} is, in a sense, a guess that the cardinality of $\{b \in A \mid \mathfrak{A} \models \gamma_t[ab]\}$ (and thus also $\{b \in A \mid \mathfrak{A} \models \delta_t[ab]\}$) is exactly i . Thus, fixing some $t \in T$, we rewrite the t -th uniform Härtig assertion of φ in $\mathcal{FO}_{\text{Pres}}^2$ as follows:

$$\bigwedge_{i \in [0, k] \cup \{\aleph_0\}} \forall x \left(\mathfrak{p}_i^t(x) \rightarrow (\exists_{[=i]} y \gamma_t \wedge \exists_{[=i]} y \delta_t) \right). \quad (\theta_{3,t})$$

We define $\theta_3 := \bigwedge_{t \in T} \theta_{3,t}$ and claim that φ is satisfiable if and only if there are some $k \in \mathbb{N}$ for which ψ_k is satisfiable.

It is easy to verify that $\psi_k \models \varphi$ for each $k \in \mathbb{N}$ thus securing the “if” direction of our claim. Thus, let us turn to the “only-if” direction and suppose $\mathfrak{A} \models \varphi$ is countable. We may, by Lemma 7.2.4, assume that there are no diverging 2-types for any character realised in \mathfrak{A} . Fixing some $t \in T$ we show the following consequence of convergence: there is some $k_t \in \mathbb{N}$ such that, for each $a \in A$, the set $\{b \in A \mid \text{and } \mathfrak{A} \models \gamma_t[ab]\}$ is of cardinality at most k_t or exactly \aleph_0 . To see this recall that there are finitely different character in \mathfrak{A} . Thus, we may list them as χ_1, \dots, χ_m . Taking any $i \in [1, m]$ and supposing that $\chi_i(\xi) = 1^{+1}$ for some 2-type ξ , we have, by convergence, that there is some $k_{\chi_i, \xi} \in \mathbb{N}$ satisfying $|\{b \in A \mid \mathfrak{A} \models \xi[ab]\}| \leq k_{\chi_i, \xi}$ for each $a \in A$ of character χ_i . Writing $k_{\chi_i, \xi} := 0$ whenever $\chi_i(\xi) \neq 1^{+1}$ we have, for each $a \in A$, that $|\{b \in A \mid \text{and } \mathfrak{A} \models \xi[ab]\}|$ is \aleph_0 or at most $k_\xi := \max_{i \in [1, m]} (k_{\chi_i, \xi})$. It is then easy to verify that $k_t := \sum_{\substack{\xi \models \gamma_t \\ \xi \in \text{ATP}_2^\Sigma}} k_\xi$ is the required bound on $\{b \in A \mid \text{and } \mathfrak{A} \models \gamma_t[ab]\}$ for each $a \in A$. By writing $k = \max_{t \in T} (k_t)$ we expand \mathfrak{A} into a model $\mathfrak{A}' \models \psi_k$ by setting $\mathfrak{A}' := \mathfrak{A}$ and $\mathfrak{A}' \models \mathfrak{p}_i^t[a]$ for each $a \in A$ and $t \in T$, and where $i = |\{b \in A \mid \mathfrak{A} \models \gamma_t[ab]\}|$

(= $|\{b \in A \mid \mathfrak{A} \models \delta_t[ab]\}|$). Clearly, either $i = \aleph_0$ or $i \leq k_t \leq k$ thus ensuring that the assignment does not go out of bounds. It is then easy to verify that $\mathfrak{A}' \models \psi_k$.

Recall that $\mathcal{FO}_{\text{Pres}}^2$ has a decidable satisfiability problem [Benedikt et al., 2024]. Thus, in the event that φ is satisfiable, a Turing machine checking satisfiability of ψ_1, ψ_2, \dots one-by-one will eventually reach a formula that is satisfiable. Hence, the language $\text{Sat}(\mathcal{C}_{\mathcal{UHA}}^2)$ is RECURSIVELY ENUMERABLE. \square

Recall that \mathcal{AFC}^3 is a subfragment of \mathcal{FO} and thus has a satisfiability problem which is CO-RECURSIVELY ENUMERABLE. Combining the lemma above with the fact that each $\mathcal{C}_{\mathcal{UHA}}^2$ -sentence can be translated into an equisatisfiable \mathcal{AFC}^3 -sentence (Lemma 7.1.2), we conclude the section having proved following theorem:

Theorem 7.2.6. *The satisfiability problem for $\mathcal{C}_{\mathcal{UHA}}^2$ is RECURSIVE.*

7.3 Reducing \mathcal{AFC}^3 to $\mathcal{C}_{\mathcal{UHA}}^2$

In this section decidability of $\text{Sat}(\mathcal{AFC}^3)$ by reducing the problem to $\text{Sat}(\mathcal{C}_{\mathcal{UHA}}^2)$. We start off with a technical observation regarding the identity relations in \mathcal{AFC}^3 .

Lemma 7.3.1. *There is an \mathcal{AFC}^3 -sentence φ_ϵ axiomatising a ternary predicate ϵ such that $\varphi_\epsilon \models \forall \mathbf{x}_3 (\epsilon(\mathbf{x}_3) \leftrightarrow x_1 = x_3)$.*

Proof. The result is immediate by setting φ_ϵ to be the following:

$$\forall \mathbf{x}_2 \epsilon(x_1 x_2 x_1) \wedge \forall \mathbf{x}_2 \exists_{[=1]} x_3 \epsilon(\mathbf{x}_3).$$

Take any $\mathfrak{A} \models \varphi_\epsilon$ and $a, b \in A$. By the first conjunct of φ_ϵ we have that $aba \in \epsilon^{\mathfrak{A}}$. Suppose now that $abc \in \epsilon^{\mathfrak{A}}$ but $a \neq c$. Then $\{d \in A \mid \mathfrak{A} \models \epsilon[abd]\}$ is of cardinality at least 2 thus contradicting the second conjunct of φ_ϵ . \square

As a consequence we will assume that the predicate ϵ (axiomatised as in φ_ϵ) is readily available in \mathcal{AFC}^3 . Say that an \mathcal{AFC}^3 -sentence is in *normal-form* if it takes the following shape:

$$\forall \mathbf{x}_3 \alpha \wedge \bigwedge_{s \in S} \forall \mathbf{x}_2 \exists_{[=M_s]} x_3 (\beta_s \wedge \neg \epsilon(\mathbf{x}_3) \wedge x_2 \neq x_3), \quad (\mathcal{AFC}^3\text{-nmf})$$

where α, β_s are quantifier-free \mathcal{AFC}^3 -formulas indexed by a finite set S and each $M_s \geq 1$. By rewriting formulas we have the following:

Lemma 7.3.2. *Suppose that φ is an \mathcal{AFC}^3 -sentence and let M_{\max} be the maximal numeric subscript (of a counting quantifier) occurring in φ . We may then compute, in polynomial time, an \mathcal{AFC}^3 -sentence ψ in normal-form such that φ and ψ are satisfiable over the same domains of size at least $M_{\max}+3$.*

Proof. Step 1: removing $\exists_{[=M]}$. Let $\varphi_0^0 := \varphi$ and take any subformula $\theta(\mathbf{x}_k)$ of the form $\exists_{[=M]}x_{k+1}\chi$ with χ not containing any quantifier of the form $\exists_{[=M']}$. We introduce two fresh k -ary symbols \mathbf{p} and \mathbf{q} , and set φ_1^0 to be φ_0^0 but with $\theta(\mathbf{x}_k)$ replaced by the formula $\mathbf{p}(\mathbf{x}_k) \wedge \mathbf{q}(\mathbf{x}_k)$ whilst also defining ψ_1^0 to be

$$\forall \mathbf{x}_k (\mathbf{p}(\mathbf{x}_k) \leftrightarrow \exists_{[\leq M]}x_{k+1}\chi) \wedge \forall \mathbf{x}_k (\mathbf{q}(\mathbf{x}_k) \leftrightarrow \exists_{[\geq M]}x_{k+1}\chi).$$

Clearly, $\varphi_1^0 \wedge \psi_1^0 \models \varphi_0^0$ whilst models $\mathfrak{A} \models \varphi_0^0$ can be extended to $\mathfrak{A}' \models \varphi_1^0 \wedge \psi_1^0$ by setting $\mathfrak{A}' := \mathfrak{A}$, $\mathbf{p}^{\mathfrak{A}'} := \{\bar{a} \in A^k \mid \mathfrak{A}, \bar{a} \models \exists_{[\leq M]}x_{k+1}\chi\}$, and $\mathbf{q}^{\mathfrak{A}'} := \{\bar{a} \in A^k \mid \mathfrak{A}, \bar{a} \models \exists_{[\geq M]}x_{k+1}\chi\}$. By processing φ_1^0 and subsequent formulas in the same way, we will eventually reach a formula φ_m^0 that has no quantifiers of the form $\exists_{[=M]}$. We proceed to the next step of the construction by setting $\varphi^1 := \varphi_m^0 \wedge \psi_1^0 \wedge \dots \wedge \psi_m^0$ and noting that φ and φ^1 are satisfiable over the same domains.

Step 2: normalisation. We now compute a formula φ^2 that is satisfiable over the same domains as φ^1 . We make sure that φ^2 takes the following form:

$$\forall \mathbf{x}_3 \alpha^2(\mathbf{x}_3) \wedge \bigwedge_{z \in Z} \forall \mathbf{x}_2 (\delta_z(\mathbf{x}_2) \rightarrow \exists_{[\bowtie_z M_z]}x_3 \beta_z^2(\mathbf{x}_3)),$$

where Z is a finite set of indices such that for each $z \in Z$ the symbols \bowtie_z are either \leq or \geq , the numeric subscripts satisfy $0 \leq M_z \leq M_{\max}+1$, and α^2 , δ_z , β_z^2 are quantifier-free \mathcal{AF}^3 -formulas.

Write $\varphi_0^1 := \varphi^1$ and assume that subformulas with standard quantifiers $\forall x \theta$ and $\exists x \chi$ are converted to their counting quantifier equivalents $\exists_{[\leq 0]}x \neg \theta$ and $\exists_{[\geq 1]}x \chi$. Additionally, replace any instance of $\exists_{[\geq 0]}x \theta$ by \top as such formulas are trivially valid. Now, let us take any formula $\theta(\mathbf{x}_k)$ of the form $\exists_{[\bowtie M]}x_{k+1}\chi$ (where \bowtie is, again, either \leq or \geq). Writing $\overline{\exists_{[\leq M]}}$ for $\exists_{[\geq M+1]}$ and $\overline{\exists_{[\geq M]}}$ for $\exists_{[\leq M-1]}$ we introduce a fresh k -ary symbol \mathbf{p} and define φ_1^1 to be φ_0^1 but with $\theta(\mathbf{x}_k)$ replaced with $\mathbf{p}(\mathbf{x}_k)$. Additionally, we define ψ_1^1 to be the following sentence:

$$\forall \mathbf{x}_k (\mathbf{p}(\mathbf{x}_k) \rightarrow \exists_{[\bowtie M]}x_{k+1}\chi) \wedge \forall \mathbf{x}_k (\neg \mathbf{p}(\mathbf{x}_k) \rightarrow \overline{\exists_{[\bowtie M]}x_{k+1}\chi}).$$

It is easy to see that $\varphi_1^1 \wedge \psi_1^1 \models \varphi_0^1$. On the other hand, if $\mathfrak{A} \models \varphi_0^1$, then by setting

$\mathfrak{A}' := \mathfrak{A}$ and $\mathfrak{p}^{\mathfrak{A}'}$:= $\{\bar{a} \in A^k \mid \mathfrak{A} \models \theta[\bar{a}]\}$ we will have obtained a model of $\varphi_1^1 \wedge \psi_1^1$.

Now, let us process φ_1^1 and subsequent formulas as above until a sentence φ_m^1 composed solely of proposition letters is reached. We will then have, after a trivial rearrangement of formulas, that the sentence $\varphi^2 := \psi_1^1 \wedge \dots \wedge \psi_m^1 \wedge \varphi_m^1$ is of the required form and satisfiable over the same domains as φ .

Step 3: reintroducing $\exists_{[=M]}$.

