# Logics with Counting
# ESSLLI 2009

Ian Pratt-Hartmann
School of Computer Science
University of Manchester
Manchester M13 9PL, UK

# Contents

# Chapter 1

# Introduction

It is well-known that first-order logic is able to express facts about how many objects have a certain property. For example, the sentences

> No professor supervises more than 3 graduate students
> Every graduate student is supervised by at most 1 professor

may be formalized as

$$\neg\exists x(\mathrm{prof}(x)\wedge \exists y_1\ldots y_4\Big(\bigwedge_{1\leq i\leq 4}(\mathrm{grad}(y_i)\wedge \mathrm{sup}(x,y_i))\wedge \bigwedge_{1\leq i<j\leq 4} y_i \not\approx y_j\Big) \tag{1.1}$$

$$\forall x(\mathrm{grad}(x)\rightarrow \forall y_1 y_2\Big(\bigwedge_{1\leq i\leq 2}\big(\mathrm{prof}(y_i)\wedge \mathrm{sup}(y_i,x)\big)\rightarrow y_1\approx y_2\Big), \tag{1.2}$$

respectively. A more succinct (and readable) formalization is possible if, in addition to the familiar quantifiers $\forall$ and $\exists$, we employ so-called *counting quantifiers*, having the forms $\exists_{\leq C}$, $\exists_{\geq C}$ or $\exists_{=C}$, where $C$ is a numerical subscript. We read $\exists_{\leq C} v\phi$ as "there exist at most $C$ $v$ such that $\phi$", $\exists_{\geq C} v\phi$ as "there exist at least $C$ $v$ such that $\phi$" and $\exists_{=C} v\phi$ as "there exist exactly $C$ $v$ such that $\phi$". With this apparatus, the above sentences may be formalized as

$$\neg\exists x(\mathrm{prof}(x)\wedge \exists_{\geq 4} y(\mathrm{grad}(y)\wedge \mathrm{sup}(x,y))) \tag{1.3}$$

$$\forall x(\mathrm{grad}(x)\rightarrow \exists_{\leq 1} y(\mathrm{prof}(y)\wedge \mathrm{sup}(y,x))). \tag{1.4}$$

The formation rules governing formulas with counting quantifiers are completely analogous to those for $\forall$ and $\exists$.

Formally, the semantics of counting quantifiers may be given in terms of satisfaction in a structure relative to a variable assignment, in the standard way. Specifically, for any structure $\mathfrak{A}$ with domain $A$, and any variable assignment

function $s$, we define

$$\mathfrak{A} \models_s \exists_{\leq C} v\phi \text{ iff } |\{s' \mid s' =_v s \text{ and } \mathfrak{A} \models_{s'} \phi\}| \leq C$$
$$\mathfrak{A} \models_s \exists_{\geq C} v\phi \text{ iff } |\{s' \mid s' =_v s \text{ and } \mathfrak{A} \models_{s'} \phi\}| \geq C$$
$$\mathfrak{A} \models_s \exists_{=C} v\phi \text{ iff } |\{s' \mid s' =_v s \text{ and } \mathfrak{A} \models_{s'} \phi\}| = C,$$

where, as usual, $s' =_v s$ indicates that the functions $s$ and $s'$ are identical except possibly in the value they assign to $v$. It is easy to verify that, under these semantics, $\exists x\phi$ is logically equivalent to $\exists_{\geq 1} x\phi$, $\forall x\phi$ is logically equivalent to $\exists_{\leq 0} x\neg\phi$, $\exists_{=C} x\phi$ is logically equivalent to $\exists_{\leq C} x\phi \wedge \exists_{\geq C} x\phi$, $\exists_{\leq C} x\phi$ is logically equivalent to $\neg\exists_{\geq C+1} x\phi$, and $\exists_{\geq 0} x\phi$ is logically true. Thus, we could if we wished eliminate all the other quantifiers in favour of $\exists_{\geq C}$ ($C > 0$). However, we continue to employ the redundant quantifiers for ease of exposition.

At first blush, it might seem strange to consider counting quantifiers at all, since they do not extend the expressive power of first-order logic: any formula involving counting quantifiers can be translated into a logically equivalent formula without counting quantifiers. Thus, for example, for any $C > 0$, we have the logical equivalence

$$\exists_{\geq C} x\phi(x) \equiv \exists x_1, \ldots, x_C \big( \bigwedge_{1 \leq i \leq C} \phi(x_i) \wedge \bigwedge_{1 \leq i < j \leq C} x_i \not\approx x_j \big),$$

where $x_1, \ldots x_C$ do not occur free in $\phi$; and similarly, *mutatis mutandis*, for the quantifiers $\exists_{\leq C}$ and $\exists_{=C}$. However, matters are different when we consider certain fragments of first-order logic, for which the addition of counting quantifiers yields proper super-fragments with interesting computational properties. These properties form the subject of these notes.

In general, by a *logic*, $\mathcal{L}$, we understand any set of expressions, known as *formulas* of $\mathcal{L}$ (or *$\mathcal{L}$-formulas*), equipped with a truth-conditional semantics. Specifically, we assume that $\mathcal{L}$ associates with any $\mathcal{L}$-formula $\phi$ a signature of non-logical primitives (constants, function-symbols and predicates of various types) occurring in $\phi$, and determines, for any structure $\mathfrak{A}$ interpreting these primitives, whether $\phi$ is satisfied in $\mathfrak{A}$. Under these assumptions, we say that $\phi$ is *satisfiable* if it is satisfied in some structure, and *finitely satisfiable* if it is satisfied in some finite structure. A set $\Phi$ of formulas is satisfiable if all its members are simultaneously satisfied in some structure, and similarly for finite satisfiability. Trivially, finite satisfiability implies satisfiability. If, for finite sets of $\mathcal{L}$-formulas, the converse implication holds, then $\mathcal{L}$ is said to have the *finite model property*.

**Definition 1.** Let $\mathcal{L}$ be a logic. The *satisfiability problem for* $\mathcal{L}$, Sat-$\mathcal{L}$, is the problem of determining whether a given finite set of $\mathcal{L}$-formulas is satisfiable; the *finite satisfiability problem for* $\mathcal{L}$, Fin-Sat-$\mathcal{L}$, is the problem of determining whether a given finite set of $\mathcal{L}$-formulas is finitely satisfiable.

Notice that Definition 1 does not assume that $\mathcal{L}$ is closed under Boolean conjunction. We assume that $\mathcal{L}$-formulas (and sets of such) may be coded, in some

standard way, as strings over a finite alphabet. The size of a formula $\phi$, denoted $\|\phi\|$, is simply the length of its encoding string. Thus, (finite sets of) formulas may be regarded as inputs to Turing machines over the relevant alphabet, so that the notion of decidability and the apparatus of computational complexity theory can be applied to Sat-$\mathcal{L}$ and Fin-Sat-$\mathcal{L}$.

A logic $\mathcal{L}$ has the finite model property if and only if the problems Sat-$\mathcal{L}$ and Fin-Sat-$\mathcal{L}$ coincide. Moreover, if $\mathcal{L}$ is a subset of first-order logic having the finite model property, then Sat-$\mathcal{L}$ (=Fin-Sat-$\mathcal{L}$) is decidable. For, given any finite set $\Phi$ of $\mathcal{L}$-formulas, we may enumerate all finite structures (over the signature of $\Phi$) and all theorems of first-order logic, in parallel, stopping when either a satisfying structure or the theorem $\neg \bigwedge \Phi$ is found. It is well-known that first-order logic lacks the finite model property, and that its satisfiability problem and finite satisfiability problem are both undecidable. However, many decidable fragments are known, and we briefly consider some well-known cases here.

The *two-variable fragment* (of first-order logic) *with equality*, denoted $\mathcal{L}^2_{\approx}$, is the set of first-order formulas involving only the two variables $x$ and $y$. For example, the sentences

Every graduate student is supervised by a professor who teaches some course

may be formalized as

$$\forall x(\mathrm{grad}(x) \rightarrow \exists y(\mathrm{sup}(y,x) \wedge \mathrm{prof}(y) \wedge \exists x(\mathrm{teach}(y,x) \wedge \mathrm{course}(x)))). \quad (1.5)$$

Notice, incidentally, that this formula has triply-nested quantifiers; nevertheless, it is in the two-variable fragment, because the variable $x$ bound by the outermost quantifier is re-used by the innermost quantifier. Despite its non-trivial expressive power, the fragment $\mathcal{L}^2_{\approx}$ can be shown to have the finite model property. Thus Sat-$\mathcal{L}^2_{\approx}$ (= Fin-Sat-$\mathcal{L}^2_{\approx}$) is a decidable problem. In fact, this problem is NEXPTIME-complete. By contrast, the three-variable fragment, $\mathcal{L}^3_{\approx}$, defined analogously, lacks the finite model property, and both Sat-$\mathcal{L}^3_{\approx}$ and Fin-Sat-$\mathcal{L}^3_{\approx}$ are undecidable.