We will transform φ^2 into φ^3 of the form:

$$\forall \mathbf{x}_3 \alpha^3(\mathbf{x}_3) \wedge \bigwedge_{z \in Z} \forall \mathbf{x}_2 \exists_{[=M_z]} x_3 \beta_z^3(\mathbf{x}_3),$$

Returning to φ^2 , for each $z \in Z$ we introduce a fresh ternary predicate symbol \mathfrak{p}_z and define the following formula:

$$\psi_z := \begin{cases} \forall \mathbf{x}_2 \exists_{[=M_z]} x_3 \mathfrak{p}_z(\mathbf{x}_3) \wedge \forall \mathbf{x}_3 ((\delta_z(\mathbf{x}_2) \wedge \beta_z^2(\mathbf{x}_3)) \rightarrow \mathfrak{p}_z(\mathbf{x}_3)) & \text{if } \bowtie_z \equiv \leq, \\ \forall \mathbf{x}_2 \exists_{[=M_z]} x_3 \mathfrak{p}_z(\mathbf{x}_3) \wedge \forall \mathbf{x}_3 (\mathfrak{p}_z(\mathbf{x}_3) \rightarrow (\delta_z(\mathbf{x}_2) \rightarrow \beta_z^2(\mathbf{x}_3))) & \text{if } \bowtie_z \equiv \geq. \end{cases}$$

After rearrangements, the formula $\varphi^3 := \forall \mathbf{x}_3 \alpha^2 \wedge \bigwedge_{z \in Z} \psi_z$ is of the required form. We claim that φ^3 is satisfiable over the same domains of size at least $M_{\max}+1$ as φ . It is immediate that $\varphi^3 \models \varphi^2$ as $\psi_z \models \forall \mathbf{x}_2 \exists_{[\bowtie_z M_z]} x_3 \beta_z^2$. For the converse direction, suppose that $\mathfrak{A} \models \varphi^2$ with $|A| \geq M_{\max}+1$. Taking $\bar{a} \in A^2$ we must have, for each $z \in Z$, that $\mathfrak{A}, \bar{a} \models \delta_z(\mathbf{x}_2) \rightarrow \exists_{[\bowtie_z M_z]} x_3 \beta_z^2(\mathbf{x}_3)$. Define $B_{\bar{a}}$ to be a subset of A of size exactly M_z satisfying the following:

$$\begin{aligned} \{b \in A \mid \mathfrak{A} \models \beta_z^2[\bar{a}b]\} &\subseteq B_{\bar{a}} & \text{if } \bowtie_z \equiv \leq \\ B_{\bar{a}} &\subseteq \{b \in A \mid \mathfrak{A} \models \beta_z^2[\bar{a}b]\} & \text{if } \bowtie_z \equiv \geq \end{aligned}$$

Since $M_z \leq M_{\max}+1$, the set $B_{\bar{a}}$ is well-defined. By setting $\mathfrak{A}' := \mathfrak{A}$ and $\mathfrak{A}' \models \mathfrak{p}_z[\bar{a}b]$ for each $b \in B_{\bar{a}}$ we will have $\mathfrak{A}', \bar{a} \models \exists_{[=M_z]} x_3 \mathfrak{p}_z(\mathbf{x}_3)$. It is then easy to verify that $\mathfrak{A}', \bar{a} \models \forall x_3 ((\delta_z(\mathbf{x}_2) \wedge \beta_z^2(\mathbf{x}_3)) \rightarrow \mathfrak{p}_z(\mathbf{x}_3))$ if $\bowtie_z \equiv \leq$, and, in case $\bowtie_z \equiv \geq$, $\mathfrak{A}', \bar{a} \models \forall x_3 (\mathfrak{p}_z(\mathbf{x}_3) \rightarrow (\delta_z(\mathbf{x}_2) \rightarrow \beta_z^2(\mathbf{x}_3)))$. Repeating this procedure for each pair $\bar{a} \in A^2$ and $z \in Z$ will yield the required results.

Step 4: adding =.

We now compute φ^4 of the following form:

$$\forall \mathbf{x}_3 \alpha^4(\mathbf{x}_3) \wedge \bigwedge_{t \in T} \forall \mathbf{x}_2 \exists_{[=M_t]} x_3 (\beta_t^4(\mathbf{x}_3) \wedge x_2 \neq x_3),$$

where α^4, β_t^4 are quantifier-free \mathcal{AF}^3 -formulas and $1 \leq M_t \leq M_{\max} + 1$ with $t \in T$.

Returning to φ^3 and taking any $z \in Z$ let \mathbf{p}_z and \mathbf{p}'_z be fresh ternary predicates. Let us fix $\mathfrak{A} \models \varphi_3$ of size $M_z + 1$ as a motivating example and say that a tuple $ab \in A^2$ exhibits self-witnessing behaviour if $\mathfrak{A} \models \beta_z^3[abb]$. We write the following two sentences to rid ourselves of self-witnesses:

$$\forall \mathbf{x}_3 \left(\beta_z^3(x_1 x_2 x_2) \rightarrow (\beta_s^3(\mathbf{x}_3) \leftrightarrow \mathbf{p}_z(\mathbf{x}_3)) \right) \wedge \forall \mathbf{x}_2 \exists_{[=M_z-1]} x_3 (\mathbf{p}_z(\mathbf{x}_3) \wedge x_2 \neq x_3), \quad (\gamma_z)$$

$$\forall \mathbf{x}_3 \left(\neg \beta_z^3(x_1 x_2 x_2) \rightarrow (\beta_s^3(\mathbf{x}_3) \leftrightarrow \mathbf{p}'_z(\mathbf{x}_3)) \right) \wedge \forall \mathbf{x}_2 \exists_{[=M_z]} x_3 (\mathbf{p}'_z(\mathbf{x}_3) \wedge x_2 \neq x_3). \quad (\delta_z)$$

In short, the first conjunct of γ_z checks for self-witnessing behaviour. In case it occurs, the second conjunct of γ_z discounts the self-witness from the existential requirement. The sentence δ_z handles the case where the self-witness does not occur. Note that if self-witnessing behaviour is not detected, then tuples satisfying \mathbf{p}_z need not satisfy β_z^3 . Similarly, in the case where self-witnessing is present, tuples satisfying \mathbf{p}'_z need not satisfy β_z^3 . It should now be clear that \mathfrak{A} can be expanded into a model \mathfrak{A}' of $\gamma_z \wedge \delta_z$ by setting:

- $\mathfrak{A}' \models \mathbf{p}_z[abc] \Leftrightarrow \mathfrak{A} \models \beta_z^3[abc]$ for $a, b, c \in A$ with $\mathfrak{A} \models \beta_z^3[abb]$,
- $\mathfrak{A}' \models \mathbf{p}_z[abc]$ for $a, b \in A$ with $\mathfrak{A} \not\models \beta_z^3[abb]$ and all elements c that belong to some $M_z - 1$ sized subset of $A \setminus \{b\}$.
- $\mathfrak{A}' \models \mathbf{p}'_z[abc] \Leftrightarrow \mathfrak{A} \models \beta_z^3[abc]$ for $a, b, c \in A$ with $\mathfrak{A} \not\models \beta_z^3[abb]$,
- $\mathfrak{A}' \models \mathbf{p}'_z[abc]$ for $a, b \in A$ with $\mathfrak{A} \models \beta_z^3[abb]$ and all elements c that belong to some M_z sized subset of $A \setminus \{b\}$.

For the converse direction, it is easy to verify that $\gamma_z \wedge \delta_z \models \exists_{[=M_z]} x_3 \beta_z^3(\mathbf{x}_3)$. Processing each $z \in Z$ in this way we arrive at a formula $\varphi^4 := \forall \mathbf{x}_3 \alpha^4(\mathbf{x}_3) \wedge \bigwedge_{z \in Z} (\gamma_z \wedge \delta_z)$ that is, after trivial rearrangements, of the required form.

Step 5: adding ϵ .

We finally compute φ^5 of the required form (\mathcal{AFC}^3 -nmf):

$$\forall \mathbf{x}_3 \alpha(\mathbf{x}_3) \wedge \bigwedge_{s \in S} \forall \mathbf{x}_2 \exists_{[=M_s]} x_3 (\beta_s(\mathbf{x}_3) \wedge \neg \epsilon(\mathbf{x}_3) \wedge x_2 \neq x_3),$$

where α, β_s are quantifier-free \mathcal{AF}^3 -formulas indexed by a finite set S and for each $s \in S$ we have $1 \leq M_s \leq M_{\max} + 1$.

Similarly as in the last step we turn to φ^4 and fix some $t \in T$. Taking any $\mathfrak{A} \models \varphi^4$ of size at least $M_z + 2$ and $a, b \in A$ we say that ab exhibits self-witnessing behaviour if $\mathfrak{A} \models \beta_t^\neq[aba]$, where $\beta_t^\neq := \beta_t(\mathbf{x}_3) \wedge x_2 \neq x_3$. Note that if $a = b$, then self-witnessing behaviour cannot occur. Setting \mathbf{p}_t and \mathbf{p}'_t to be fresh ternary predicates, we safeguard against self-witnesses by writing the following:

$$\forall \mathbf{x}_3 \left(\beta_t^\neq(x_1 x_2 x_1) \rightarrow (\beta_t^\neq(\mathbf{x}_3) \leftrightarrow \mathbf{p}_t(\mathbf{x}_3)) \right) \wedge \forall \mathbf{x}_2 \exists_{[=M_t-1]} x_3 (\mathbf{p}_t(\mathbf{x}_3) \wedge \neg \epsilon(\mathbf{x}_3) \wedge x_2 \neq x_3), \quad (\gamma_t)$$

$$\forall \mathbf{x}_3 \left(\neg \beta_t^\neq(x_1 x_2 x_1) \rightarrow (\beta_s^\neq(\mathbf{x}_3) \leftrightarrow \mathbf{p}'_t(\mathbf{x}_3)) \right) \wedge \forall \mathbf{x}_2 \exists_{[=M_t]} x_3 (\mathbf{p}'_t(\mathbf{x}_3) \wedge \neg \epsilon(\mathbf{x}_3) \wedge x_2 \neq x_3). \quad (\delta_t)$$

The above sentences work similarly to those computed in step 4. We define $\mathfrak{A}' := \mathfrak{A}$ to be a structure interpreting the new symbols by setting the following:

- $\mathfrak{A}' \models \mathbf{p}_t[abc] \Leftrightarrow \mathfrak{A} \models \beta_t^\neq[abc]$ for $a, b, c \in A$ with $\mathfrak{A} \models \beta_t^\neq[aba]$,
- $\mathfrak{A}' \models \mathbf{p}_t[abc]$ for $a, b \in A$ with $\mathfrak{A} \not\models \beta_t^\neq[aba]$ and all elements c that belong to some $M_t - 1$ sized subset of $A \setminus \{a, b\}$.
- $\mathfrak{A}' \models \mathbf{p}'_t[abc] \Leftrightarrow \mathfrak{A} \models \beta_t^\neq[abc]$ for $a, b, c \in A$ with $\mathfrak{A} \not\models \beta_t^\neq[aba]$,
- $\mathfrak{A}' \models \mathbf{p}'_t[abc]$ for $a, b \in A$ with $\mathfrak{A} \models \beta_t^\neq[aba]$ and all elements c that belong to some M_t sized subset of $A \setminus \{a, b\}$.

We first verify that $\mathfrak{A}' \models \gamma_t$. Take any $a, b \in A$. If ab does not exhibit self-witnessing behaviour, then the first conjunct of δ_t is inactive. The second conjunct of δ_t is then satisfied by our assignment. If, on the other hand, ab has a self-witness, then there are $M_t - 1$ elements $c \neq a$ such that $\mathfrak{A}' \models \mathbf{p}_t[abc]$. In fact, $c \neq b$ by $\mathfrak{A} \models \beta_t^\neq[abc]$. Thus, the second conjunct of δ_t is satisfied.

Moving on to δ_t , take any $a, b \in A$. The interesting case is where ab does not exhibit self-witnessing behaviour. Then there are M_s elements $c \in A$ having $c \neq a$. Note, again, that $c \neq b$ by $\mathfrak{A} \models \beta_t^\neq[abc]$.

The converse direction is trivial. By repeating this construction for each $t \in T$ we will have that the sentence $\varphi^5 := \forall \mathbf{x}_3 \alpha^4 \wedge \bigwedge_{t \in T} (\gamma_t \wedge \delta_t)$ is of the required form after trivial rearrangements. \square

For the rest of the section we fix φ to be some \mathcal{AFC}^3 -sentence in normal-form (\mathcal{AFC}^3 -nmf). We define the adjacent closure of φ , denoted $\text{acl}(\varphi)$, as:

$$\forall x_1 x_2 \bigwedge \{ \alpha^f \mid f : [1, 3] \rightarrow [1, 2] \} \wedge \bigwedge_{s \in S} \forall x_1 \exists_{[=M_s]} x_2 (\beta_s(x_1 x_1 x_2) \wedge x_1 \neq x_2). \quad (\text{acl}(\varphi))$$

Renaming x_1 to x and x_2 to y it is easy to see that $\text{acl}(\varphi)$ is a \mathcal{C}^2 -sentence. It is easy to extend Lemma 5.2.5 to \mathcal{AFC}^3 and show that $\varphi \models \text{acl}(\varphi)$.

In the context of \mathcal{AFC}^3 , an adjacent 3-type ξ over some given signature Σ is a maximal consistent² set of adjacent atomic literals in variables x_1, x_2, x_3 and with symbols from $\Sigma \cup \{=, \epsilon\}$. Given a structure \mathfrak{A} , each $\bar{a} \in A^3$ realises a unique adjacent 3-type in \mathfrak{A} denoted by $\text{atp}^{\mathfrak{A}}[\bar{a}]$. Consider now the existential requirements of φ . We write:

$$\beta := \bigvee_{s \in S} \beta_s(\mathbf{x}_3) \wedge \neg \epsilon(\mathbf{x}_3) \wedge x_2 \neq x_3, \text{ along with}$$

$\beta^{-1} := \beta(x_3 x_2 x_1)$, and $\widehat{\beta} := \beta \wedge \beta^{-1}$. We say that an adjacent 3-type ξ is a *ray* if $\xi \models \beta$. A ray ξ is *unidirectional* if $\xi \not\models \widehat{\beta}$ and *bidirectional* otherwise. Given $\mathfrak{A} \models \varphi$ and $abc \in A^3$ with $\xi := \text{atp}^{\mathfrak{A}}[abc]$ we think of ab as *emitting* ξ and cb as *absorbing* it. (Note the order in which the latter tuple is presented).

Take any $\mathfrak{A} \models \varphi$ and $b \in A$. Say that \mathfrak{A} is *b-semichromatic* if there are no elements $a, c \in A$ such that $\text{tp}^{\mathfrak{A}}(ab) = \text{tp}^{\mathfrak{A}}(cb)$ and $\mathfrak{A} \models \widehat{\beta}[abc]$. We say that \mathfrak{A} is *b-chromatic* if it is *b-semichromatic* and there are no distinct $a, c, c' \in A$ such that $\mathfrak{A} \models \widehat{\beta}[abc] \wedge \widehat{\beta}[abc']$ and $\text{atp}^{\mathfrak{A}}[cb] = \text{atp}^{\mathfrak{A}}[c'b]$. Plainly put, *b-semichromaticity* says that, if ab sends a bidirectional ray to cb , then the two pairs differ in the 2-type they realise. Additionally, *b-chromaticity* stipulates that, if ab sends a bidirectional ray to cb and $c'b$, then the recipients differ in the 2-types they realise. A structure \mathfrak{A} is *chromatic* if it is *b-chromatic* for each $b \in A$. The following is an adaptation of [Pratt-Hartmann, 2010, Lemma 4]:

Lemma 7.3.3. *Suppose $\mathfrak{A} \models \varphi$. Then \mathfrak{A} can be expanded into a chromatic model.*

²Recall that $\epsilon(xyz)$ iff $x=z$. Thus, adjacent types as $\{\neg \epsilon(xyz), x=y, y=z\}$ are inconsistent.

Proof. Writing $k := \lceil \log_2(M^2+M+1) \rceil$ we build a structure $\mathfrak{B} := \mathfrak{A}$ that interprets fresh binary predicates $\mathbf{p}_1, \dots, \mathbf{p}_k$. We proceed by picking any $b \in A$ and defining a graph $G = (A, E \cup F)$ as follows:

$$\begin{aligned} E &:= \{a, c \in A \mid \mathfrak{A} \models \widehat{\beta}[abc]\}, \\ F &:= \{c, c' \in A \mid c \neq c' \text{ and } \mathfrak{A} \models \widehat{\beta}[abc] \wedge \widehat{\beta}[abc'] \text{ for some } a \in A\}. \end{aligned}$$

Since each pair of elements in \mathfrak{A} has at most M existential witnesses, we have that the degree of G does not exceed M^2+M . The graph G may then be properly coloured using M^2+M+1 colours (see [Pratt-Hartmann, 2023, Lemma 8.5]). Writing P for the power set of $\{\mathbf{p}_1, \dots, \mathbf{p}_k\}$ we see that $|P| = 2^k \geq M^2+M+1$. Thus, we may fix $f : A \rightarrow P$ to be some colouring of G . Setting $ab \in \mathbf{p}^{\mathfrak{B}}$ for each $a \in A$ and $\mathbf{p} \in f(a)$ we have that \mathfrak{B} is b -chromatic. Now, apply the operations above for some other $b' \in B$ thus obtaining a colouring f' . Notice that the colours $f'(b')$ and $f(b')$ need not coincide. But, for distinct pairs, satisfaction of $\mathbf{p}(x_1x_2)$ is unrelated to that of $\mathbf{p}(x_2x_1)$ for any $\mathbf{p} \in \{\mathbf{p}_1, \dots, \mathbf{p}_k\}$. Clearly, b -chromaticity is retained by applying the colouring procedure for b' . Thus, by running the construction for each $b \in B$, we produce the required chromatic structure \mathfrak{B} . \square

For the rest of the section fix $M := \sum_{s \in S} M_s$. Now, take some structure \mathfrak{A} , a 2-type ζ realised in \mathfrak{A} and $b \in A$. We say that ζ is *differentiated* for b in \mathfrak{A} if the set $\{a \in A \mid \mathbf{tp}^{\mathfrak{A}}[ab] = \zeta\}$ has cardinality *outside* of the range $[2, 3M]$. A structure \mathfrak{A} is differentiated for a set of 2-types Z if each 2-type $\zeta \in Z$ is differentiated for all $b \in A$. Lastly, say that \mathfrak{A} is differentiated if \mathfrak{A} is differentiated for all 2-types over $\text{sig}(\mathfrak{A})$. We provide an adaptation of [Pratt-Hartmann, 2010, Lemma 5] which is quite a bit more involved than the original:

Lemma 7.3.4. *Any structure \mathfrak{A} can be expanded into a differentiated model.*

Proof. Set $\ell := \lceil \log_2(3M) \rceil$ and let $\mathbf{q}_1, \dots, \mathbf{q}_\ell$ be binary predicates that are not part of $\text{sig}(\mathfrak{A})$. Now, say that a 2-type ζ is *primordial* if

$$\zeta \models \neg \mathbf{q}_1(x_1x_2) \wedge \dots \wedge \neg \mathbf{q}_\ell(x_1x_2) \text{ and } \zeta \models \neg \mathbf{q}_1(x_2x_1) \wedge \dots \wedge \neg \mathbf{q}_\ell(x_2x_1).$$

Similarly, we say that ζ is *evolved* if

$$\zeta \models \mathbf{q}_1(x_1x_2) \vee \dots \vee \mathbf{q}_\ell(x_1x_2) \text{ but } \zeta \models \neg \mathbf{q}_1(x_2x_1) \wedge \dots \wedge \neg \mathbf{q}_\ell(x_2x_1).$$

Note that, if ζ is primordial, then so is ζ^{-1} . On the other hand, if ζ is evolved, then ζ^{-1} is neither primordial nor evolved. We call 2-types such as ζ^{-1} from the latter case *revolved*. That is to say, a 2-type ζ is revolved if ζ^{-1} is evolved. We differentiate \mathfrak{A} in two steps. First, we construct a model \mathfrak{B} which is differentiated for primordial and evolved 2-types. Having done that, we differentiate revolved 2-types of \mathfrak{B} thus forming a *fully* differentiated model \mathfrak{C} .

Let $\mathfrak{B}_0 := \mathfrak{A}$ be a structure having, without loss of generality, $\mathfrak{q}_1^{\mathfrak{B}}, \dots, \mathfrak{q}_\ell^{\mathfrak{B}}$ all be empty sets. It should be clear that all 2-types realised in \mathfrak{B}_0 are primordial. We define a process for obtaining a transfinite sequence of structures $(\mathfrak{B}_\lambda)_{\lambda \geq 0}$. Picking any ordinal $\lambda \geq 0$ let X_λ be the set of pairs $\langle b, \zeta \rangle$ formed by taking elements $b \in A$ and primordial 2-types ζ satisfying $\mathfrak{B}_\lambda, b \models \exists_{[\leq 1]} x \zeta(xy)$. Similarly, let Y_λ be defined in the same way as X_λ but for evolved 2-types ζ . We will maintain the following properties throughout our construction for ordinals λ :

1. the domain of \mathfrak{B}_λ is A . Moreover, for all \bar{a} over A and all $\text{sig}(\mathfrak{A})$ -formulas ψ we have $\mathfrak{A} \models \psi[\bar{a}]$ if and only if $\mathfrak{B}_\lambda \models \psi[\bar{a}]$;
2. for $\kappa < \lambda$, if \mathfrak{B}_κ is differentiated for primordial 2-types, then $\mathfrak{B}_\lambda = \mathfrak{B}_\kappa$, otherwise $X_\kappa \subsetneq X_\lambda$;
3. for all $b \in A$ and all evolved 2-types ζ , we have $\langle b, \zeta \rangle \in Y_\lambda$.

Take any ordinal $\lambda \geq 0$. We start by setting $\mathfrak{B}_{\lambda+1} := \mathfrak{B}_\lambda$. If \mathfrak{B}_λ is differentiated for primordial and evolved 2-types there is nothing to do. If this is not the case we proceed as follows. Pick an element $b \in A$ and a primordial 2-type ζ such that ζ is undifferentiated for b in \mathfrak{B}_λ . Then, define f to be any injection with domain $U_\zeta := \{a \in A \mid \text{tp}^{\mathfrak{B}_\lambda}[ab] = \zeta\}$ and the co-domain being the powerset of $\{\mathfrak{q}_1, \dots, \mathfrak{q}_\ell\}$. For all $a \in U_\zeta$ and $\mathfrak{q} \in f(a)$ add the tuple ab to $\mathfrak{q}^{\mathfrak{B}_{\lambda+1}}$. Clearly, elements \bar{a} over A satisfy the same $\text{sig}(\mathfrak{A})$ -formulas in $\mathfrak{B}_{\lambda+1}$ as they do in \mathfrak{B}_λ . Thus, condition 1 is secured. Notice now, that $\langle b, \zeta \rangle \in X_{\lambda+1} \setminus X_\lambda$. Since all 2-types created are either evolved or revolved, we have secured condition 2. To see that condition 3 holds first notice that $\text{tp}^{\mathfrak{B}_{\lambda+1}}[ba]$ is revolved for all $a \in U_\zeta$. Thus, $\langle a, \text{tp}^{\mathfrak{B}_{\lambda+1}}[ba] \rangle \notin Y_\kappa$ for any $\kappa \leq \lambda+1$. Secondly, taking any $a \in U_\zeta$ we see that there is no other element $c \in A$ such that $\text{tp}^{\mathfrak{B}_{\lambda+1}}[ab] = \text{tp}^{\mathfrak{B}_{\lambda+1}}[cb]$. Thus, $\langle b, \text{tp}^{\mathfrak{B}_{\lambda+1}}[ab] \rangle \in Y_{\lambda+1}$ for each $a \in U_\zeta$ as required.

If λ is a limit ordinal define $\mathfrak{B}_\lambda := \mathfrak{B}_0$ with $\mathfrak{q}^{\mathfrak{B}_\lambda} := \bigcup_{\kappa < \lambda} \mathfrak{q}^{\mathfrak{B}_\kappa}$ for each \mathfrak{q} amongst $\mathfrak{q}_1, \dots, \mathfrak{q}_\ell$. It is easy to see that $X_\lambda = \bigcup_{\kappa < \lambda} X_\kappa$ whilst $Y_\lambda = Y_\kappa$ for each $\kappa < \lambda$ thus establishing conditions 1, 2 and 3 for \mathfrak{B}_λ .

Having outlined the construction, let us take $(\mathfrak{B}_\lambda)_{\lambda \geq 0}$ to be any sequence of structures obtained as described above. Writing ω_1 for the first uncountable ordinal, we claim that \mathfrak{B}_{ω_1} is differentiated for primordial and evolved 2-types. To see this notice that, by condition 2 of the construction, the class of sets $C := \{X_\lambda \mid \lambda \geq 0\}$ is well-ordered by the subset relation \subseteq . Thus, the order-type of $\langle C, \subseteq \rangle$ is an ordinal, say, λ . Since $X_\kappa \subseteq A \times \text{ATP}_2^{\text{sig}(\mathfrak{B}_0)}$ is countable for any $\kappa \geq 0$, we must have that λ is countable as well. By condition 2, $X_\kappa \subsetneq X_\iota$ for $\kappa < \iota < \lambda$. Thus, again by condition 2, we must have $\mathfrak{B}_\kappa = \mathfrak{B}_\lambda$ for each $\kappa > \lambda$. But this implies that $\mathfrak{B}_{\omega_1} = \mathfrak{B}_\lambda$ is differentiated for primordial 2-types. It is immediate by condition 3 that \mathfrak{B}_{ω_1} is differentiated for evolved 2-types as well.

We now differentiate evolved 2-types of $\mathfrak{B} := \mathfrak{B}_{\omega_1}$ thus forming the advertised fully differentiated structure \mathfrak{C} . It should be clear that each 2-type of \mathfrak{B} is either primordial, evolved, or evolved. Now, recalling that $\ell = \lceil \log_2(3M) \rceil$, let $\mathfrak{w}_1, \dots, \mathfrak{w}_\ell$ be a series of fresh binary predicates. We set $\mathfrak{C} := \mathfrak{B}$ and perform the following action for each $b \in A$ and all evolved 2-types ζ realised in \mathfrak{B} in parallel. Write $U_\zeta := \{a \in A \mid \text{tp}^\mathfrak{B}[ab] = \zeta\}$. In case $|U_\zeta| \in [2, 3M]$ define f to be any injection between U_ζ and the powerset of the new symbols $\{\mathfrak{w}_1, \dots, \mathfrak{w}_\ell\}$. Lastly, add ab to $\mathfrak{w}^\mathfrak{C}$ for each $a \in U_\zeta$ and $\mathfrak{w} \in f(a)$.

We claim that \mathfrak{C} , when constructed as above, is differentiated. To see this let us take any element $b \in A$ and any 2-type ζ realised in \mathfrak{B} . If ζ is primordial, then it is differentiated for b in \mathfrak{B} by condition 2. Since no primordial 2-types are created or altered in the construction of \mathfrak{C} , we must have that ζ is differentiated for b in \mathfrak{C} as well. Suppose now that ζ is a evolved 2-type. Taking any $a \in A$ such that $\text{tp}^\mathfrak{B}[ab] = \zeta$ we claim that $\text{tp}^\mathfrak{C}[ab] \neq \text{tp}^\mathfrak{C}[cb]$ for all $c \in A \setminus \{a\}$. Clearly, if $\text{tp}^\mathfrak{B}[cb] = \zeta$ then cb and ab will satisfy different differentiation predicates $\mathfrak{w}_1, \dots, \mathfrak{w}_\ell$ in \mathfrak{C} . Otherwise, if $\text{tp}^\mathfrak{B}[cb] \neq \zeta$, we cannot have $\text{tp}^\mathfrak{C}[cb] = \text{tp}^\mathfrak{C}[ab]$ by providing interpretations to fresh symbols $\mathfrak{w}_1, \dots, \mathfrak{w}_\ell$. Lastly, suppose that ζ is evolved. By condition 3 in the construction of \mathfrak{B} we have that there is at most one $a \in A$ for which $\text{tp}^\mathfrak{B}[ab] = \zeta$. But since there is no $c \in A \setminus \{a\}$ such that $\text{tp}^\mathfrak{B}[cb] = \zeta$ we again have that $\text{tp}^\mathfrak{C}[ab] = \text{tp}^\mathfrak{C}[cb]$ is unachievable by interpreting $\mathfrak{w}_1, \dots, \mathfrak{w}_\ell$. \square

Taking any model of φ we may assume that it is chromatic and differentiated. Indeed, any model of φ can be expanded into a chromatic model as described in Lemma 7.3.3. It should be clear that if 2-types ζ and η over a common signature differ, then it is impossible to create $\zeta' \supseteq \zeta$ and $\eta' \supseteq \eta$ such that $\zeta' = \eta'$ by introducing new symbols. Thus, by differentiating a chromatic model of φ as

done in Lemma 7.3.4, we obtain a chromatic and differentiated model of φ . For the rest of the section we fix Σ to be the signature of φ along with all the “new” predicates mentioned in Lemmas 7.3.3 and 7.3.4.