But what happens if we add counting quantifiers to $\mathcal{L}^2_{\approx}$? Define the *two-variable fragment with counting*, denoted $\mathcal{C}^2$, to be the set of first-order formulas involving only the two variables $x$ and $y$, but with the counting quantifiers $\exists_{\leq C}$, $\exists_{\geq C}$ and $\exists_{=C}$ (for every $C \geq 0$) allowed. Thus, for example, the formulas (1.3) and (1.4) are in $\mathcal{C}^2$. It is not difficult to see that $\mathcal{C}^2$ is strictly more expressive than $\mathcal{L}^2_{\approx}$. In fact, $\mathcal{C}^2$ lacks the finite model property, so that the problems Sat-$\mathcal{C}^2$ and Fin-Sat-$\mathcal{C}^2$ do not coincide. Nevertheless, it can be shown that these problems are still both decidable. Indeed, as we shall see below, they both have the same computational complexity as Sat-$\mathcal{L}^2_{\approx}$.

A second well-known decidable fragment of first-order logic is the so-called *guarded fragment*, denoted $\mathcal{G}$. A formal treatment will be given below, but, roughly, a guarded formula is one in which all quantification is restricted to the patterns

$$\forall \bar{v}(\gamma \rightarrow \psi) \qquad\qquad \exists \bar{v}(\gamma \wedge \psi),$$

where $\gamma$ is an atomic formula featuring all the free variables in $\psi$ together with $\bar{v}$. The formula (1.5) above is guarded. As an example of a non-guarded formula, we need look no further than

Every graduate student admires every professor

$$\forall x(\mathrm{grad}(x) \to \forall y(\mathrm{prof}(y) \to \mathrm{adm}(x, y)));$$

for the sub-formula $\mathrm{adm}(x, y)$ has a free variable, namely, $x$, which is not an argument of the atom $\mathrm{prof}(y)$. The guarded fragment was originally introduced as a generalization of modal logic, and exhibits much of the latter's well-behavedness. In particular, it has the finite model property, so that $\mathcal{G}$-Sat (=$\mathcal{G}$-Fin-Sat) is decidable. Just as with first-order logic, so too with the guarded fragment, it makes sense to consider its finite-variable sub-fragments. We denote by $\mathcal{G}^k$ the fragment of $\mathcal{G}$ involving at most $k$ variables. It transpires that bounding the number of variables reduces the computational complexity of the satisfiability problem.

What happens if we add counting quantifiers to $\mathcal{G}$? Define the *$k$-variable guarded fragment with counting*, denoted $\mathcal{GC}^k$, to be the result of extending $\mathcal{G}^k$ by allowing the counting quantifiers $\exists_{\leq C}$, $\exists_{\geq C}$ and $\exists_{=C}$, as long as they occur in the pattern $Qv(\gamma \wedge \psi)$, where $\gamma$ is an atomic formula featuring all the free variables of $\psi$ together with $v$. Here, the results are mixed. Even for $k \geq 2$, $\mathcal{GC}^k$ lacks the finite model property. Nevertheless, Sat-$\mathcal{GC}^2$ and Fin-Sat-$\mathcal{GC}^2$ are both decidable, again with the same computational complexity as Sat-$\mathcal{G}^2$. However, for $k \geq 3$, Sat-$\mathcal{GC}^k$ and Fin-Sat-$\mathcal{GC}^k$ are undecidable.

Since we have mentioned results on computational complexity, this is perhaps an appropriate point at which to clarify two details regarding the sizes of formulas: one unimportant, the other crucial. The first concerns the encoding of non-logical primitives. For most interesting fragments, signature elements cannot, strictly speaking, be coded as a single symbol, but rather as a string of symbols whose length is logarithmic in the total number of symbols appearing in the formula. However, doing so has no effect on the computational complexity of the fragments considered here, and the issue may be safely disregarded. The second detail concerns the encoding of numerical subscripts in counting quantifiers. Such subscripts are, strictly speaking, *numerals*, rather than numbers: strings of binary digits encoding numbers in the standard way. (It makes no essential difference whether we use base 2 or some other base.) Thus, in treating, as we shall, quantifier subscripts as if they were *numbers*, rather than numerals, we are employing a notational and textual shortcut. In this connection, however, it is important to understand that *size* of a positive numerical subscript $C$ is not $C$, but rather $\lfloor \log_2 C \rfloor + 1$, where $\lfloor x \rfloor$ denotes the largest integer less than or equal to $x$. Put another way, the numbers featuring in numerical subscripts are in general exponentially large compared with the formulas containing them.

# Chapter 2

# Finite-variable fragments of first-order logic

We explained in Chapter 1 that the counting quantifiers are of interest primarily in terms of their computational effect on certain decidable fragments of first-order logic. The purpose of this chapter is to lay the groundwork for our investigation of counting quantifiers by introducing some of these fragments.

For $k > 1$, the *$k$-variable fragment of first-order logic with equality*, denoted $\mathcal{L}^k_{\approx}$, is the set of first-order formulas over a purely relational signature involving no variables other than $x_1, \ldots, x_k$. The *$k$-variable fragment of first-order logic* (*without equality*), denoted $\mathcal{L}^k$, is the set $\mathcal{L}^k_{\approx}$-formulas not involving the equality predicate $\approx$. For ease of reading, we write $x$, $y$, $z$ instead of $x_1$, $x_2$, $x_3$.

## 2.1 Fragments with one variable

We begin with the 1-variable fragments $\mathcal{L}^1$ and $\mathcal{L}^1_{\approx}$, whose treatment is nearly trivial. Since we are working with a purely relational signature (no individual constants of function-symbols), the only equality atom in $\mathcal{L}^1_{\approx}$ is $x \approx x$, which may be equivalently replaced by $\top$. Therefore, we consider only $\mathcal{L}^1$. In fact, we may as well assume that all predicates have arity 0 or 1, since predicates of higher arity evidently do not essentially increase expressive power in $\mathcal{L}^1$. For the remainder this section, then, we assume a signature of nullary and unary predicates only.

**Lemma 1.** *Let $\phi$ be an $\mathcal{L}^1$-formula. We can construct, in time bounded by a polynomial function of $\|\phi\|$, an $\mathcal{L}^1$-formula*

$$\psi := \forall x \alpha \wedge \bigwedge_{1 \leq h \leq m} \exists x \beta_h, \tag{2.1}$$

*where $1 \leq m \leq \|\phi\|$, and $\alpha, \beta_1, \ldots, \beta_m$ are quantifier-free $\mathcal{L}^1$-formulas, such that $\phi$ and $\psi$ are satisfiable over the same domains.*

*Proof.* By prefixing an existential quantifier if necessary, we may assume that $\phi$ is a sentence. Set $\phi_0 := \forall x\phi$.

If $\phi_0$ has a subformula $\theta = \exists x\chi$, with $\chi$ quantifier-free, let $p$ be a new nullary predicate, let $\phi_1 = \phi[p/\theta]$, and let

$$\psi_1 = \exists x(p \to \chi) \wedge \forall x(\chi \to p).$$

It is easy to see that $\phi_0$ and $\phi_1 \wedge \psi_1$ are satisfiable over the same domains. For, on the one hand, $\phi_1 \wedge \psi_1$ entails $\phi_0$, and on the other, any model $\mathfrak{A}$ of $\phi_0$ may be expanded to a model of $\phi_1 \wedge \psi_1$ by interpreting $p$ to be true if and only if $\mathfrak{A} \models \theta$. Similarly, if $\phi_0$ has a proper subformula $\theta = \forall x\chi$, with $\chi$ quantifier-free, define $\phi_1$ and $\psi_1$ analogously, subject to the obvious adjustments. Now process $\phi_1$ in the same way, and continue until some formula $\phi_m$ is reached having the form $\forall x\chi$, with $\chi$ quantifier-free. Thus, $\phi_0$ and

$$\phi_m \wedge \psi_m \wedge \psi_{m-1} \wedge \cdots \wedge \psi_1$$

are satisfied over the same domains. Re-arrangement of conjuncts yields the desired formula $\psi$. □

**Lemma 2.** *Let $\phi$ be a formula in $\mathcal{L}^1$. If $\phi$ is satisfiable, then it is satisfiable over a domain of size at most $\|\phi\|$.*

*Proof.* By Lemma 1, we may assume $\phi$ to be of the form (2.1). If $\mathfrak{A} \models \phi$, let $a_1, \ldots a_m \in A$ be witnesses for the respective conjuncts $\exists x\beta_h$. Let $\mathfrak{B}$ be the restriction of $\mathfrak{A}$ to $\{a_1, \ldots, a_m\}$. It is then obvious that $\mathfrak{B} \models \phi$. □

**Theorem 1.** *The problem Sat-$\mathcal{L}^1$ is* NPTime*-complete.*

*Proof.* Membership in NP follows from Lemma 2. NP-hardness is immediate, since $\mathcal{L}^1$ includes propositional logic. □

We remark that allowing individual constants to appear in the fragment has no effect on the above results (though it clutters the exposition); the details are routine and left to the reader.