Now, recall that an adjacent 3-type ξ is a ray if $\xi \models \beta$. We say that ξ is an *anti-ray* if neither ξ nor ξ^{-1} is a ray. Given 2-types ζ and η , and recalling that $\hat{\alpha} = \alpha(x_1x_2x_3) \wedge \alpha(x_3x_2x_1)$, we write $\zeta \stackrel{\mathcal{L}}{\sim} \eta$ if there is no anti-ray ξ such that:

$$\xi \models \hat{\alpha} \wedge \zeta(x_1x_2) \wedge \eta(x_3x_2).$$

The following lemma is a generalisation of [Pratt-Hartmann, 2010, Lemma 7]. In essence, the lemma states that if, in any $\mathfrak{A} \models \varphi$, the 2-types ζ and η are realised sufficiently many times (relative to some $b \in A$), then there must be an anti-ray:

Lemma 7.3.5. *Take $\mathfrak{A} \models \varphi$ with $b \in A$ and suppose ζ and η are 2-types realised in \mathfrak{A} . Then, if $U_\zeta := \{a \in A \mid \text{tp}^{\mathfrak{A}}(ab) = \zeta\}$ and $V_\eta := \{c \in A \mid \text{tp}^{\mathfrak{A}}(cb) = \eta\}$ are both of cardinality at least $3M+1$, there is some $a \in U_\zeta$ and $c \in V_\eta$ such that $\text{atp}^{\mathfrak{A}}[abc]$ is an anti-ray that contains the literal $\neg \mathbf{e}(\mathbf{x}_3)$.*

Proof. Pick (not necessarily disjoint) sets $U'_\zeta \subseteq U_\zeta$ and $V'_\eta \subseteq V_\eta$ of cardinality exactly $3M+1$. We build two bipartite graphs $G = (U'_\zeta, V'_\eta, E_G)$ and $H = (U'_\zeta, V'_\eta, E_H)$ as follows:

$$E_G := \{ac \in A^2 \mid a \neq c\},$$

$$E_H := \{ac \in A^2 \mid a \neq c \text{ and } \mathfrak{A} \models \beta[abc] \text{ or } \mathfrak{A} \models \beta^{-1}[abc]\}.$$

Clearly, $E_H \subseteq E_G$. Notice that $|E_G| \geq (3M+1)^2 - (3M+1) = 3M(3M+1)$ whilst $|E_H| \leq 2M(3M+1)$. Recalling that $M \geq 1$ we have that the set $E_G \setminus E_H$ is non-empty. Taking any $ac \in E_G \setminus E_H$, we have that $\text{atp}^{\mathfrak{A}}[abc]$ is an anti-ray. Since $a \neq c$ we have that it contains $\neg \mathbf{e}(\mathbf{x}_3)$ as required. \square

We proceed by recalling that φ takes the following form:

$$\forall \mathbf{x}_3 \alpha \wedge \bigwedge_{s \in S} \forall \mathbf{x}_2 \exists_{[=M_s]} x_3 (\beta_s \wedge \neg \mathbf{e}(\mathbf{x}_3) \wedge x_2 \neq x_3).$$

A *star* over Σ and φ is then a partial function σ mapping rays over Σ to cardinalities in $[1, M]$ and satisfying the following conditions:

S0: for each $\xi \in \text{dom}(\sigma)$ we have³ $\xi \models \neg \mathbf{e}(\mathbf{x}_3) \wedge x_2 \neq x_3$,

³This is immediate by the definition of a ray.

S1: there is some 2-type ζ , denoted by $\mathbf{tp}(\sigma)$, s.t. $\zeta = \xi|_{[1,2]}$ for all $\xi \in \text{dom}(\sigma)$,

$$\text{S2: } \sum_{\xi \in \text{dom}(\sigma)}^{\xi \models \beta_s} \sigma(\xi) = M_s \text{ for each } s \in S,$$

S3: $\xi \models \widehat{\alpha}$ for each $\xi \in \text{dom}(\sigma)$.

When considering adjacent 3-types $\xi \notin \text{dom}(\sigma)$ we take $\sigma(\xi)$ to be 0. Given a model \mathfrak{A} of φ , every pair of elements $ab \in A^2$ realises a unique star $\mathbf{str}^{\mathfrak{A}}[ab] = \sigma$ defined for adjacent 3-types ξ as follows:

$$\sigma(\xi) := \begin{cases} |\{c \in A \mid \mathfrak{A} \models \xi[abc]\}| & \text{if } \xi \text{ is a ray and } \mathfrak{A} \models \xi[abc] \text{ for some } c \in A, \\ \text{undefined} & \text{otherwise.} \end{cases}$$

It is easy to verify that $\mathbf{str}^{\mathfrak{A}}[ab]$ satisfies **S0–S3**. In fact, if the model \mathfrak{A} is chromatic, we have that each star σ realised in \mathfrak{A} satisfies the following additional properties:

S4: for each bidirectional $\xi \in \text{dom}(\sigma)$ we have $\sigma(\xi) = 1$,

S5: for each bidirectional $\xi \in \text{dom}(\sigma)$ we have $\xi^{-1}|_{[1,2]} \neq \mathbf{tp}(\sigma)$,

S6: if $\xi, \mu \in \text{dom}(\sigma)$ are distinct bidirectional rays, then $\xi^{-1}|_{[1,2]} \neq \mu^{-1}|_{[1,2]}$.

To verify that these properties hold, let us fix some $b \in A$. By b -semichromaticity there are no elements $a, c \in A$ such that $\mathbf{tp}^{\mathfrak{A}}[ab] = \mathbf{tp}^{\mathfrak{A}}[cb]$ and $\mathfrak{A} \models \widehat{\beta}[abc]$. Thus, **S5** is satisfied by every star in \mathfrak{A} . To establish **S4** and **S6** recall that, by b -chromaticity, there are no elements $a, c, c' \in A$ with $\mathbf{tp}^{\mathfrak{A}}(cb) = \mathbf{tp}^{\mathfrak{A}}(c'b)$ and $\mathfrak{A} \models \widehat{\beta}[abc] \wedge \widehat{\beta}[abc']$. Since every model of φ has a chromatic expansion (Lemma 7.3.3), we will have that stars satisfy **S4–S6** as well. We write **STR** for the set of all stars computed over Σ and φ and satisfying **S0–S6**. Say that a star σ *emits* a ray ξ if $\xi \in \text{dom}(\sigma)$. Similarly, we will say that σ *absorbs* an adjacent 3-type ξ if $\xi^{-1}|_{[1,2]} = \mathbf{tp}(\sigma)$. Before proceeding any further, we provide a quick reference guide (Table 7.3.1) to terminology we used or will define in the future.

Define Σ^+ to be the signature $\Sigma \cup \{\mathbf{r}_\sigma \mid \sigma \in \mathbf{STR}\}$, where each \mathbf{r}_σ is a fresh binary predicate. Intuitively, the satisfaction of \mathbf{r}_σ will indicate that a pair is *promised* to realise the star σ . We compute a $\mathcal{C}_{\mathcal{UHA}}^2$ -sentence ψ over Σ^+ such that

- if $\mathfrak{A} \models \psi$, then \mathfrak{A} can be reconstructed into a model \mathfrak{A}' of φ with the property $\mathfrak{A} \models \mathbf{r}_\sigma[ab] \Leftrightarrow \mathbf{str}^{\mathfrak{A}'}[ab] = \sigma$ for distinct $a, b \in A$.

In regards to 2-types ζ, η and the adjacent 3-type ξ

β	$\bigvee_{s \in S} \beta_s(\mathbf{x}_3) \wedge \neg \mathbf{e}(\mathbf{x}_3) \wedge x_2 \neq x_3$
β^{-1}	$\beta(x_3 x_2 x_1)$
$\widehat{\beta}$	$\beta \wedge \beta^{-1}$
ray	an adjacent 3-type ξ s.t. $\xi \models \beta$
bidirectional ray	a ray ξ s.t. $\xi \models \widehat{\beta}$
unidirectional ray	a ray ξ s.t. $\xi \not\models \widehat{\beta}$
anti-ray	an adjacent 3-type ξ s.t. $\xi \not\models \beta \vee \beta^{-1}$
$\zeta \stackrel{\mathcal{L}}{\sim} \eta$	there is no anti-ray ξ s.t. $\xi \models \widehat{\alpha} \wedge \zeta(x_1 x_2) \wedge \eta(x_3 x_2)$
ATP_i^Σ	the set of all adjacent i -types over $\Sigma \cup \{\mathbf{e}, =\}$

In regards to star σ , 2-type η , and adjacent 3-type ξ

$\text{tp}(\sigma)$	the 2-type ζ s.t. $\zeta = \xi _{[1,2]}$ for each $\xi \in \text{dom}(\sigma)$
σ emits ξ	$\xi \in \text{dom}(\sigma)$
σ absorbs ξ	$\xi^{-1} _{[1,2]} = \text{tp}(\sigma)$
$\sigma _\eta$	σ restricted to rays ξ satisfying $\xi^{-1} _{[1,2]} = \eta$
$\sigma \llbracket \eta \rrbracket$	number of rays ξ emitted by σ satisfying $\xi^{-1} _{[1,2]} = \eta$
$\sigma \stackrel{\beta}{\leftrightarrow} \eta$	there is a bidirectional ray $\xi \in \text{dom}(\sigma)$ s.t. $\xi^{-1} _{[1,2]} = \eta$
STR	the set of all stars compatible with S0–S6

Relative to c, b in \mathfrak{A}

cb realises σ	$\text{str}^\mathfrak{A}[cb] = \sigma$
cb realises $\sigma _\eta$	$\mathfrak{A}, cb \models \exists_{[=\sigma(\xi)]} x_3 \xi$ for all $\xi \in \text{dom}(\sigma _\eta)$
cb is promised σ	$\mathfrak{A} \models \mathbf{r}_\sigma[cb]$

Table 7.3.1: Quick reference guide for Sec 7.3.

- if \mathfrak{A} is a chromatic, differentiated model of φ , then the structure \mathfrak{A}^+ obtained by setting $\mathfrak{A}^+ \models \mathbf{r}_\sigma[ab] \iff \mathbf{str}^{\mathfrak{A}}[ab] = \sigma$ for each $ab \in A^2$ is a model of ψ .

In essence, ψ will be a list of facts about chromatic, differentiated models of φ . For motivational purposes we fix \mathfrak{A} to be any such model of φ and proceed as follows. Define $\mathfrak{A}^+ := \mathfrak{A}$ with $\mathbf{r}_\sigma^{\mathfrak{A}^+} = \{ab \in A^2 \mid \mathbf{str}^{\mathfrak{A}}[ab] = \sigma\}$ for each $\sigma \in \mathbf{STR}$. Notice that, for each $ab \in A^2$, there is exactly one $\sigma \in \mathbf{STR}$ for which $\mathfrak{A}^+ \models \mathbf{r}_\sigma[ab]$. We secure this fact with the following sentence:

$$\forall xy \left(\bigvee_{\sigma \in \mathbf{STR}} (\mathbf{r}_\sigma(xy) \wedge (\mathbf{tp}(\sigma))(xy)) \wedge \bigwedge_{\substack{\sigma \neq \tau \\ \sigma, \tau \in \mathbf{STR}}} (\neg \mathbf{r}_\sigma(xy) \vee \neg \mathbf{r}_\tau(xy)) \right). \quad (\psi_1)$$

To motivate the rest of the construction, let us fix some $b \in A$. Taking some $\zeta \in \mathbf{ATP}_2^\Sigma$ recall that, by b -differentiation, the number of elements $a \in A$ for which $\mathbf{tp}^{\mathfrak{A}}[ab] = \zeta$ is outside the range $[2, 3M]$. We thus have that \mathfrak{A}^+ models:

$$\bigwedge_{\zeta \in \mathbf{ATP}_2^\Sigma} \forall y \left(\exists_{[\leq 1]} x \zeta(xy) \vee \exists_{[\geq 3M+1]} x \zeta(xy) \right). \quad (\psi_2)$$

Proceeding similarly, take some bidirectional ray $\xi \in \mathbf{ATP}_3^\Sigma$ and write:

$$\begin{aligned} U &:= \{\sigma \in \mathbf{STR} \mid \xi \in \text{dom}(\sigma) \text{ and } \mathbf{str}^{\mathfrak{A}}[ab] = \sigma \text{ for some } a \in A\}, \\ V &:= \{\tau \in \mathbf{STR} \mid \xi^{-1} \in \text{dom}(\tau) \text{ and } \mathbf{str}^{\mathfrak{A}}[cb] = \tau \text{ for some } c \in A\}. \end{aligned}$$

We claim that $U' := \{a \in A \mid \mathbf{str}^{\mathfrak{A}}[ab] \in U\}$ and $V' := \{c \in A \mid \mathbf{str}^{\mathfrak{A}}[cb] \in V\}$ are of the same cardinality. To see this, take any $a \in U'$. By **S4**, we have $(\mathbf{str}^{\mathfrak{A}}[ab])(\xi) = 1$. Thus, there is a single element $c \in A$ having $\mathfrak{A} \models \xi[abc]$. In fact, $c \in V'$. Now, taking any $c \in V'$ we have, again by **S4**, that $(\mathbf{str}^{\mathfrak{A}}[cb])(\xi^{-1}) = 1$. Taking $a \in A$ to satisfy $\mathfrak{A} \models \xi^{-1}[cba]$ we have $a \in U'$. Clearly, this establishes a bijection between the sets U' and V' . The following secures the observed equicardinality using uniform H\"artig assertions:

$$\bigwedge_{\xi \in \widehat{\beta}} \forall y \mathbf{I}(x, x) \left(\bigvee_{\sigma \in \mathbf{STR}}^{\xi \in \text{dom}(\sigma)} \mathbf{r}_\sigma(xy), \bigvee_{\tau \in \mathbf{STR}}^{\xi^{-1} \in \text{dom}(\tau)} \mathbf{r}_\tau(xy) \right). \quad (\psi_3)$$

Taking σ to be a star and ζ to be a 2-type over Σ let us write

$$\sigma[[\zeta]] := \sum_{\xi \in \text{dom}(\sigma)}^{\xi^{-1}|_{[1,2]} = \zeta} \sigma(\xi).$$

That is to say, $\sigma[[\zeta]]$ is the number of adjacent 3-types ξ having $\xi^{-1}|_{[1,2]} = \zeta$ that are emitted by σ . We proceed eliminating stars which cannot be realised in any model of φ . Still keeping $b \in A$ fixed take some $\zeta \in \text{ATP}_2^\Sigma$ for which there is no element $a \in A$ having $\text{tp}^{\mathfrak{A}}[ab] = \zeta$. Then there is no $c \in A$ having $(\text{str}^{\mathfrak{A}}[cb])[[\zeta]] \geq 1$. Thus, \mathfrak{A}^+ is a model of:

$$\bigwedge_{\zeta \in \text{ATP}_2^\Sigma} \forall y \left(\left(\exists_{[=0]} x \zeta(xy) \right) \rightarrow \left(\forall x \bigvee_{\sigma \in \text{STR}}^{\sigma[[\zeta]] = 0} \mathbf{r}_\sigma(xy) \right) \right). \quad (\psi_4)$$

On the other hand, suppose that there is a unique element $a \in A$ for which $\text{tp}^{\mathfrak{A}}[ab] = \zeta$. Then, for all $c \in A$, the tuple cb emits at most a single ray ξ such that $\xi^{-1}|_{[1,2]} = \zeta$. It is then immediate that \mathfrak{A}^+ models:

$$\bigwedge_{\zeta \in \text{ATP}_2^\Sigma} \forall y \left(\left(\exists_{[=1]} x \zeta(xy) \right) \rightarrow \left(\forall x \bigvee_{\sigma \in \text{STR}}^{\sigma[[\zeta]] \leq 1} \mathbf{r}_\sigma(xy) \right) \right). \quad (\psi_5)$$

Taking any star σ and a 2-type ζ over Σ let us write

$$\sigma \rightrightarrows \zeta \text{ iff there is some } \xi \in \text{dom}(\sigma) \text{ s.t. } \xi \models \widehat{\beta} \text{ and } \xi^{-1}|_{[1,2]} = \zeta.$$

Still keeping the assumption that $a \in A$ is the unique element having $\text{tp}^{\mathfrak{A}}[ab] = \zeta$ let us write $\sigma := \text{str}^{\mathfrak{A}}[ab]$. Taking any $\eta \in \text{ATP}_2^\Sigma$ we let V_η be the set of elements $c \in A$ satisfying $\text{tp}^{\mathfrak{A}}[cb] = \eta$. Now take V_η^- to be the set of elements $c \in V_\eta$ such that $\mathfrak{A} \models \beta[abc]$. Clearly, $|V_\eta| \geq |V_\eta^-| = \sigma[[\eta]]$. Now, if for some $c \in V_\eta^-$ we have $\mathfrak{A} \models \beta[cba]$, then $\mathfrak{A} \models \widehat{\beta}[cba]$. But, by b -differentiation, there can only one such element $c \in V_\eta$. Thus, we must have $(\text{str}^{\mathfrak{A}}[c'b])[[\zeta]] = 0$ for all $c' \in V_\eta^- \setminus \{c\}$. Informally the above can be read as ab needing enough absorption sites amongst

the tuples cb , where $c \in V_\eta$. The following is then modeled by \mathfrak{A}^+ :

$$\bigwedge_{\zeta, \eta \in \text{ATP}_2^\Sigma} \bigwedge_{\substack{\sigma \in \text{STR} \\ \text{tp}(\sigma) = \zeta}} \forall y \left(\left(\exists_{[=1]} x \zeta(xy) \wedge \exists x \mathbf{r}_\sigma(xy) \right) \rightarrow \left(\exists_{[\geq \sigma[\zeta]]} x \bigvee_{\substack{\tau \in \text{STR} \\ \text{tp}(\tau) = \eta}}^{\tau[\zeta] = 0 \text{ or } \tau \Rightarrow \zeta} \mathbf{r}_\tau(xy) \right) \right). \quad (\psi_6)$$

Let us now take ζ and η to be any 2-types over Σ . Suppose that there is no anti-ray that entails the formula $\hat{\alpha} \wedge \zeta(x_1x_2) \wedge \eta(x_3x_2)$. That is to say, $\zeta \stackrel{\mathcal{L}}{\sim} \eta$. Keeping $b \in A$ fixed we write:

$$U_\zeta := \{a \in A \mid \text{tp}^{\mathfrak{A}}[ab] = \zeta\}, \\ V_\eta := \{c \in A \mid \text{tp}^{\mathfrak{A}}[cb] = \eta\}.$$

Since \mathfrak{A} is differentiated, we have that both sets are of cardinality at most 1 or at least $3M+1$. But, by Lemma 7.3.5, at least one of U_ζ, V_η must be of cardinality at most 1. Thus, \mathfrak{A}^+ models the following:

$$\bigwedge_{\zeta, \eta \in \text{ATP}_2^\Sigma}^{\zeta \stackrel{\mathcal{L}}{\sim} \eta} \forall y \left(\left(\exists_{[\leq 1]} x \zeta(xy) \right) \vee \left(\exists_{[\leq 1]} x \eta(xy) \right) \right). \quad (\psi_7)$$

Let us keep $\zeta \stackrel{\mathcal{L}}{\sim} \eta$ together with U_ζ and V_η as before. Take U_ζ^- to be the set of elements $a \in U_\zeta$ for which ab does not send unidirectional rays to cb for any $c \in V_\eta$. But then, by definition of $\zeta \stackrel{\mathcal{L}}{\sim} \eta$, $\mathfrak{A} \models \beta[cba]$ for each $c \in V_\eta$ and $a \in U_\zeta^-$. Since no pair of elements in \mathfrak{A} emits more than M rays, we must then have $|U_\zeta^-| \leq M$ or $|V_\eta| = 0$. It is then easy to verify that \mathfrak{A}^+ is a model of:

$$\bigwedge_{\zeta, \eta \in \text{ATP}_2^\Sigma}^{\zeta \stackrel{\mathcal{L}}{\sim} \eta} \forall y \left(\left(\exists_{[\leq M]} x \bigvee_{\substack{\sigma \in \text{STR} \\ \text{tp}(\sigma) = \zeta}}^{\sigma[\eta] = 0, \text{ or } \sigma[\eta] = 1 \text{ and } \sigma \Rightarrow \eta} \mathbf{r}_\sigma(xy) \right) \vee \left(\exists_{[=0]} x \eta(xy) \right) \right). \quad (\psi_8)$$

Let us dissect the case $|U_\zeta^-| \in [1, M]$. As argued before, $\mathfrak{A} \models \beta[cba]$ for each $c \in V_\eta$ and $a \in U_\zeta^-$. Thus, $(\text{str}^{\mathfrak{A}}[cb])[\zeta] = |U_\zeta^-|$ for each $c \in V_\eta$. It is then

immediate that \mathfrak{A}^+ is a model of the following:

$$\bigwedge_{\zeta, \eta \in \text{ATP}_2^\Sigma} \zeta \stackrel{\circ}{\sim} \eta \quad \bigwedge_{k=1}^{k \leq M} \forall y \left(\left(\exists_{[=k]} x \quad \bigvee_{\substack{\sigma \in \text{STR} \\ \text{tp}(\sigma) = \eta}}^{\sigma[[\zeta]=0, \text{ or} \\ \sigma[[\zeta]=1 \text{ and } \sigma \Rightarrow \zeta]} \mathbf{r}_\sigma(xy) \right) \rightarrow \right. \\ \left. \forall x \left(\zeta(xy) \rightarrow \bigvee_{\substack{\sigma \in \text{STR} \\ \text{tp}(\sigma) = \zeta}}^{\sigma[[\eta]=k]} \mathbf{r}_\sigma(xy) \right) \right). \quad (\psi_9)$$

We conclude the definition of the $\mathcal{C}_{\mathcal{UHA}}^2$ -sentence $\psi := \psi_1 \wedge \dots \wedge \psi_9 \wedge \text{acl}(\varphi)$ having achieved the following:

Lemma 7.3.6. *Suppose $\mathfrak{A} \models \varphi$ is chromatic and differentiated. Then, $\mathfrak{A}^+ \models \psi$.*

As promised, we now proceed with the converse direction:

Lemma 7.3.7. *Suppose $\mathfrak{A} \models \psi$. Then, \mathfrak{A} can be reconstructed into $\mathfrak{A}' \models \varphi$.*

Proof. Let \mathfrak{A}^- be the Σ -reduct of \mathfrak{A} . We assume, without loss of generality, that \mathfrak{A}^- is of primitive height 2; that is to say if $\mathfrak{p} \in \Sigma$, then $\bar{w} \in \mathfrak{p}^{\mathfrak{A}^-}$ implies that \bar{w} is a word over some binary alphabet. We think of primitive triples over A as having undefined adjacent 3-types and construct $\mathfrak{A}' := \mathfrak{A}^-$ by specifying them. In particular, we follow the ‘‘advice’’ of $\mathbf{r}_\sigma \in \Sigma^+ \setminus \Sigma$ and guarantee that $\text{str}^{\mathfrak{A}'}[ab] = \sigma \Leftrightarrow \mathfrak{A} \models \mathbf{r}_\sigma[ab]$ for all $\sigma \in \text{STR}$ and primitive $ab \in A^2$. Note that the above notion is well-defined as, by ψ_1 , we have that there is exactly one $\sigma \in \text{STR}$ such that $\mathfrak{A} \models \mathbf{r}_\sigma[ab]$. We call σ the star promised to ab .

Step 1: Bidirectional rays.

For the whole of the construction let us fix some $b \in A$. We extend \mathfrak{A}' in stages, firstly, by assigning bidirectional rays for primitive triples of elements. To this end, take ξ to be some bidirectional ray over Σ and write:

$$U := \{a \in A \mid ab \text{ is promised a star } \sigma \text{ s.t. } \sigma(\xi) = 1\}, \\ V := \{c \in A \mid cb \text{ is promised a star } \tau \text{ s.t. } \sigma(\xi^{-1}) = 1\}.$$

By ψ_3 the sets U and V are of equal cardinality. Thus, we may take $f : U \rightarrow V$ to be any bijection and set $\text{atp}^{\mathfrak{A}'}[abf(a)] := \xi$ for each $a \in U$. To see that this is well-defined notice that $U \neq V$ by **S5**. Moreover, since $\widehat{\beta} \models x_1 \neq x_2 \wedge x_2 \neq x_3$,

we have that $b \notin U \cup V$. Thus, the triples considered in this procedure are, as required, primitive. We claim that no clashes are introduced by repeating the assignment for another bidirectional ray, say $\mu \neq \xi$. To see this, take some primitive $ab \in A^2$. Supposing that $\mu, \xi \in \text{dom}(\sigma)$, where σ is the star promised to ab , we have, by S6, that $\mu|_{[2,3]} \neq \xi|_{[2,3]}$. Thus, if for some $c \in A$ the tuple abc was assigned the bidirectional ray ξ in \mathfrak{A}' , we have that $\text{tp}^{\mathfrak{A}'}[cb] \neq \mu|_{[2,3]}$ thus eliminating the possibility of a double assignment.

Assume now, that the procedure above has been carried out for all bidirectional rays. Property S4 gives us that, for each $a \in A$, the tuple ab in \mathfrak{A}' emits all the bidirectional rays of the star promised to it.

Step 2: Assigning unidirectional rays.

Taking a star σ and a 2-type η let us write $\sigma|_{\eta}$ for the partial function defined on rays ξ as follows:

$$\sigma|_{\eta}(\xi) := \begin{cases} \sigma(\xi) & \text{if } \xi^{-1}|_{[1,2]} = \eta, \\ \text{undefined} & \text{otherwise.} \end{cases}$$

In the sequel we will say that a pair of elements $ab \in A^2$ realise the partial star $\sigma|_{\eta}$ just in case⁴ $\sigma|_{\eta}(\xi) = |\{c \in A \mid \text{atp}^{\mathfrak{A}'}[abc] = \xi\}|$ for each $\xi \in \text{dom}(\sigma|_{\eta})$. Clearly, if ab realises all partial stars $\{\sigma|_{\eta} \mid \eta \in \text{ATP}_2^{\Sigma}\}$, then ab realises the star σ .

This part of the construction is executed by taking 2-types over Σ in pairs. We thus fix ζ and η to be not necessarily distinct 2-types, both of which contain the literal $x_1 \neq x_2$, and write

$$\begin{aligned} U_{\zeta} &:= \{a \in A \mid \mathfrak{A}, \text{tp}^{\mathfrak{A}^-}(ab) = \zeta\}, \\ V_{\eta} &:= \{c \in A \mid \mathfrak{A}, \text{tp}^{\mathfrak{A}^-}(cb) = \eta\}. \end{aligned}$$

Clearly, neither $\text{tp}^{\mathfrak{A}^-}[bb] = \zeta$ nor $\text{tp}^{\mathfrak{A}^-}[bb] = \eta$, thus $b \notin U_{\zeta} \cup V_{\eta}$. The current step is further divided into 4 cases based on the properties of ζ , η , U_{ζ} and V_{η} .