Consider the sub-fragment $\mathcal{S}$ of $\mathcal{L}^1$ consisting of all formulas having the following forms

$$\begin{array}{ll} \exists x(p(x) \wedge q(x)) & \exists x(p(x) \wedge \neg q(x)) \\ \forall x(p(x) \to q(x)) & \forall x(p(x) \to \neg q(x)), \end{array} \tag{2.2}$$

where $p$ and $q$ are unary predicates. This fragment is of some historical and linguistic interest, because we can think of unary predicates as corresponding to common nouns, and the formulas (2.2) to English sentences of the forms

$$\begin{array}{ll} \mathsf{Some}\ p\ \mathsf{is\ a}\ q & \mathsf{Some}\ p\ \mathsf{is\ not\ a}\ q \\ \mathsf{Every}\ p\ \mathsf{is\ a}\ q & \mathsf{No}\ p\ \mathsf{is\ a}\ q, \end{array} \tag{2.3}$$

respectively. This is, in essence, the language of the syllogistic set out in Aristotle's *Prior analytics*. Inspection of the above forms shows that, if $\phi$ is an

$\mathcal{S}$-formula, then there is an $\mathcal{S}$-formula $\bar{\phi}$ logically equivalent to its negation. Thus, validity and satisfiability are dual in $\mathcal{S}$ in the usual way: an argument with premises $\Phi$ and conclusion $\phi$ is valid if and only if the set $\phi \cup \{\bar{\phi}\}$ is not satisfiable. Note, however that $\mathcal{S}$ is not closed under conjunction. (Thus, it is generally only interesting to consider the satisfiability of *sets* of $\mathcal{S}$-formulas.)

**Theorem 2.** *The satisfiability problem for $\mathcal{S}$ is* NLogSpace*-complete.*

*Proof.* The problem 2-SAT is the defined as follows. Given a set $\Gamma$ of propositional clauses, none of which contains more than two literals, determine whether $\Gamma$ is satisfiable. It is well-known that 2-SAT is NLogSpace-complete (see, e.g. Papadimitriou [27], pp. 398.) It is straightforward to show that the satisfiability problem for $\mathcal{S}$ and the problem 2-SAT can be reduced to each other. The theorem then follows from the fact that NLogSpace is closed under reductions. □

Adding individual constants to $\mathcal{S}$ makes no difference to Theorem 2, a detail we leave to the reader to verify. Adding counting quantifiers, by contrast, does make a difference, and we consider this matter in Chapter 2.

## 2.2 Fragments with two variables

In this section, we show that the fragments $\mathcal{L}^2$ and $\mathcal{L}^2_{\approx}$ have the finite model property, and that their satisfiability (= finite satisfiability) problems are NExpTime-complete. In dealing with these fragments, we may as well assume that all predicates have arity at most 2, since predicates of higher arity evidently do not essentially increase expressive power in $\mathcal{L}^2$. For the remainder of this section, then, we assume a signature of nullary, unary and binary predicates only.

**Lemma 3** (Scott normal form)**.** *Let $\phi$ be a formula in $\mathcal{L}^2_{\approx}$. We can construct, in time bounded by a polynomial function of $\|\phi\|$, an $\mathcal{L}^2_{\approx}$-formula*

$$\psi := \forall x \forall y(\alpha \vee x \approx y) \wedge \bigwedge_{1 \leq h \leq m} \forall x \exists y(\beta_h(x, y) \wedge x \not\approx y), \tag{2.4}$$

*where $1 \leq m \leq \|\phi\|$, and $\alpha, \beta_1, \ldots, \beta_m$ are quantifier-free $\mathcal{L}^2$-formulas, such that $\phi$ and $\psi$ are satisfiable over the same domains containing at least 2 elements.*

*Proof.* We proceed as for Lemma 1. By prefixing existential quantifiers if necessary, we may assume that $\phi$ is a sentence. Set $\phi_0 := \forall x \phi$. In the sequel, let $u$, $v$ be the variables $x$, $y$, in either order.

Suppose $\phi_0$ has a subformula $\theta(v) = \exists u \chi$, with $\chi$ quantifier-free. Let $p$ be a new unary predicate, let $\phi_1$ be $\phi[p(v)/\theta(v)]$, and let

$$\psi_1 := \forall v \exists u(p(v) \rightarrow \chi) \wedge \forall v \forall u(\chi \rightarrow p(v)).$$

It is easy to see that $\phi_0$ and $\phi_1 \wedge \psi_1$ are satisfiable over the same domains. For, on the one hand, $\phi_1 \wedge \psi_1$ entails $\phi_0$, and on the other, any model $\mathfrak{A}$ of $\phi_0$

may be expanded to a model of $\phi_1 \wedge \psi_1$ by interpreting $p$ to be satisfied by an element $a \in A$ if and only if $\mathfrak{A} \models \theta[a]$. Similarly, if $\phi_0$ has a proper subformula $\theta(v) = \forall x\chi$, with $\chi$ quantifier-free, define $\phi_1$ and $\psi_1$ analogously, subject to the obvious adjustments. Now process $\phi_1$ in the same way, and continue until some formula $\phi_m$ is reached having the form $\forall x p(x)$—which we may re-write as $\forall x \forall y p(x)$. Thus, $\phi_0$ and

$$\phi_m \wedge \psi_m \wedge \psi_{m-1} \wedge \cdots \wedge \psi_1$$

are satisfied over the same domains. Re-arrangement of conjuncts yields a formula

$$\psi' := \forall x \forall y \alpha'(x,y) \wedge \bigwedge_{1 \leq h \leq m} \forall x \exists y (\beta'_h(x,y)),$$

where $\alpha'(x,y)$ and the $\beta'_h(x,y)$ are quantifier-free $\mathcal{L}^2_{\approx}$-formulas.

It remains only to reform the occurrences of $\approx$ in $\alpha'(x,y)$ and the $\beta'_h(x,y)$. Restricting attention to domains containing at least 2 elements, we have the following logical equivalences:

$$\forall x \forall y \alpha'(x,y) \equiv \forall x \forall y ((\alpha'(x,x) \wedge \alpha'(x,y)) \vee x \approx y)$$
$$\forall x \exists y \beta'_h(x,y) \equiv \forall x \exists y ((\beta'_h(x,x) \vee \beta'_h(x,y)) \wedge x \not\approx y).$$

Now, if $\theta$ is any $\mathcal{L}^2_{\approx}$-formula, denote by $\theta^*$ the result of replacing all atoms of the forms $x \approx x$ and $y \approx y$ in $\theta$ by $\top$, and all atoms of the forms $x \approx y$ and $y \approx x$ by $\bot$. Thus, we have the logical equivalences

$$\theta \vee x \approx y \equiv \theta^* \vee x \approx y \qquad \theta \wedge x \not\approx y \equiv \theta^* \wedge x \not\approx y.$$

Setting

$$\alpha := \big(\alpha'(x,y) \wedge \alpha'(x,x)\big)^*$$
$$\beta_h := \big(\beta'_h(x,y) \vee \beta'_h(x,x)\big)^*,$$

for all $h$ $(1 \leq h \leq m)$, and then

$$\psi := \forall x \forall y (\alpha(x,y) \vee x \approx y) \wedge \bigwedge_{1 \leq h \leq m} \forall x \exists y (\beta_h(x,y) \wedge x \not\approx y),$$

we see that $\psi$ and $\psi'$ are logically equivalent over domains containing at least 2 elements. Thus, $\psi$ has the properties required for the lemma. □

At this point, we introduce some familiar concepts which will feature prominently in the sequel. Fix some purely relational signature $\Sigma$. A *literal* (over $\Sigma$) is an atomic formula or the negation of an atomic formula. A *1-type* (over $\Sigma$) is a maximal consistent set of equality-free literals over $\Sigma$ involving only the variable $x$. A *2-type* (over $\Sigma$) is a maximal consistent set of equality-free literals over $\Sigma$ involving only the variables $x$ and $y$. Reference to $\Sigma$ is suppressed where clear from context. If $\mathfrak{A}$ is any structure interpreting $\Sigma$, and $a \in A$, then there

exists a unique 1-type $\pi(x)$ over $\Sigma$ such that $\mathfrak{A} \models \pi[a]$; we denote $\pi$ by $\mathrm{tp}^{\mathfrak{A}}[a]$. If, in addition, $b \in A$ is distinct from $a$, then there exists a unique 2-type $\tau(x, y)$ over $\Sigma$ such that $\mathfrak{A} \models \tau[a, b]$; we denote $\tau$ by $\mathrm{tp}^{\mathfrak{A}}[a, b]$. We do not define $\mathrm{tp}^{\mathfrak{A}}[a, b]$ if $a = b$. If $\pi$ is a 1-type, we say that $\pi$ is *realized* in $\mathfrak{A}$ if there exists $a \in A$ with $\mathrm{tp}^{\mathfrak{A}}[a] = \pi$. If $\tau$ is a 2-type, we say that $\tau$ is *realized* in $\mathfrak{A}$ if there exist distinct $a, b \in A$ with $\mathrm{tp}^{\mathfrak{A}}[a, b] = \tau$.

**Notation 1.** Let $\tau$ be a 2-type over a purely relational signature $\Sigma$. The result of transposing the variables $x$ and $y$ in $\tau$ is also a 2-type, denoted $\tau^{-1}$; and the set of literals in $\tau$ not featuring the variable $y$ is a 1-type, denoted $\mathrm{tp}_1(\tau)$. We write $\mathrm{tp}_2(\tau)$ for the 1-type $\mathrm{tp}_1(\tau^{-1})$.

Note that $\mathrm{tp}_2(\tau)$ is the result of taking the set of literals in $\tau$ not featuring the variable $x$, and then replacing $y$ throughout by $x$.