Step 2.1: When a 2-type is not realised.

Let the 2-types ζ , η and sets U_{ζ} , V_{η} be as described above. For the current step suppose one of the sets, say U_{ζ} , is empty. Taking any $c \in V_{\eta}$ we have, by ψ_4 , that

⁴Recall that we allow $\text{atp}^{\mathfrak{A}'}[abc]$ to be undefined.

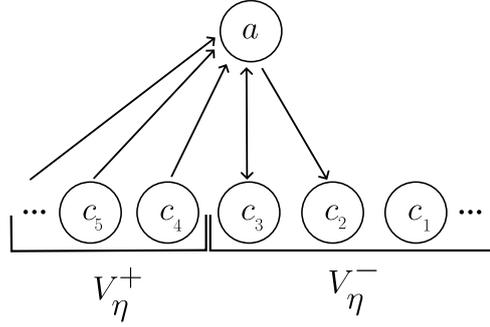


Figure 7.3.1: Possible configuration after Step 2.3 (with element b omitted). Unidirectional lines indicate unidirectional rays. Bidirectional lines are bidirectional rays. Note that elements (e.g. c_1) need not receive or emit rays.

cb is promised some star σ s.t. $\sigma[\zeta] = 0$. Thus, in this case, there is nothing to do as cb already realises the partial star $\sigma|_\zeta$ in \mathfrak{A}' . The case $V_\eta = \emptyset$ is analogous.

Step 2.2: When $\zeta = \eta$ is realised exactly once.

Assume that $\zeta = \eta$ whilst keeping the sets U_ζ and V_η as before. Clearly, $U_\zeta = V_\eta$. Proceeding with the assumption that $|U_\zeta| = 1$ we claim that there is nothing to do as, taking $a \in U_\zeta$ and writing σ for the star promised to ab , we have that ab already realises the partial star $\sigma|_\zeta$. Indeed, if $\sigma[\zeta] \geq 1$, then, by ψ_6 , we have that $|U_\zeta| \geq 2$ or $\sigma \rightrightarrows \zeta$. Immediately, the former contradicts our initial assumption that $|U_\zeta| = 1$. But it cannot be the latter as, by S5, there is no bidirectional ray $\xi \in \text{dom}(\sigma)$ having $\xi^{-1}|_{[1,2]} = \text{atp}(\sigma) = \zeta$. Thus, $\sigma[\zeta] = 0$ as required.

Step 2.3: When $\zeta \neq \eta$ with ζ realised exactly once.

Let us keep the 2-type ζ and η , and the sets U_ζ and V_η . This time, however, assume that $\zeta \neq \eta$ whilst $|U_\zeta| = 1$ and $|V_\eta| \geq 1$. Since U_ζ and V_η are disjoint, every tuple abc formed from $a \in U_\zeta$ and $c \in V_\eta$ is primitive. Taking any $c \in V_\eta$, the sentence ψ_5 ensures that the star τ promised to cb has $\tau[\zeta] \leq 1$. We may thus divide the set V_η into:

$$\begin{aligned} V_\eta^- &:= \{c \in V \mid \tau \text{ is promised to } cb \text{ and } \tau[\zeta] = 0 \text{ or } \tau \rightrightarrows \zeta\}, \text{ and} \\ V_\eta^+ &:= \{c \in V \mid \tau \text{ is promised to } cb \text{ and } \tau[\zeta] = 1, \text{ but not } \tau \rightrightarrows \zeta\}. \end{aligned}$$

Let us take any $c \in V_\eta$ and write τ for the star promised to cb . Supposing,

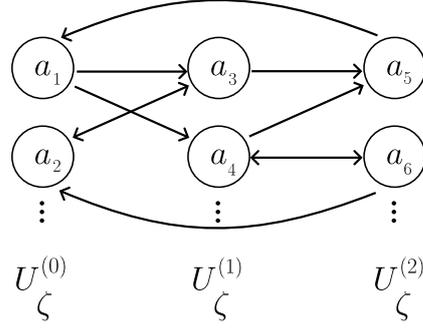


Figure 7.3.2: Possible configuration after Step 2.4 (with element b omitted). Unidirectional lines indicate unidirectional rays. Bidirectional lines are bidirectional rays. Note that unidirectional rays are only emitted from left to right (with elements of $U_\zeta^{(2)}$ wrapping around).

first, that $c \in V_\eta^+$, let ξ be the unique ray satisfying $\sigma(\xi) \neq 0$. By definition, ξ is unidirectional. Thus, we are free to set $\mathbf{atp}^{\mathfrak{A}'}[cba] := \xi$ for the singleton $a \in U_\zeta$. Having done this, the tuple cb now realises the partial star $\tau|_\eta$, whilst leaving the star realised by ab unaltered.

Now, let us suppose that $c \in V_\eta^-$. We claim that cb already realises the partial star $\tau|_\zeta$. This is immediate if $\tau[\zeta] = 0$. In case $\tau[\zeta] = 1$, then, by definition, we have $\tau \rightrightarrows \zeta$. But, by step 1, it is already the case that $\mathbf{atp}^{\mathfrak{A}'}[cba] = \xi$, where a is the singleton of U_ζ and ξ is the unique (bidirectional) ray satisfying $\tau|_\zeta(\xi) \neq 0$.

We may process each $c \in V_\eta$ as given above and conclude that cb realises the partial star $\tau|_\zeta$. Note that the star realised by ab , where $a \in U_\zeta$ is untouched.

We are left to deal with the singleton $a \in U_\zeta$ and the star σ promised to ab . By ψ_6 , $|V_\eta^-| \geq \sigma[\eta]$. Thus, there are enough absorption sites for rays of $\sigma|_\eta$. Note that, as we saw in the previous paragraph, if $\sigma \rightrightarrows \eta$, then there already is an element $c' \in V_\eta^-$ for which $\mathbf{atp}^{\mathfrak{A}'}[abc']$ is a bidirectional ray. Thus, for each $c \in V_\eta^- \setminus \{c\}$ we assign $\mathbf{atp}^{\mathfrak{A}'}[abc]$ to either be undefined or a unidirectional ray of $\text{dom}(\sigma_\eta)$. Since $|V_\eta^-| \geq \sigma[\eta]$, we may pick an assignment which results in ab realising the partial star $\sigma|_\eta$. Since only unidirectional rays are set here, we conclude that, for $c \in V_\eta^-$, the star realised by cb remains unaltered.

A possible configuration obtained at this step is shown in Figure 7.3.1.

Step 2.4: When ζ, η are both realised at least $3M+1$ times.

Lastly, keeping the 2-type ζ and η , and the sets U_ζ and V_η as before, we consider the case where $|U_\zeta| \geq 3M+1$ and $|V_\eta| \geq 3M+1$. First, we partition U_ζ into mutually disjoint sets $U_\zeta^{(0)}, U_\zeta^{(1)}, U_\zeta^{(2)}$ and partition V_η into mutually disjoint sets $V_\eta^{(0)}, V_\eta^{(1)}, V_\eta^{(2)}$. Since $|U_\zeta| \geq 3M+1$ and $|V_\eta| \geq 3M+1$, we may assume that the newly formed sets are each of cardinality at least M . We assume, for simplicity, that $\zeta \neq \eta$. The case $\zeta = \eta$ is handled similarly but by having $U_\zeta^{(i)} = V_\zeta^{(i)}$ for all $i \in [0, 2]$. We proceed with an adaptation of the circular witnessing technique encountered in Lemmas 2.2.1 and 5.4.5.

Let us fix some $i \in [0, 2]$ and define $+_3$ to be addition modulo 3. Taking any $a \in U_\zeta^{(i)}$ and writing σ for the star promised to ab we have $\sigma \llbracket \eta \rrbracket \leq M \leq |V_\eta^{(i+31)}|$. That is to say, there are enough absorption sites in $V_\eta^{(i+31)}$ for rays required by $\sigma|_\eta$. If it is *not* the case that $\sigma \Rightarrow \eta$, then we set $\text{atp}^{2'}[abc]$ for $c \in V_\eta^{(i+31)}$ to be undefined or a unidirectional ray from $\text{dom}(\sigma|_\eta)$. Again, we may assign the rays in such a way that results in ab realising the partial star $\sigma|_\eta$. Turning to the case where $\sigma \Rightarrow \eta$, we have, by our construction in step 1, that there already is some element $c' \in V_\eta$ for which $\text{atp}^{2'}[abc']$ is set to be a bidirectional ray. Thus, we need only deal with unidirectional rays of σ of which there are $\sigma \llbracket \eta \rrbracket - 1$. Clearly, there are enough elements in $V_\eta^{(i+31)} \setminus \{c'\}$ to do just that. Since only unidirectional rays are assigned here, the stars realised by cb of $c \in V_\eta$ are unaltered. By repeating this process for each $i \in [0, 2]$ and all $a \in U_\zeta^{(i)}$ we will have that ab realises the partial star $\sigma|_\eta$, where σ is the star promised to ab . It should be clear that the triples considered here are primitive. Indeed, even if $\zeta = \eta$ we have $U_\zeta^{(i)} \neq V_\eta^{(i+31)}$. Taking any $i \in [0, 2]$, $c \in V_\eta^{(i)}$ and $a \in U_\zeta^{(i+31)}$ we see that $\text{atp}^{2'}[cba]$ was not assigned an adjacent 3-type by the above procedure. Since $|U_\zeta^{(i+31)}| \geq M$ we repeat the above operations for each $i \in [0, 2]$ and $c \in V_\eta^{(i)}$ it being understood that witnesses are picked from $U_\zeta^{(i+31)}$. Then, for each $c \in V_\eta$, we have that cb realises the partial star $\tau|_\zeta$, where τ is the star promised to cb . The stars of ab for $a \in U_\zeta$ are, again, unaltered by this procedure.

A possible configuration obtained after this step is presented in Figure 7.3.2. In particular, possible relations between the elements of U_ζ are shown.

Step 3: Assigning anti-rays.

Suppose now that the steps 1–2.4 are replicated for each $b \in A$ and every unordered pair of 2-types ζ, η . Thus, for each primitive $ab \in A^2$ we have

$\text{str}^{\mathfrak{A}'}[ab] = \sigma$, where σ is the star promised to ab . It might be the case, however, that some primitive triple of elements $abc \in A^3$ has not been assigned an adjacent 3-type. In that case we proceed as follows. Let us write $\zeta := \text{atp}^{\mathfrak{A}'}[ab]$ and $\eta := \text{atp}^{\mathfrak{A}'}[cb]$. If it is *not* the case that $\zeta \mathcal{L} \eta$, then there is an anti-ray ξ having $\xi \models \widehat{\alpha} \wedge \zeta(x_1x_2) \wedge \eta(x_3x_2)$. We claim, in addition, that ξ can be picked such that $\xi \models \neg \mathbf{e}(\mathbf{x}_3)$. This is immediate if $\zeta \neq \eta$ as it implies $x_1 \neq x_3$. On the other hand, if $\zeta = \eta$, then either $a = c$ or $|\{c \in A \mid \text{tp}^{\mathfrak{A}'}[cb] = \zeta\}| \geq 3M+1$. The former case contradicts primitiveness of abc ; as for the latter we have $\xi \models \neg \mathbf{e}(\mathbf{x}_3)$ by referencing Lemma 7.3.5. Thus, by setting $\text{atp}^{\mathfrak{A}'}[abc] := \xi$ we provide an adjacent 3-type to abc without altering the stars of ab and cb in \mathfrak{A} .

Now, keeping $abc \in A^3$ as before, suppose that $\zeta \mathcal{L} \eta$. We claim that $\text{atp}^{\mathfrak{A}'}[abc]$ has already been set to a ray, thus contradicting our initial assumption. Write:

$$\begin{aligned} U_\zeta &:= \{a \in A \mid \mathfrak{A}, \text{tp}^{\mathfrak{A}'}(ab) = \zeta\}, \\ V_\eta &:= \{c \in A \mid \mathfrak{A}, \text{tp}^{\mathfrak{A}'}(cb) = \eta\}. \end{aligned}$$

By ψ_7 , we have $|U_\zeta| \leq 1$ or $|V_\eta| \leq 1$. Clearly, if $\zeta = \eta$, then abc is not primitive. Suppose then, that the 2-types are distinct and, without loss of generality, that $|U_\zeta| = 1$ whilst $|V_\eta| \geq 1$. Fix a to be the singleton element of U_ζ and define

$$\begin{aligned} V_\eta^- &:= \{c' \in V \mid \text{str}^{\mathfrak{A}'}[c'b] \text{ and } (\text{str}^{\mathfrak{A}'}[c'b])[\zeta] = 0 \text{ or } \text{str}^{\mathfrak{A}'}[c'b] \Rightarrow \zeta\}, \text{ and} \\ V_\eta^+ &:= \{c' \in V \mid \text{str}^{\mathfrak{A}'}[c'b] \text{ and } (\text{str}^{\mathfrak{A}'}[c'b])[\zeta] = 1, \text{ but not } \text{str}^{\mathfrak{A}'}[c'b] \Rightarrow \zeta\}. \end{aligned}$$

Notice that, by ψ_8 and the assumption that U_ζ is non-empty, we have $|V_\eta^-| \leq M$ and thus, by ψ_9 , $(\text{str}^{\mathfrak{A}'}[ab])[\eta] = |V_\eta^-|$. In case $c \in V_\eta^-$, then, by step 2.3 of our construction, $\text{atp}^{\mathfrak{A}'}[abc]$ was set to be a ray. On the other hand, if $c \in V_\eta^+$ then, by the same step, we have set $\text{atp}^{\mathfrak{A}'}[cba]$ to be a ray. In both cases $\text{atp}^{\mathfrak{A}'}[abc]$ is already defined thus contradicting our initial assumption.