**Remark 1.** *If $\tau$ is any 2-type over a purely relational signature $\Sigma$, $\mathfrak{A}$ is a structure interpreting $\Sigma$, and $a$, $b$ are distinct elements of $A$ such that $\mathrm{tp}^{\mathfrak{A}}[a, b] = \tau$, then $\mathrm{tp}^{\mathfrak{A}}[b, a] = \tau^{-1}$, $\mathrm{tp}^{\mathfrak{A}}[a] = \mathrm{tp}_1(\tau)$ and $\mathrm{tp}^{\mathfrak{A}}[b] = \mathrm{tp}_2(\tau)$.*

A terminological note: in books on model theory, the word "type" is standardly used to refer to a maximal consistent set of *formulas* (over some signature) featuring a fixed collection of variables—including formulas involving quantifiers. What *we* are calling types here are known, in that nomenclature, as "rank-0 types". In the sequel, however, we only ever have occasion to refer to rank-0 types, and so we continue to use the more abbreviated terminology.

**Remark 2.** *If $\Sigma$ features only unary and binary predicates, and $|\Sigma| = s$, then there are exactly $2^s$ 1-types over $\Sigma$ and at most $2^{4s}$ 2-types.*

**Lemma 4.** *Let $\phi$ be a satisfiable $\mathcal{L}^2_{\approx}$-formula, and let $n = \|\phi\|$. Then $\phi$ has a model of size at most $3n.2^n$.*

*Proof.* If $\phi$ is satisfiable over a 1-element domain, there is nothing to prove. Moreover, we may assume without loss of generality that the signature of $\phi$ features only unary and binary predicates, since any nullary predicates can be replaced by $\top$ or $\bot$ according to their truth-value in some model. Now let $\psi$ be the formula (2.4) constructed in Lemma 3, and suppose $\mathfrak{A} \models \psi$. Let the signature of $\psi$ be $\Sigma^*$; thus $m \leq n$ and $|\Sigma^*| \leq n$. It suffices to construct a model $\mathfrak{B}$ of $\psi$ with $|B| \leq 3n.2^n$. In the rest of the proof, we employ the symbols $\alpha$, $m$ $\beta_1, \ldots, \beta_m$ as they appear in (2.4).

If $a \in A$, we say that $a$ is a *king* if $a$ is the only element $a' \in A$ such that $\mathrm{tp}^{\mathfrak{A}}[a'] = \mathrm{tp}^{\mathfrak{A}}[a]$. Let $K$ be the set of kings. Evidently, $|K| \leq 2^n$. We now define a set of elements $C \subseteq A$, called the *court*, as follows: every king is a member of court; and, for each king $a$, and every $h$ $(1 \leq h \leq m)$, we select one element $b \in A \setminus \{a\}$ such that $\mathfrak{A} \models \beta_h[a, b]$, and let $b$ be a member of court. (These elements need not be distinct, and may themselves be kings.) Evidently, $|C| \leq (m + 1)|K|$. Let $\mathfrak{C}$ be the structure on $C$ induced by $\mathfrak{A}$. Notice that, for any element $a \in A$, and any $h$ $(1 \leq h \leq m)$, there exists an element $b \in A \setminus \{a\}$

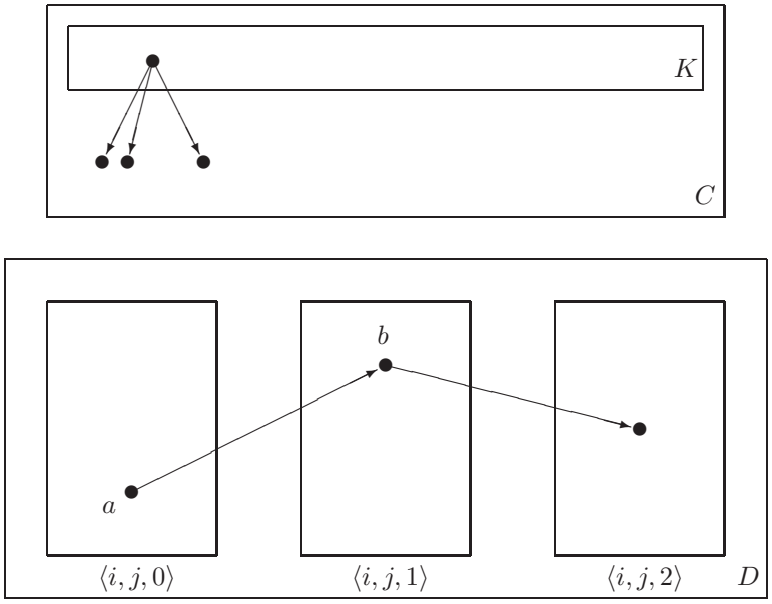


Figure 2.1: The small model property for $\mathcal{L}^2_{\approx}$, showing the absence of clashes in Stage 3 of the proof of Lemma 4

(depending on $a$ and $h$) such that $\mathfrak{A} \models \beta_h[a, b]$. As we might say, $b$ satisfies the 'existential requirement' $\exists y(\beta_h(x, y) \wedge x \not\approx y)$ of $a$.

Let $\pi_1, \ldots, \pi_L$ be the 1-types realized in $\mathfrak{A}$ by more than one element (i.e. realized in $\mathfrak{A}$, but not by kings). Evidently, $L \leq 2^n - |K|$. Let $D$ be the set of integer triples

$$\{\langle i, h, k\rangle | 1 \leq i \leq L,\ 1 \leq h \leq m \text{ and } 0 \leq k \leq 2\},$$

and assume without loss of generality that $C$ and $D$ are disjoint. Now let $B = C \cup D$. Hence $|B| \leq (m + 1)|K| + 3m(2^n - |K|) \leq 3n2^n$. We define a structure on $\mathfrak{B}$ satisfying $\psi$. In defining this structure (Fig. 2.1), we ensure that all the existential requirements of successive elements of $B$ are satisfied. We proceed in three stages.

**Stage 1:** Import the structure $\mathfrak{C}$ onto the elements $C \subseteq B$ (so that $\mathfrak{A}$ and $\mathfrak{B}$ agree on the court). Consider any king $a$; and let $1 \leq h \leq m$. By construction of $C$, there exists $b \in C \setminus \{a\}$ such that $\mathfrak{A} \models \beta_h[a, b]$, whence $\mathfrak{B} \models \beta_h[a, b]$. Thus, all kings have their existential requirements satisfied.

**Stage 2:** First, fix the 1-type of any element $\langle i, h, k\rangle \in D$ by setting

$$\mathrm{tp}^{\mathfrak{B}}[\langle i, h, k\rangle] = \pi_i.$$

Now consider any $a \in C$ which is not a king, and any $h$ such that $1 \leq h \leq m$. Choose some $b \in A \setminus \{a\}$ such that $\mathfrak{A} \models \beta_h[a, b]$. If $b$ is a king, then $\mathrm{tp}^{\mathfrak{B}}[a, b] = \mathrm{tp}^{\mathfrak{A}}[a, b]$ from Stage 1, so that we already have $\mathfrak{B} \models \beta_h[a, b]$. If $b$ is not a king, then its 1-type is $\pi_i$ for some $i$ $(1 \leq i \leq L)$, so let $b' = \langle i, h, 0\rangle \in D$, and set $\mathrm{tp}^{\mathfrak{B}}[a, b'] = \mathrm{tp}^{\mathfrak{A}}[a, b]$. Since $\mathrm{tp}^{\mathfrak{A}}[b] = \pi_i = \mathrm{tp}^{\mathfrak{B}}[b']$, there is no clash with any of the 1-types already assigned to elements of $\mathfrak{B}$; hence this assignment is legitimate. Proceeding in this way, we can satisfy the existential requirements of all the elements in $C$.

**Stage 3:** Consider any $a = \langle i, h', k\rangle \in D$, and any $h$ such that $1 \leq h \leq m$. Let $a' \in A$ be such that $\mathrm{tp}^{\mathfrak{A}}[a'] = \pi_i$, and pick $b' \in A \setminus \{a\}$ such that $\mathfrak{A} \models \beta_h[a', b']$. If $b'$ is a king, let $b = b'$; otherwise, let $i'$ be such that $\mathrm{tp}^{\mathfrak{A}}[b'] = \pi_{i'}$, and let $b = \langle i', h, k + 1\rangle$. Set $\mathrm{tp}^{\mathfrak{B}}[a, b] = \mathrm{tp}^{\mathfrak{A}}[a', b']$ (we show presently that we are free to do this). Proceeding in this way, we can satisfy the existential requirements for all $a \in D$.

It is immediate that the 2-type assignments made in this stage cannot clash with any previously assigned 1-types. However, we must show that the assignment to $\mathrm{tp}^{\mathfrak{B}}[a, b]$ described above made in this Stage cannot clash with any previously-made assignment to $\mathrm{tp}^{\mathfrak{B}}[b, a]$. To see that this is so, suppose first that $b'$ is a king. Then $b = b'$, and any previous assignment to $\mathrm{tp}^{\mathfrak{B}}[b, a]$ must have occurred in Stage 1. But that is impossible, because, in Stage 1 kings had their 2-types fixed only only with respect to other members of court. Suppose, on the other hand, that $b'$ is not a king, so that $b \in D$, and any previous assignment to $\mathrm{tp}^{\mathfrak{B}}[b, a]$ must have occurred in Stage 3. Writing $a = \langle i_1, h_1, k_1\rangle$ and $b = \langle i_2, h_2, k_2\rangle$, we then have $k_2 = k_1 + 1 \mod 3$ and $k_1 = k_2 + 1 \mod 3$, which is impossible (see Fig 2.1, where $k_1 = 0$). Hence, none of the 2-type assignments made in Stage 3 clashes with any other.