The last case of our assignment (or lack there of) is illustrated in Figure 7.3.3.

Correctness of the construction.

Let us now verify that $\mathfrak{A}' \models \varphi$. Recall that φ is of the form:

$$\forall \mathbf{x}_3 \alpha \wedge \bigwedge_{s \in S} \forall \mathbf{x}_2 \exists_{[=M_s]} x_3 (\beta_s \wedge \neg \mathbf{e}(\mathbf{x}_3) \wedge x_2 \neq x_3).$$

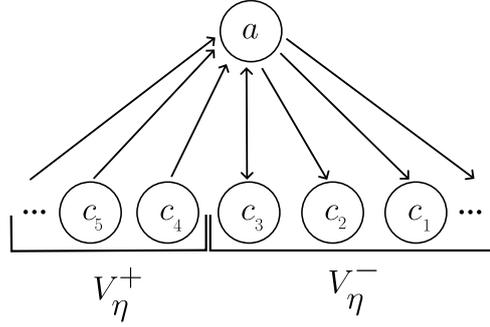


Figure 7.3.3: Assigning anti-rays in step 3 when $\zeta \lesssim \eta$ (with element b omitted). Unidirectional lines indicate unidirectional rays. Bidirectional lines are bidirectional rays. When this occurs there is nothing more to set.

Taking any tuple $abc \in A^3$ we claim that $\mathfrak{A}' \models \alpha[abc]$. Suppose, first, that abc is not primitive and let \bar{w} be its primitive generator. Then, by $\text{acl}(\varphi)$, $\mathfrak{A} \models \alpha[\bar{w}^f]$ for every $f : [1, 3] \rightarrow [1, 2]$ and thus $\mathfrak{A} \models \alpha[abc]$. Since α does not feature symbols outside $\Sigma \cup \{=, \mathbf{e}\}$, we have $\mathfrak{A}^- \models \alpha[abc]$. Additionally, when constructing \mathfrak{A}' from \mathfrak{A}^- , we did not alter adjacent 3-types of non-primitive triples. Thus, $\mathfrak{A}' \models \alpha[abc]$ as required. If abc is primitive, then $\xi := \text{atp}^{\mathfrak{A}'}[abc]$ is either an anti-ray satisfying $\xi \models \hat{\alpha}$ or a ray. If ξ is indeed a ray, we have that ξ or ξ^{-1} is emitted by a star. But then, by **S3**, we have $\xi \models \hat{\alpha}$ as required.

To see that the existential requirements are met pick some $ab \in A^2$. Suppose, first, that $a = b$. Then, by $\text{acl}(\varphi)$, we have $\mathfrak{A}, a \models \exists_{[=M_s]} x_2 (\beta_s(x_1 x_1 x_2) \wedge x_1 \neq x_2)$ for each $s \in S$. But this is equivalent to saying $\mathfrak{A}, aa \models \exists_{[=M_s]} x_3 (\beta_s(\mathbf{x}_3) \wedge \neg \mathbf{e}(\mathbf{x}_3) \wedge x_2 \neq x_3)$. Using similar arguments as before, we already have that the existential requirements are met by aa in \mathfrak{A}' . If $a \neq b$, then we see that $\sigma := \text{str}^{\mathfrak{A}'}[ab]$ is the star promised to ab . Fixing any $s \in S$ we have, by **S2**, that $\sum_{\xi \in \text{dom}(\sigma)}^{\xi \models \beta_s} \sigma(\xi) = M_s$. Thus, $\mathfrak{A}, ab \models \exists_{[=M_s]} x_3 (\beta_s(\mathbf{x}_3) \wedge \neg \mathbf{e}(\mathbf{x}_3) \wedge x_2 \neq x_3)$ as required. \square

We are now at a position to present a decision procedure for \mathcal{AFC}^3 -sentences such as φ . First, writing M_{\max} for the maximal numeric subscript in φ , guess a $\text{sig}(\varphi)$ -structure \mathfrak{A} of size at most $M_{\max} + 2$. If $\mathfrak{A} \models \varphi$ accept. Otherwise, compute a normal-form formula φ' from φ as given in Lemma 7.3.2. It is guaranteed that φ' is satisfiable over domains of size at least $M_{\max} + 3$ if and only if φ is. Now, from φ' , compute the $\mathcal{C}_{\mathcal{UHA}}^2$ -sentence ψ as given just before Lemma 7.3.6. By Theorem 7.2.6 the satisfiability status of ψ can be effectively checked. Combining Lemmas 7.3.6 and 7.3.7 we see that ψ is satisfiable if and only if φ' is. Hence:

Theorem 7.3.8. *The satisfiability problem for \mathcal{AFC}^3 is RECURSIVE.*

Notice that the construction outlined in Lemmas 7.3.6 and 7.3.7 does not alter the domains of the original models. As a consequence the original formula φ has a finite model of size at least $M_{\max}+3$ if and only if the computed formula ψ has one as well. Since $\text{FinSat}(\mathcal{AFC}^3)$ is Σ_1^0 -complete (Theorem 6.1.2), we have:

Corollary 7.3.9. *The finite satisfiability problem for $\mathcal{C}_{\mathcal{UHLA}}^2$ is Σ_1^0 -complete.*

Chapter 8

Conclusions

Our contribution in this body of work is twofold. We began with by considering the satisfiability problem for the fluted fragment. In Chapter 4 we showed that \mathcal{FLPC} has a (finite) satisfiability problem that is TOWER-complete (Theorem 4.3.5). When restricted to $\ell \geq 2$ variables, we showed that the complexity drops to $(\ell-1)$ -NEXPTIME (Theorem 4.3.4). A crucial aspect of our proof is a new observation concerning satisfiable fluted sentences. In Lemma 4.2.1 we demonstrated that 2-variable fluted sentences with periodic counting admit homogenous models; i.e. structures in which elements of the same fluted 1-type have the same fluted 1-profile (read: behave similarly). An analogous property was established for the multivariable fragment in Lemma 4.3.1. The significance of homogeneity in the studies of fluted languages does not end here. Having corollary 4.4.1 it is likely that (some form of) homogeneity is present in extensions of fluted languages and thus description logics lacking role inverses (e.g. \mathcal{ALCHOQ}). The exact reach of our results is left for future studies.

Notwithstanding our new techniques, we are still left with open problems concerning the complexity of (finite) satisfiability for the fluted fragment. It was shown in [Pratt-Hartmann et al., 2019] that the fluted fragment with $\ell \geq 2$ variables has a (finite) satisfiability problem which is $\lfloor \ell/2 \rfloor$ -NEXPTIME-hard. At the moment of writing, this is the best complexity lower-bound for satisfiability problems of \mathcal{FL} , \mathcal{FLPC} and even \mathcal{AF} . Immediately, we have the complexity lower- and upper- bound does not coincide for $\text{Sat}(\mathcal{FLPC}^\ell)$ and $\text{FinSat}(\mathcal{FLPC}^\ell)$ for ℓ as small as 3. The same problem persists for $\text{Sat}(fl^k)$ ($= \text{FinSat}(fl^k)$) with $k \geq 5$. It is not clear whether the upper- or lower-bound (or both) is lax.

Our second result is that concerning the adjacent fragment of first-order logic

– a new language introduced in Chapter 5. We showed that satisfiable \mathcal{AF} -sentences always have finite models and, in Theorem 5.5.4, established TOWER-completeness for $\text{Sat}(\mathcal{AF})$ ($= \text{FinSat}(\mathcal{AF})$). When considering the ℓ -variable fragment we again have that $\text{Sat}(\mathcal{AF})$ ($= \text{FinSat}(\mathcal{AF})$) is $\lfloor \ell/2 \rfloor$ -NEXPTIME-hard, but in $(\ell-1)$ -NEXPTIME for $\ell \geq 2$ (lower-bound inherited from \mathcal{FL} , upper-bound in Theorem 5.5.4). When equality is disallowed, the complexity bounds are the same as for the fluted fragment (also without equality). More specifically, we showed in Theorem 5.4.7 that $\text{Sat}(\mathcal{af}^k)$ ($= \text{FinSat}(\mathcal{af}^k)$) is in $(k-2)$ -NEXPTIME for $k \geq 3$. It is again unknown which bound is non-optimal.

We have identified that the adjacent fragment is a maximal argument-sequence logic that has a decidable (finite) satisfiability problem (Theorem 5.7.2). Additionally, with Corollary 5.5.5 we showed that \mathcal{AF} with binary relations that have predetermined semantics has a decidable (finite) satisfiability problem if and only if the (finite) satisfiability problem is decidable for \mathcal{FO}^2 with the same relations. We thus turned to syntactic extensions for \mathcal{AF} in the form of counting and periodic counting quantifiers. With Theorem 6.1.2 we showed that $\text{FinSat}(\mathcal{AFC}^\ell)$, $\text{FinSat}(\mathcal{AFP}^\ell)$, and $\text{Sat}(\mathcal{AFC}^\ell)$ are undecidable for $\ell \geq 3$. This contrasts our results in Chapter 4 along with decidability results for (finite) satisfiability problems of \mathcal{C}^2 [Pratt-Hartmann, 2010] and $\mathcal{FO}_{\text{Pres}}^2$ [Benedikt et al., 2024]. With Theorem 6.2.4 we established undecidability of $\text{Sat}(\mathcal{AF}^k)$ for $k \geq 4$. Perhaps our most intriguing result is that concerning $\text{Sat}(\mathcal{AF}^3)$. With Theorem 7.3.8 we showed that the problem is decidable, albeit, without a complexity-theoretic upper-bound. Finding this bound is currently open. One of the major difficulties in establishing a decision procedure is that it must fail to determine finite satisfiability. Nonetheless, we provided a satisfactory bound on the *limits of decision*.

Note that the undecidability reductions found in Chapter 6 produce unguarded formulas. Additionally, the undecidability reduction used for the 3-variable guarded fragment with counting in [Grädel, 1999] does not produce adjacent formulas. It is, as of writing, unknown whether the guarded adjacent fragment with counting quantifiers has a decidable (finite) satisfiability problem. The importance of this problem is elevated when having description logics in mind. A positive decidability result would allow us to extend description logics with role inverses and counting (e.g. \mathcal{ALCHIQ}) in a multi-variable fashion – a feature that is unobtainable (without losing decidability of satisfiability) when considering multi-variable guarded formulas as extensions to description logics.

Another interesting problem applicable to formal languages \mathcal{L} is that of *spectra*; that is: the problem of determining the sizes of finite models that sentences of \mathcal{L} can have. (See [Durand et al., 2012] for a survey of the problem). It is known that the spectra of \mathcal{FO} -sentences is exactly the class NEXPTIME [Jones and Selman, 1974]. It is likely that the spectra of \mathcal{AFC} is similar. For one, using our construction in Section 6.2, one can axiomatise grids with rows and columns corresponding to tape contents and time passed in a run of a non-deterministic Turing machine. The only difficulty left is encoding the input string on the tape. Nonetheless, we believe the following to be likely:

Conjecture 8.0.1. *For each set of strings S over $\{0, 1\}$ that is recognised by a non-deterministic Turing machine in time $f(n) = 2^n$, there is an \mathcal{AFC}^4 -sentence with the spectrum $\{n^2 \mid n \text{ is the integer value of some } \bar{b} \in S\}$.*

Note that our speculation above does not apply to the spectra of \mathcal{AFC}^3 as the encoding of grids found in Section 6.2 operates on four variables. We thus consider alternative modes of computation. Let us take $D(x_1, \dots, x_\ell) = 0$ to be a Diophantine equation in the variables x_1, \dots, x_ℓ for some $n \geq 1$. Now, taking any polynomial $p(x)$ define $S_{D,p}$ to be the following set of natural numbers:

$$\{a \in \mathbb{N} \mid \exists b_2, \dots, b_\ell \in \mathbb{N} \text{ s.t. } D(a, b_2, \dots, b_m) = 0 \text{ and } b_2, \dots, b_m \leq 2^{p(\|a\|)}\},$$

where $\|a\|$ denotes the number of bits needed to encode the value of a . Following the definition of L. Adleman and K. Manders in [Adleman and Manders, 1976], the class \mathbf{D} is then the class of sets below:

$$\{S_{D,p} \mid D(x_1 \dots, x_\ell) = 0 \text{ is a Diophantine equation and } p(x) \text{ is a polynomial}\}.$$

It is known that $\mathbf{D} \subseteq \text{NPTIME}$, the converse (i.e. $\text{NPTIME} \subseteq \mathbf{D}$) is still an open problem. Now, recall that in Section 6.1 we constructed a sentence that has a model if and only if a given set of *simple* Diophantine equations has a solution. We believe our reduction could be used to show the following:

Conjecture 8.0.2. *For each Diophantine equation $D(x_1 \dots, x_\ell) = 0$ there is an \mathcal{AFC}^3 -sentence with the spectrum $\{n^k \mid n \in S_{D,p}\}$, where $k \geq 1$ and $p(x) = x$.*

One technical challenge to overcome is that of transforming Diophantine equations into *simple* form as this introduces new variables which might have large values. We do not indulge further leaving spectra problems for future research.