**Stage 4:** To complete the definition of $\mathfrak{B}$, suppose $\mathrm{tp}^{\mathfrak{B}}[a, b]$ was not assigned in any of the Stages 1–3. Certainly, then, $a$ and $b$ cannot both be kings; in particular, if $\mathrm{tp}^{\mathfrak{B}}[a] = \mathrm{tp}^{\mathfrak{B}}[b]$, it follows that this common 1-type is realized more than once in $\mathfrak{A}$. Therefore, we can always find distinct elements $a'$, $b'$ of $A$ such that $\mathrm{tp}^{\mathfrak{B}}[a] = \mathrm{tp}^{\mathfrak{A}}[a']$ and $\mathrm{tp}^{\mathfrak{B}}[b] = \mathrm{tp}^{\mathfrak{A}}[b']$. Now make the 2-type assignment $\mathrm{tp}^{\mathfrak{B}}[a, b] = \mathrm{tp}^{\mathfrak{A}}[a', b']$. This obviously does not clash with any previously-made 1-type assignments, and completes the definition of $\mathfrak{B}$.

Since all existential requirements are satisfied following Stages 1–3, we have

$$\mathfrak{B} \models \bigwedge_{1 \leq h \leq m} \forall x \exists y (\beta_h(x, y) \wedge x \not\approx y).$$

And since all 2-types realized in $\mathfrak{B}$ are also realized in $\mathfrak{A}$, we have

$$\mathfrak{B} \models \forall x \forall y (\alpha \vee x \approx y).$$

□

**Corollary 1.** *The problem Sat-$\mathcal{L}^2_{\approx}$ (=Fin-Sat-$\mathcal{L}^2_{\approx}$) is in NEXPTIME.*

There is in fact a matching lower bound to Corollary 1, namely

**Lemma 5.** *The problem Sat-$\mathcal{L}^2$ is NEXPTIME-hard.*

The proof is omitted from these notes. Hence we have

**Theorem 3.** *The problems Sat-$\mathcal{L}^2$ and Sat-$\mathcal{L}^2_{\approx}$ are NEXPTIME-complete.*

*Proof.* Corollary 1 and Lemma 5. □

## 2.3 Fragments with three or more variables

This section is omitted from these notes.

## 2.4 Bibliographic notes

The locus classicus for the Syllogistic—the fragment here denoted $\mathcal{S}$—is Aristotle [2]. The study of the syllogistic dominated logic until the end of the 19th Century. For a comprehensive complexity-theoretic account of the syllogistic and some of its extensions, see Pratt-Hartmann and Moss [29].

The fragment $\mathcal{L}^2_{\approx}$ has an interesting history. Lemma 3, which is due to Scott [34], reduces the problem Sat-$\mathcal{L}^2_{\approx}$ to the satisfiability problem for the so-called Gödel fragment with equality: the set of first-order formulas in prenex-form having quantifier prefix matching $\exists^*\forall\forall\exists^*$. Essentially the same argument reduces Sat-$\mathcal{L}^2$ to the satisfiability problem for the Gödel fragment *without* equality. Gödel [9] had earlier shown that the Gödel fragment (without equality) has the finite model property, and is thus decidable. Unfortunately, Gödel also claimed that adding equality would not affect this result, a claim which was only later shown to be false by Goldfarb [10]. Relying on Gödel's incorrect assertion, Scott claimed to have a proof that $\mathcal{L}^2_{\approx}$ is decidable. (What Scott actually showed was the decidability for $\mathcal{L}_2$ only.) That the full two-variable fragment does indeed have the finite model property was eventually established by Mortimer [23]. The tight NEXPTIME complexity bound was first obtained by Grädel, Kolaitis and Vardi [12], whose proof Lemma 4 repeats.

# Chapter 3

# The One Variable Fragment with Counting

In this chapter, we consider the 1-variable fragment with counting quantifiers, $\mathcal{C}^1$. The main result is that $\mathcal{C}^1$ has the finite model property, and Sat-$\mathcal{C}^1$ is NP-complete. We also consider an NP-complete, proper fragment of $\mathcal{C}^1$, of historical interest.

Recall that $\mathcal{C}^1$ involves a purely relational signature: i.e. there are no individual constants of function symbols. Since, therefore, the only occurrences of $\approx$ in $\mathcal{C}^1$ are in the trivial atom $x \approx x$, we may assume that $\approx$ does not occur in $\mathcal{C}^1$. In fact, we may further assume that the signature contains no predicate of arity greater than 1, since such predicates evidently do not essentially increase expressive power in the presence of only one variable. In this chapter, then, we assume a signature of nullary and unary predicates only.

## 3.1 The one-variable fragment with counting

**Lemma 6.** *Let $\phi$ be a $\mathcal{C}^1$-formula. We can generate, in time bounded by a polynomial function of $\|\phi\|$, a formula*

$$\psi = p_m \wedge \bigwedge_{1 \leq h \leq m} \big(p_h \leftrightarrow \exists_{\bowtie_h C_h} x \beta_h\big), \tag{3.1}$$

*satisfiable over the same domains as $\phi$, where $1 \leq m \leq \|\phi\|$, the $p_h$ are nullary predicates, the $\beta_h$ are quantifier-free $\mathcal{L}^1$-formulas, the $\bowtie_h$ are any of the symbols $\leq$, $\geq$ or $=$, and the $C_h$ are either 0, 1 or occur as a quantifier subscript in $\phi$.*

*Proof.* Let $\phi_0 = \exists_{\geq 1} x \phi$. Thus, $\phi_0$ is satisfied in the same structures as $\phi$.

Suppose $\phi_0$ has a subformula $\theta = \exists_{\leq D} u \chi$, with $\chi$ quantifier-free. Let $p_1$ be a new nullary predicate, let $\phi_1 = \phi[p_1/\theta]$, and let

$$\psi_1 := p_1 \leftrightarrow \exists_{\leq D} v \chi,$$

so that $\phi_1 \wedge \psi_1$ entails $\phi_0$. On the other hand, any model $\mathfrak{A}$ of $\phi_0$ may be expanded to a model of $\phi_1 \wedge \psi_1$ by interpreting $p_1$ to be true in $\mathfrak{A}$ just in case $\mathfrak{A} \models \theta$. Similarly, if $\phi_0$ has a subformula $Qu\chi$, where $Q$ is any of $\exists_{\geq D}$, $\exists_{=D}$, $\forall$, or $\exists$, and $\chi$ quantifier-free, define $\phi_1$ and $\psi_1$ analogously, subject to the obvious adjustments. Now process $\phi_1$ in the same way, and continue until some formula $\phi_m$ is reached which is a single proposition letter $p_m$. Thus, $\phi_0$ and

$$\psi := p_m \wedge \psi_1 \wedge \cdots \wedge \psi_m$$

are satisfied over the same domains. But $\psi$ has the form (3.1), as required. □

It is easy to derive a small model property for $\mathcal{C}^1$ from the above normal form. For convenience, we consider first a special case.

**Lemma 7.** *Let $\phi$ be a $\mathcal{C}^1$-formula of the form*

$$\bigwedge_{1 \leq i \leq m} \exists_{\geq C_i} x \alpha_i \wedge \bigwedge_{1 \leq j \leq m'} \exists_{\leq D_j} x \beta_j, \tag{3.2}$$

*with the $\alpha_i$ and $\beta_j$ quantifier-free. If $\phi$ has a model, then it has a model of size at most $C_1 + \cdots + C_m$.*

*Proof.* For all $i$ $(1 \leq i \leq m)$, select distinct elements $a_{i,1}, \ldots, a_{i,C_i}$ satisfying $\alpha_i$ in $\mathfrak{A}$. (Note that the $a_{i,j}$ and $a_{i',j'}$ are not required to be distinct for $i \neq i'$.) Let $A' = \{a_{i,j} \mid 1 \leq i \leq m,\ 1 \leq j \leq C_i\}$, and let $\mathfrak{A}'$ be the restriction of $\mathfrak{A}$ to $A'$. It is obvious that $\mathfrak{A}' \models \phi$ and that $A' \leq C_1 + \cdots + C_m$. □

**Corollary 2.** *Any satisfiable $\mathcal{C}^1$-formula $\phi$ has a model over a domain of size at most $2^n$ where $n = \|\phi\|$.*

*Proof.* From Lemmas 6 and 7. □

The bound of Lemma 7 is, essentially, the best possible, as the following example shows.

**Example 1.** Let $\phi_k$ be the formula $\exists_{2^k-1} x \top$, $k = 1, 2, \ldots$. Since quantifier subscripts are coded as binary numerals, we have $\|\phi_k\| = k + 3$; however, the smallest model of $\phi_k$ has $2^k - 1$ elements.

Thus, although both $\mathcal{L}^1$ and $\mathcal{C}^1$ have a small model property, the size bound is linear in the former case, but exponential in the latter. This is significant, because, in Chapter 2, we used the linear size bounded to show that Sat-$\mathcal{L}^1$ is in NP: showing that $\phi$ is satisfiable is simply a matter of guessing a structure of size at most $\|\phi\|$ and verifying that it is a model of $\phi$. Evidently, no such simple-minded approach will work for Sat-$\mathcal{C}^1$: the smallest satisfying structures are too big.