Bibliography

- [Adleman and Manders, 1976] Adleman, L. and Manders, K. (1976). Diophantine complexity. In *17th Annual Symposium on Foundations of Computer Science (sfcs 1976)*, pages 81–88.
- [Andréka et al., 1998] Andréka, H., Németi, I., and van Benthem, J. (1998). Modal Languages and Bounded Fragments of Predicate Logic. *Journal of Philosophical Logic*, 27(3):217–274.
- [Barwise, 1967] Barwise, K. J. (1967). *Infinitary Logic and Admissible Sets*. Ph.d. thesis, Stanford University.
- [Bednarczyk, 2021] Bednarczyk, B. (2021). Exploiting forwardness: Satisfiability and query-entailment in forward guarded fragment. In Faber, W., Friedrich, G., Gebser, M., and Morak, M., editors, *Logics in Artificial Intelligence - 17th European Conference, JELIA 2021, Virtual Event, May 17-20, 2021, Proceedings*, volume 12678 of *Lecture Notes in Computer Science*, pages 179–193. Springer.
- [Bednarczyk et al., 2017] Bednarczyk, B., Charatonik, W., and Kieroński, E. (2017). Extending two-variable logic on trees. In Goranko, V. and Dam, M., editors, *26th EACSL Annual Conference on Computer Science Logic, CSL 2017, August 20-24, 2017, Stockholm, Sweden*, volume 82 of *LIPICs*, pages 11:1–11:20. Schloss Dagstuhl - Leibniz-Zentrum für Informatik.
- [Bednarczyk and Jaakkola, 2022] Bednarczyk, B. and Jaakkola, R. (2022). Towards a Model Theory of Ordered Logics: Expressivity and Interpolation. In Szeider, S., Ganian, R., and Silva, A., editors, *47th International Symposium on Mathematical Foundations of Computer Science (MFCS 2022)*, volume 241 of *Leibniz International Proceedings in Informatics (LIPICs)*, pages 15:1–15:14, Dagstuhl, Germany. Schloss Dagstuhl – Leibniz-Zentrum für Informatik.

- [Bednarczyk et al., 2023] Bednarczyk, B., Kojelis, D., and Pratt-Hartmann, I. (2023). On the Limits of Decision: the Adjacent Fragment of First-Order Logic. In Etessami, K., Feige, U., and Puppis, G., editors, *50th International Colloquium on Automata, Languages, and Programming (ICALP 2023)*, volume 261 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 111:1–111:21, Dagstuhl, Germany. Schloss Dagstuhl – Leibniz-Zentrum für Informatik.
- [Bednarczyk et al., 2024] Bednarczyk, B., Kojelis, D., and Pratt-Hartmann, I. (2024). The adjacent fragment and quine’s limits of decision.
- [Benedikt et al., 2024] Benedikt, M., Kostylev, E., and Tan, T. (2024). Two Variable Logic with Ultimately Periodic Counting. *SIAM Journal on Computing*, 53(4):884–968.
- [Benedikt et al., 2020] Benedikt, M., Kostylev, E. V., and Tan, T. (2020). Two Variable Logic with Ultimately Periodic Counting. In Czumaj, A., Dawar, A., and Merelli, E., editors, *47th International Colloquium on Automata, Languages, and Programming (ICALP 2020)*, volume 168 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 112:1–112:16, Dagstuhl, Germany. Schloss Dagstuhl – Leibniz-Zentrum für Informatik.
- [Berger, 1966] Berger, R. (1966). *The Undecidability of the Domino Problem*. Memoirs of the American Mathematical Society.
- [Chandra et al., 1981] Chandra, A. K., Kozen, D. C., and Stockmeyer, L. J. (1981). Alternation. *Journal of the ACM*, 28(1):114–133.
- [Charatonik et al., 2014] Charatonik, W., Kieroński, E., and Mazowiecki, F. (2014). Decidability of weak logics with deterministic transitive closure. In Henzinger, T. A. and Miller, D., editors, *Joint Meeting of the Twenty-Third EACSL Annual Conference on Computer Science Logic (CSL) and the Twenty-Ninth Annual ACM/IEEE Symposium on Logic in Computer Science (LICS), CSL-LICS ’14, Vienna, Austria, July 14 - 18, 2014*, pages 29:1–29:10. ACM.
- [Church, 1936] Church, A. (1936). A note on the entscheidungsproblem. *Journal of Symbolic Logic*, 1(1):40–41.

- [Durand et al., 2012] Durand, A., Jones, N. D., Makowsky, J. A., and More, M. (2012). Fifty years of the spectrum problem: Survey and new results. *The Bulletin of Symbolic Logic*, 18(4):505–553.
- [Fürer, 1983] Fürer, M. (1983). The computational complexity of the unconstrained limited domino problem (with implications for logical decision problems). *Logic and Machines: Decision Problems and Complexity*.
- [Gödel, 1933] Gödel, K. (1933). Zum Entscheidungsproblem des logischen Funktionenkalküls. 40(1):433–443.
- [Goldfarb, 1984] Goldfarb, W. D. (1984). The unsolvability of the gödel class with identity. *The Journal of Symbolic Logic*, 49(4):1237–1252.
- [Grädel, 1999] Grädel, E. (1999). On the Restraining Power of Guards. *The Journal of Symbolic Logic*, 64(4):1719–1742.
- [Grädel et al., 1997a] Grädel, E., Kolaitis, P. G., and Vardi, M. Y. (1997a). On the decision problem for two-variable first-order logic. *The Bulletin of Symbolic Logic*, 3(1):53–69.
- [Grädel et al., 1997b] Grädel, E., Otto, M., and Rosen, E. (1997b). Two-variable logic with counting is decidable. In *Proceedings of Twelfth Annual IEEE Symposium on Logic in Computer Science*, pages 306–317.
- [Grädel et al., 1999] Grädel, E., Otto, M., and Rosen, E. (1999). Undecidability results on two-variable logics. *Archive for Mathematical Logic*, 38(4):313–354.
- [Grädel, 1999] Grädel, E. (1999). On the restraining power of guards. *The Journal of Symbolic Logic*, 64(4):1719–1742.
- [Harel, 1984] Harel, D. (1984). A simple highly undecidable domino problem. In *Proceedings of the conference On Logic and Computation*.
- [Herzig, 1990] Herzig, A. (1990). A new decidable fragment of first order logic. In *Abstracts of the 3rd Logical Biennial Summer School and Conference in honour of S. C. Kleene*, Varna, Bulgaria.
- [Hilbert and Ackermann, 1928] Hilbert, D. and Ackermann, W. (1928). *Grundzüge der theoretischen Logik*. Springer, Berlin.

- [Hustadt et al., 2004] Hustadt, U., Schmidt, R. A., and Georgieva, L. (2004). A survey of decidable first-order fragments and description logics. *Journal of Relational Methods in Computer Science*, 1(3):251–276.
- [Jaakkola, 2021] Jaakkola, R. (2021). Ordered fragments of first-order logic. In Bonchi, F. and Puglisi, S. J., editors, *46th International Symposium on Mathematical Foundations of Computer Science, MFCS 2021, August 23-27, 2021, Tallinn, Estonia*, volume 202 of *LIPICs*, pages 62:1–62:14. Schloss Dagstuhl - Leibniz-Zentrum für Informatik.
- [Jeřábek, 2016] Jeřábek, E. (2016). Division by zero. *Archive for Mathematical Logic*, 55(7):997–1013.
- [Jones and Selman, 1974] Jones, N. D. and Selman, A. L. (1974). Turing machines and the spectra of first-order formulas. *The Journal of Symbolic Logic*, 39(1):139–150.
- [Keisler, 1971] Keisler, H. J. (1971). *Model theory for infinitary logic : logic with countable conjunctions and finite quantifiers*. Studies in logic and the foundations of mathematics ; v. 62. North-Holland Pub. Co., Amsterdam.
- [Keisler and Knight, 2004] Keisler, H. J. and Knight, J. F. (2004). Barwise: Infinitary logic and admissible sets. *The Bulletin of Symbolic Logic*, 10(1):4–36.
- [Kieronski, 2005] Kieronski, E. (2005). Results on the Guarded Fragment with Equivalence or Transitive Relations. In Ong, C. L., editor, *Computer Science Logic, 19th International Workshop, CSL 2005, 14th Annual Conference of the EACSL, Oxford, UK, August 22-25, 2005, Proceedings*, volume 3634 of *Lecture Notes in Computer Science*, pages 309–324. Springer.
- [Kieroński and Otto, 2012] Kieroński, E. and Otto, M. (2012). Small substructures and decidability issues for first-order logic with two variables. *Journal of Symbolic Logic*, 77(3):729–765.
- [Kieroński and Tendera, 2009] Kieroński, E. and Tendera, L. (2009). On finite satisfiability of two-variable first-order logic with equivalence relations. In *Proceedings of the 24th Annual IEEE Symposium on Logic in Computer Science, LICS 2009, 11-14 August 2009, Los Angeles, CA, USA*, pages 123–132. IEEE Computer Society.

- [Kojelis, 2025] Kojelis, D. (2025). On Homogenous Models of Fluted Languages. In *Under submission to CSL'25: Proceedings of the 33rd EACSL Annual Conference on Computer Science Logic*.
- [Mortimer, 1975] Mortimer, M. (1975). On languages with two variables. *Mathematical Logic Quarterly*, 21(1):135–140.
- [Pacholski et al., 2000] Pacholski, L., Szwoast, W., and Tendera, L. (2000). Complexity results for first-order two-variable logic with counting. *SIAM J. Comput.*, 29(4):1083–1117.
- [Papadimitriou, 1981] Papadimitriou, C. H. (1981). On the complexity of integer programming. *J. ACM*, 28(4):765–768.
- [Pratt-Hartmann, 2005] Pratt-Hartmann, I. (2005). Complexity of the two-variable fragment with counting quantifiers. *Journal of Logic, Language and Information*, 14(3):369–395.
- [Pratt-Hartmann, 2010] Pratt-Hartmann, I. (2010). The two-variable fragment with counting revisited. In Dawar, A. and de Queiroz, R. J. G. B., editors, *Logic, Language, Information and Computation, 17th International Workshop, WoLLIC 2010, Brasilia, Brazil, July 6-9, 2010. Proceedings*, volume 6188 of *Lecture Notes in Computer Science*, pages 42–54. Springer.
- [Pratt-Hartmann, 2021] Pratt-Hartmann, I. (2021). Fluted logic with counting. In Bansal, N., Merelli, E., and Worrell, J., editors, *48th International Colloquium on Automata, Languages, and Programming, ICALP 2021, July 12-16, 2021, Glasgow, Scotland (Virtual Conference)*, volume 198 of *LIPICs*, pages 141:1–141:17. Schloss Dagstuhl - Leibniz-Zentrum für Informatik.
- [Pratt-Hartmann, 2023] Pratt-Hartmann, I. (2023). *Fragments of First-Order Logic*. Oxford University Press.
- [Pratt-Hartmann, 2024] Pratt-Hartmann, I. (2024). Walking on Words. In Inenaga, S. and Puglisi, S. J., editors, *35th Annual Symposium on Combinatorial Pattern Matching, CPM 2024, June 25-27, 2024, Fukuoka, Japan*, volume 296 of *LIPICs*, pages 25:1–25:17. Schloss Dagstuhl - Leibniz-Zentrum für Informatik.

- [Pratt-Hartmann et al., 2016] Pratt-Hartmann, I., Szwast, W., and Tendera, L. (2016). Quine’s Fluted Fragment is Non-Elementary. In Talbot, J.-M. and Regnier, L., editors, *25th EACSL Annual Conference on Computer Science Logic (CSL 2016)*, volume 62 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 39:1–39:21, Dagstuhl, Germany. Schloss Dagstuhl – Leibniz-Zentrum für Informatik.
- [Pratt-Hartmann et al., 2019] Pratt-Hartmann, I., Szwast, W., and Tendera, L. (2019). The fluted fragment revisited. *Journal of Symbolic Logic*, 84(3):1020–1048.
- [Pratt-Hartmann and Tendera, 2019] Pratt-Hartmann, I. and Tendera, L. (2019). The Fluted Fragment with Transitivity. In Rossmanith, P., Heggernes, P., and Katoen, J.-P., editors, *44th International Symposium on Mathematical Foundations of Computer Science (MFCS 2019)*, volume 138 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 18:1–18:15, Dagstuhl, Germany. Schloss Dagstuhl – Leibniz-Zentrum für Informatik.
- [Pratt-Hartmann and Tendera, 2022] Pratt-Hartmann, I. and Tendera, L. (2022). The fluted fragment with transitive relations. *Annals of Pure and Applied Logic*, 173(1):103042.
- [Pratt-Hartmann and Tendera, 2023] Pratt-Hartmann, I. and Tendera, L. (2023). Adding Transitivity and Counting to the Fluted Fragment. In Klin, B. and Pimentel, E., editors, *31st EACSL Annual Conference on Computer Science Logic (CSL 2023)*, volume 252 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 32:1–32:22, Dagstuhl, Germany. Schloss Dagstuhl – Leibniz-Zentrum für Informatik.
- [Purdy, 1996] Purdy, W. C. (1996). Fluted formulas and the limits of decidability. *The Journal of Symbolic Logic*, 61(2):608–620.
- [Purdy, 2002] Purdy, W. C. (2002). Complexity and Nicety of Fluted Logic. *Stud Logica*.
- [Quine, 1969] Quine, W. V. O. (1969). On the Limits of Decision. In *Proceedings of the 14th International Congress of Philosophy*, volume III, pages 57–62. University of Vienna.

- [Quine, 1976] Quine, W. V. O. (1976). Algebraic logic and predicate functors. In *The Ways of Paradox*, pages 283–307. Harvard University Press, Cambridge, MA, revised and enlarged edition.
- [Schwentick and Zeume, 2012] Schwentick, T. and Zeume, T. (2012). Two-Variable Logic with Two Order Relations. *Logical Methods in Computer Science*, Volume 8, Issue 1.
- [Scott, 1962] Scott, D. (1962). A decision method for validity of sentences in two variables. *Journal of Symbolic Logic*, 27:477.
- [Turing, 1936] Turing, A. (1936). On computable numbers, with an application to the entscheidungsproblem. *Proceedings of the London Mathematical Society*, 42(1):230–265.