It turns out, however, that models of $\mathcal{C}^1$-formulas may admit of more compact descriptions. Fix a signature $\Sigma$ consisting only of the unary predicates

$p_1, \ldots, p_l$. Recalling the notion of 1-type introduced in Chapter 2, there are evidently $L = 2^l$ 1-types over $\Sigma$, which we may list, in some arbitrary order as

$$\pi_1, \ldots, \pi_L.$$

Henceforth, we keep this ordering fixed. Let $\mathfrak{A}$ and $\mathfrak{B}$ be structures interpreting $\Sigma$. It is obvious that $\mathfrak{A}$ and $\mathfrak{B}$ are isomorphic if and only if, for every 1-type $\phi$, the sets $\{a \in A : \mathfrak{A} \models \pi[a]\}$ and $\{b \in B : \mathfrak{B} \models \pi[b]\}$ have the same cardinality. (After all, what it means for two sets to have the same cardinality is that there is a bijection between them.) That is to say, any structure $\mathfrak{A}$ interpreting $\Sigma$ can be characterized, up to isomorphism, by the sequence

$$\nu(\mathfrak{A}) = (\nu_1, \ldots \nu_L),$$

where, for all $j$ $(1 \leq j \leq L)$, $\nu_j$ is the cardinality of the set of elements whose 1-type in $\mathfrak{A}$ is $\pi_j$.

Might representing the structures in this way give us a nondeterministic algorithm for Sat-$\mathcal{C}^1$ that runs in polynomial time? By Lemma 7, we may confine our attention to structures $\mathfrak{A}$ for which every entry in $\nu(\mathfrak{A})$ is at most $2^n$, and hence can be written as a string of at most $n + 1$ bits. The problem, however, is that the number $L$ of 1-types we need to consider is exponentially large.

**Example 2.** Consider the formulas

$$\phi_k := \exists_{=2^k} x \top \wedge \bigwedge_{1 \leq i \leq k} \exists_{=2^{k-1}} x p_i(x), \tag{3.3}$$

for $k > 0$. Each $\phi_k$ has a model in which all the $2^k$ 1-types over the signature $p_1, \ldots, p_k$ are realized exactly once. However, $\phi_k$ also has a model in which only two 1-types—namely, $\{p_1(x), \ldots, p_k(x)\}$ and $\{\neg p_1(x), \ldots, \neg p_k(x)\}$—are realized exactly $2^{k-1}$ times.

Example 2 raises the following question. If a $\mathcal{C}^1$-formula $\phi$ has a model—and hence a finite model—what is the smallest number $N$ such that $\phi$ has a finite model $\mathfrak{A}$ with every component of $\nu(\mathfrak{A})$ bounded by $N$?

## 3.2 Systems of linear equations and inequalities

Consider the $\mathcal{C}^1$-formula

$$\phi := \bigwedge_{1 \leq i \leq m} \exists_{\bowtie_i C_1} \gamma_i, \tag{3.4}$$

where $m > 0$, each $\bowtie_i$ is one of $\leq$, $=$ or $\geq$, and each $\gamma_i$ is a quantifier-free $\mathcal{L}_1$-formula. (Such formulas turn out to be sufficiently general for our purposes.) For all $i$ $(1 \leq i \leq m)$ and $j$ $(1 \leq j \leq L)$, define

$$a_{i,j} = \begin{cases} 1 \text{ if } \models \pi_j \rightarrow \gamma_i \\ 0 \text{ otherwise.} \end{cases}$$

Suppose $\mathfrak{A}$ is a finite structure interpreting the signature of $\phi$. It is obvious that $\mathfrak{A} \models \phi$ if and only if $\nu(\mathfrak{A})$ is simultaneous non-negative integer solution of the system of linear equations and inequalities

$$\begin{array}{cccc} a_{1,1}x_1+ & \ldots+ & a_{1,L}x_L & \bowtie_1 \; C_1 \\ \vdots & & \vdots & \vdots \; \vdots \\ a_{m,1}x_1+ & \ldots+ & a_{m,L}x_L & \bowtie_m \; C_m, \end{array} \tag{3.5}$$

Conversely, given any solution $\nu_1, \ldots \nu_L$ of (3.5) over $\mathbb{N}$ with at least one non-zero value, we can easily construct a model $\mathfrak{A}$ of $\phi$ such that $\nu(\mathfrak{A}) = (\nu_1, \ldots \nu_L)$.

In view of the correspondence between models of $\mathcal{C}^1$-formulas and solutions of systems of linear equations, we need to study the latter. We begin by recalling some classic results.

**Definition 2.** Let $\mathcal{E}$ be a system of linear inequalities of the form (3.5), where the $C_i$ and the $a_{i,j}$ are integers, and the symbols $\bowtie_i$ are one of $\leq$, $=$ or $\geq$. We take the *size* of $\mathcal{E}$, denoted $\|\mathcal{E}\|$, to be measured in the obvious way, assuming binary encoding of integers. *Integer programming feasibility* is the problem of determining, for a given $\mathcal{E}$, whether $\mathcal{E}$ has a solution over $\mathbb{N}$. *Linear programming feasibility* is the problem of determining, for a given $\mathcal{E}$, whether $\mathcal{E}$ has a solution over $\mathbb{Q}$.

**Theorem 4.** *Let $\mathcal{E}$ be a system of linear equations and inequalities of the form* (3.5)*, where the $C_i$ and the $a_{i,j}$ are integers. If $\mathcal{E}$ has a solution over $\mathbb{N}$, then it has a solution in which every value is bounded by a fixed exponential function of $\|\mathcal{E}\|$.*

Theorem 4 will be used in Chapter 4. We mention one striking corollary here.

**Corollary 3.** *Integer programming feasibility is NP complete, and remains so even when all coefficients are required to be either 0 or 1.*

*Proof.* Membership in NP follows from Theorem 4. For if $\mathcal{E}$ is feasible with $\|\mathcal{E}\| = n$, then there is a solution in which all values are bounded by an exponential function of $n$. Using binary encoding, such a solution can be written down and checked in time bounded by a polynomial function of $n$. NP-hardness follows by an easy encoding of SAT. □

By contrast to Corollary 3, we have:

**Theorem 5.** *Linear programming feasibility is in PTIME.*

This result will be used in Chapter 5.

Example 2 directed our attention to models of $\mathcal{C}^1$-formulas in which *as few 1-types as possible are realized*; we now return to this topic. To set the scene, we first recall the following textbook result. Denote the set of *non-negative rationals* by $\mathbb{Q}^+$.

**Lemma 8.** *Let $\mathcal{E}$ be a system of $m$ linear equations with rational coefficients. If $\mathcal{E}$ has a solution over $\mathbb{Q}^+$, then $\mathcal{E}$ has a solution over $\mathbb{Q}^+$ with at most $m$ non-zero entries.*

*Proof.* We can write $\mathcal{E}$ as $\mathbf{Ax} = \mathbf{c}$, where $\mathbf{A}$ is a rational matrix with $m$ rows and, say, $L$, columns, and $\mathbf{c}$ is a rational column vector of length $m$. If $\mathbf{b}$ is any solution of $\mathcal{E}$ in $\mathbb{Q}^+$ with $k > m$ non-zero entries, the $k$ columns of $\mathbf{A}$ corresponding to these non-zero entries must be linearly dependent. Thus, there exists a non-zero rational vector $\mathbf{b}'$ with zero-entries wherever $\mathbf{b}$ has zero-entries, such that $\mathbf{Ab}' = \mathbf{0}$. But then it is easy to find a rational number $\varepsilon$ such that $\mathbf{b} + \varepsilon\mathbf{b}'$ is a solution of $\mathcal{E}$ in $\mathbb{Q}^+$ with fewer than $k$ non-zero entries. □

The question naturally arises as to the whether this bound is available when solutions are sought in $\mathbb{N}$, rather than $\mathbb{Q}^+$. The following result shows that it is not.

**Lemma 9.** *Fix $m \geq 6$. Let $\mathbf{A}$ be the $m \times (m+1)$-matrix given by*

$$\mathbf{A} = \left(\begin{array}{ccccccc|ccccccc} 1 & 1 & 1 & 0 & 0 & 0 & 0 & 0 & \dots & & & & & 0 \\ 0 & 1 & 1 & 1 & 0 & 0 & 0 & 0 & \dots & & & & & 0 \\ 0 & 0 & 1 & 1 & 1 & 0 & 0 & 0 & \dots & & & & & 0 \\ 0 & 0 & 0 & 1 & 1 & 1 & 0 & 0 & \dots & & & & & 0 \\ 0 & 0 & 0 & 0 & 1 & 1 & 1 & 0 & \dots & & & & & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & 1 & 1 & \dots & & & & & 0 \\ \vdots & & & & & & & & & & & & & \vdots \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & \dots & 0 & 1 & 1 & 1 & 1 \\ \hline 1 & 1 & 0 & 1 & 0 & 0 & 1 & 0 & \dots & & & & & 0 \end{array}\right),$$

*in which a pattern of three 1s is shifted right across the first $(m-1)$ rows, and the last row contains the seven entries shown on the left followed by $(m-6)$ 0s. Let $\mathbf{c}$ be the column vector of length $m$ given by*

$$\mathbf{c} = (3, 3, \dots, 3, 4)^T$$

*consisting of $(m-1)$ 3s and a single 4. Then the unique solution of the system of equations $\mathbf{Ax} = \mathbf{c}$ over $\mathbb{N}$ is the column vector $(1, \dots, 1)^T$ consisting of $(m+1)$ 1s.*

*Proof.* Evidently, $\mathbf{A}(1, \dots, 1)^T = \mathbf{c}$. Conversely, suppose $\mathbf{b} = (b_1, \dots, b_{m+1})^T$ is any solution of $\mathbf{Ax} = \mathbf{c}$ in $\mathbb{N}$. From the first row of $\mathbf{A}$, $b_1 + b_2 + b_3 = 3$, whence $b_1, b_2, b_3$ are either (i) the integers $0, 0, 3$ in some order, or (ii) the integers $0, 1, 2$ in some order or (iii) the integers $1, 1, 1$. By considering rows 2 to $m-1$ of $\mathbf{A}$, it is then easy to see that, in every case, these three values must recur, in the same order, to the end of the vector: that is, $\mathbf{b}$ must have the form

$$(b_1, b_2, b_3, b_1, b_2, b_3, b_1, \dots)^T.$$

From the last row of $\mathbf{A}$, then, $3b_1 + b_2 = 4$. Thus, $b_1, b_2, b_3$ are certainly not $3, 0, 0$, in any order. Suppose, then, $b_1, b_2, b_3$ are $0, 1, 2$, in some order. If $b_1 = 0$,

then $3b_1 + b_2$ is at most 2; if $b_1 = 1$, then $3b_1 + b_2$ equals either 3 or 5; and if $b_1 = 2$, then $3b_1 + b_2$ is at least 6. Thus, $b_1, b_2, b_3$ are not $0, 1, 2$, in any order, whence $\mathbf{b} = (1, \ldots, 1)^T$ as required. □

It follows that the argument of Lemma 8 cannot work for solutions over $\mathbb{N}$. Fortunately, however, an alternative approach is available.

**Definition 3.** A *Boolean equation* is any equation of the form $a_1x_1 + \cdots a_nx_n = C$, where each $a_i$ $(1 \leq i \leq n)$ is either 0 or 1, and $C$ is a natural number.

**Lemma 10.** *Let $\mathcal{E}$ be a system of $m$ Boolean equations in $L$ variables. If $\mathcal{E}$ has a solution over $\mathbb{N}$, then $\mathcal{E}$ has a solution over $\mathbb{N}$ with at most $m \log(L+1)$ non-zero entries.*

*Proof.* We write $\mathcal{E}$ as $\mathbf{Ax} = \mathbf{c}$, where $\mathbf{A}$ is a matrix of 0s and 1s with $m$ rows and $L$ columns, $\mathbf{c}$ is a column vector over $\mathbb{N}$ of length $m$, and $\mathbf{x} = (x_1, \ldots x_L)^T$. If $\mathcal{E}$ has a solution over $\mathbb{N}$, let $\mathbf{b} = (b_1, \ldots, b_L)^T$ be such a solution with a minimal number $k$ of non-zero entries. We show that

$$k \leq m \log(L+1). \tag{3.6}$$

This condition is trivially satisfied if $k = 0$, so assume $k > 0$. Furthermore, by renumbering the variables if necessary, we may assume without loss of generality that $b_j > 0$ for all $j$ $(1 \leq j \leq k)$. Now, if $I \subseteq \{1, \ldots, k\}$, define $\mathbf{v}_I$ to be the $m$-element column vector $(v_1, \ldots, v_m)^T$, where

$$v_i = \sum_{j \in I} \mathbf{A}_{i,j}.$$

That is, $\mathbf{v}_I$ is the sum of those columns of $\mathbf{A}$ indexed by elements of $I$. Since each $v_i$ $(1 \leq i \leq m)$ is a natural number satisfying

$$v_i \leq L, \tag{3.7}$$

the number of vectors $\mathbf{v}_I$ (as $I$ varies over subsets of $\{1, \ldots, k\}$) is certainly bounded by $(L+1)^m$. So suppose, for contradiction, that $k > m \log(L+1)$. Then $2^k > (L+1)^m$, whence there must exist distinct subsets $I$, $I'$ of $\{1, \ldots, k\}$ such that $\mathbf{v}_I = \mathbf{v}_{I'}$. Setting $J = I \setminus I'$ and $J' = I' \setminus I$, it is evident that $J$ and $J'$ are distinct (and disjoint), again with $\mathbf{v}_J = \mathbf{v}_{J'}$. By interchanging $J$ and $J'$ if necessary, we may assume that $J \neq \emptyset$. Now define, for all $j$ $(1 \leq j \leq L)$:

$$b'_j = \begin{cases} b_j - 1 \text{ if } j \in J \\ b_j + 1 \text{ if } j \in J' \\ b_j \text{ otherwise,} \end{cases}$$

and write $\mathbf{b}' = (b'_1, \ldots, b'_L)^T$. Since $J$ and $J'$ are disjoint, the cases do not overlap; and since the $b_j$ are all positive $(1 \leq j \leq k)$, the $b'_j$ all lie in $\mathbb{N}$. Moreover,

$$\mathbf{Ab}' = \mathbf{Ab} - \mathbf{v}_J + \mathbf{v}_{J'} = \mathbf{Ab}.$$

Since $J$ is nonempty, $\min\{b'_j | j \in J\}$, is strictly smaller than $\min\{b_j | j \in J\}$. Generating $\mathbf{b}''$, $\mathbf{b}'''$, etc. in this way (using the same $J$ and $J'$) will thus eventually result in a vector—say, $\mathbf{b}^*$—with strictly fewer non-zero entries than $\mathbf{b}$, but with $\mathbf{A}\mathbf{b}^* = \mathbf{A}\mathbf{b}$—a contradiction. □

We now strengthen Lemma 10 to obtain a bound which does not depend on $L$.

**Lemma 11.** *Let $\mathcal{E}$ be a system of $m$ Boolean equations. If $\mathcal{E}$ has a solution over $\mathbb{N}$, then $\mathcal{E}$ has a solution over $\mathbb{N}$ with at most $\frac{5}{2}m\log m + 1$ non-zero entries.*

*Proof.* The case $m = 1$ is trivial: if $\mathcal{E}$ has a solution, then it has a solution with at most one non-zero entry. So assume henceforth that $m > 1$.

In the proof of Lemma 10, the inequality (3.7) can evidently be strengthened to

$$v_i \leq k.$$

Proceeding exactly as for Lemma 10, we obtain, in place of (3.6), the inequality

$$k \leq m\log(k+1).$$

Hence, for $k$ positive, we have

$$\frac{k}{\log(k+1)} \leq m. \tag{3.8}$$

Now the left-hand side of (3.8) is greater than or equal to unity, and since the function $x \mapsto x\log x$ is monotone increasing for $x \geq e^{-1}$, we can apply it to both sides of (3.8) to obtain

$$kZ(k) \leq m\log m, \tag{3.9}$$

where, for all $k > 0$,

$$Z(k) = \frac{\log k - \log\log(k+1)}{\log(k+1)}.$$

It is straightforward to check that $Z$ is monotone increasing on the positive integers, and that $Z(k) \to 1$ as $k \to \infty$. (Indeed, for $x > 0$, the function $x \mapsto \log x/\log(x+1)$ is monotone increasing with limit 1 as $x$ tends to $\infty$; and for $x \geq 2^e - 1$, the function $x \mapsto \log\log(x+1)/\log(x+1)$ is monotone decreasing with limit 0.)

We may now establish that $k \leq \frac{5}{2}m\log m + 1$. Calculation shows that $1/Z(7) \approx 2.4542 < \frac{5}{2}$. Therefore, since $Z$ is monotone increasing, (3.9) yields, for $k \geq 7$, the inequalities $k \leq m\log m/Z(k) \leq m\log m/Z(7) < \frac{5}{2}m\log m$. Obviously, if $k \leq 6$, we have $k \leq \frac{5}{2}m\log m + 1$, since $m \geq 2$ by assumption. □

The proof of Lemma 11 actually shows a little more than advertised: for *any* real $c > 1$, there exists a $d$ such that, if $\mathcal{E}$ is a system of $m$ Boolean equations with a solution over $\mathbb{N}$, then $\mathcal{E}$ has a solution over $\mathbb{N}$ with at most $cm\log m + d$ non-zero entries. (As $c$ approaches unity, the required value of $d$ given by the

above proof quickly becomes astronomical.) It follows that none of these bounds is optimal, in the sense of being achieved infinitely often.

Returning to the main business of this chapter, we have:

**Lemma 12.** *The satisfiability (= finite satisfiability) problem for $\mathcal{C}^1$ is in NP.*

*Proof.* Let $\phi$ be a given $\mathcal{C}^1$-formula. Compute the formula $\psi$ given in (3.1), Lemma 6, and guess an assignment of the truth-values for any nullary predicates occurring on $\psi$. Carrying out routine logical simplification, we obtain a formula of the form

$$\bigwedge_{1 \leq h \leq m} \exists_{\bowtie_h C_h} \gamma_h,$$

giving rise to a system $\mathcal{E}$ of linear inequalities such that $\mathcal{E}$ has a non-zero solution over $\mathbb{N}$ if and only if $\psi$ is satisfiable. By Lemma 11, $\mathcal{E}$ has a non-zero solution over $\mathbb{N}$ if and only if it has a non-zero solution with $k \leq \frac{5}{2} m \log m$ non-zero values. Now simply guess a $k \leq \frac{5}{2} m \log m$ and set of indices $J = \{j_1, \ldots, j_k\}$ in the range between 1 and $L$, and let $\mathcal{E}'$ be the system of equations and inequalities obtained from $\mathcal{E}$ by ignoring those terms $a_{i,j} x_j$ for which $j \notin J$. The system $\mathcal{E}'$ can be computed in polynomial time; and $\phi$ is satisfiable if and only if $\mathcal{E}'$ has a non-zero solution over $\mathbb{N}$. The result then follows by Corollary 3. □

## 3.3 The numerically definite syllogistic

Define the fragment $\mathcal{N}^1$ to be the set of $\mathcal{C}^1$-formulas of the forms

$$\begin{array}{ll} \exists_{\geq C} x(p(x) \wedge q(x)) & \exists_{\geq C} x(p(x) \wedge \neg q(x)) \\ \exists_{\leq C} x(p(x) \wedge q(x)) & \exists_{\leq C} x(p(x) \wedge \neg q(x)), \end{array} \tag{3.10}$$

where $p$ and $q$ are unary predicates. This fragment is of some historical and linguistic interest, because we can think of unary predicates as corresponding to common nouns, and the formulas (3.10) to English sentences of the forms

$$\begin{array}{ll} \mathsf{At\ least}\ C\ p\ \mathsf{are}\ q & \mathsf{At\ least}\ C\ p\ \mathsf{are\ not}\ q \\ \mathsf{At\ most}\ C\ p\ \mathsf{are}\ q & \mathsf{At\ most}\ C\ p\ \mathsf{are\ not}\ q, \end{array} \tag{3.11}$$

respectively. (We have simplified the presentation here by ignoring the issue of singular/plural agreement, which has no logical or computational significance.) We call the fragment of English defined by these sentence-forms the *numerically definite syllogistic*, loosely following the terminology of de Morgan.

The sentence $\mathsf{Some}\ p\ \mathsf{are}\ q$ may be equivalently written $\mathsf{At\ least}\ 1\ p\ \mathsf{is\ a}\ q$, and the sentence $\mathsf{All}\ p\ \mathsf{are}\ q$ may be equivalently—if somewhat unidiomatically—written $\mathsf{At\ most}\ 0\ p\ \mathsf{are\ not}\ q$. Thus, the numerically definite syllogistic generalizes the ordinary syllogistic mentioned in Chapter 2. We remark that the sentence $\mathsf{There\ are\ at\ least}\ C\ p$ may be equivalently written $\mathsf{At\ least}\ C\ p\ \mathsf{are}\ p$; and similarly for $\mathsf{There\ are\ at\ most}\ C\ p$. So these sentences too are expressible in $\mathcal{N}^1$.

We now proceed to establish a lower complexity bounds for $\mathcal{N}^1$. Recall that a *graph* is a pair $G = (V, E)$ where $V$ is a finite set (the 'nodes' of $G$) and $E$ is a set of 2-element subsets of $V$ (the 'edges' of $G$). Note that graphs, in this sense, have no 'loops' or 'multiple edges'. A 3-*colouring* of $G$ is a function $t$ mapping the nodes of $G$ to the set $\{0, 1, 2\}$ such that no edge of $G$ joins two nodes mapped to the same value. We say that $G$ is 3-*colourable* if a 3-colouring of $G$ exists. The problem of deciding whether a given graph $G$ is 3-colourable is well-known to be NP-hard. We reduce it to $\mathcal{C}^1$-satisfiability.

**Lemma 13.** *The satisfiability problem for $\mathcal{N}^1$ is NP-hard.*

*Proof.* By reduction of graph 3-colourability. Let $G = (V, E)$ be a graph, and assume without loss of generality that $V = \{1, \ldots, n\}$. For all $i$ $(1 \leq i \leq n)$ and $k$ $(0 \leq k < 3)$, let $p_i^k$ be a fresh unary predicate. Think of $p_i^k(x)$ as saying: "$x$ is a colouring of $G$ in which node $i$ has colour $k$". Let $\Phi_G$ be the set of $\mathcal{C}^1$-formulas consisting of

$$\exists_{\leq 3} x(p(x) \wedge p(x)) \tag{3.12}$$

$$\{\exists_{\leq 0} x(p_i^j(x) \wedge p_i^k(x)) \mid 1 \leq i \leq n, 0 \leq j < k < 3\} \tag{3.13}$$

$$\{\exists_{\geq 1} x(p_i^k(x) \wedge p(x)) \mid 1 \leq i \leq n, 0 \leq k < 3\} \tag{3.14}$$

$$\{\exists_{\leq 0} x(p_i^k(x) \wedge p_j^k(x)) \mid (i, j) \text{ is an edge of } G, 0 \leq k < 3\} \tag{3.15}$$

We prove that $\Phi_G$ is satisfiable if and only if $G$ is 3-colourable.

Suppose $\mathfrak{A} \models \Phi_G$. By (3.12), $|p^{\mathfrak{A}}| \leq 3$. Fix any $i$ $(1 \leq i \leq n)$. No $a \in p^{\mathfrak{A}}$ satisfies any two of the predicates $p_i^0$, $p_i^1$, $p_i^2$, by (3.13); on the other hand, each of these predicates is satisfied by at least one element of $p^{\mathfrak{A}}$, by (3.14); therefore, $|p^{\mathfrak{A}}| = 3$, and each element $a$ of $p^{\mathfrak{A}}$ satisfies exactly one of the predicates $p_i^0$, $p_i^1$, $p_i^2$. Now fix any $a \in p^{\mathfrak{A}}$, and, for all $i$ $(1 \leq i \leq n)$, define $t_a(i)$ to be the unique $k$ $(1 \leq k < 3)$ such that $\mathfrak{A} \models p_i^k[a]$, by the above argument. The formulas (3.15) then ensure that $t_a$ defines a colouring of $G$. Conversely, suppose that $t : \{1, \ldots, n\} \to \{0, 1, 2\}$ defines a colouring of $G$. Let $\mathfrak{A}$ be a structure with domain $A = \{0, 1, 2\}$; let all three elements satisfy $p$; and, for all $k \in A$, let $p_i^k$ be satisfied by the single element $k + t(i)$ (where the addition is modulo 3). It is routine to verify that $\mathfrak{A} \models \Phi_G$. □

**Theorem 6.** *The satisfiability problem for any fragment between the numerically definite syllogistic and the 1-variable fragment with counting is NP-complete.*

*Proof.* Lemmas 12 and 13. □

Notice that the only numerical subscripts mentioned in the encoding of graph colourability in Lemma 13 are 0, 1 and 3. It follows that this lemma continues to hold even if if the size of a positive numerical quantifier subscript $C$ is taken to be $C$, rather than $\lfloor \log C \rfloor + 1$. As we say, the problem Sat-$\mathcal{N}^1$ is *strongly NP-complete*: it is NP-complete, and remains so under non-succinct coding of numerical inputs.

## 3.4 Bibliographic notes

The system $\mathcal{N}$, which we have called the *numerically definite syllogistic*, corresponds loosely to the logic investigated by de Morgan [6]. The semantic apparatus in terms of which satisfiability (or, dually, validity) is understood today was of course not available to de Morgan, who presents a collection of *numerically definite syllogisms*—in effect, 2-premise proof-rules for the fragment $\mathcal{N}$. De Morgan's system is certainly incomplete with respect to the semantics given here. Similar systems (likewise incomplete) have been proposed by Murphree [24] and Hacker and Parry [14]. For impossibility results on complete systems of numerical syllogisms, see Pratt-Hartmann [32], Sec. 5, and [33].

Theorem 4 was proved by Borosh and Treybig [4]. Theorem 5 was proved by Khachiyan [17].

The problem PSAT was shown to be in NP by Georgakopoulos *et al.* [8]. (For a good textbook treatment, see Paris [28], Chapter 10.) Theorem 6 was first proved, independently, by Kuncak and Rinard [19] and Pratt-Hartmann [32]; however, the fundamental combinatorial idea behind the proof actually appears in Eisenbrand and Shmonina [7].

# Chapter 4

# The Two-Variable Fragment with Counting

## 4.1 Normal forms

In this chapter, we consider the 2-variable fragment with counting quantifiers, over a signature of unary and binary predicates, $\mathcal{C}^2$. The main result of this chapter is that the problems Sat-$\mathcal{C}^2$ and Fin-Sat-$\mathcal{C}^2$ are both NEXPTIME-complete.

The theorems and lemmas in this chapter all continue to hold for signatures also featuring individual constants, nullary predicates or predicates of arity 3 or more. However, the addition of these non-logical primitives to $\mathcal{C}^2$ would complicate the exposition, and would not essentially increase expressive power. Consequently, we ignore them in the sequel. By contrast, adding function-symbols to the signature would invalidate the key results of this chapter. In the presence of counting quantifiers, equality is definable by the formula $\forall x e(x, x) \wedge \forall x \exists_{\leq 1} y e(x, y)$. Consequently, we shall assume that the equality predicate $\approx$ is available in $\mathcal{C}^2$.

The Löwenheim-Skolem-Tarski theorem guarantees that, if a first-order formula has a model, then it has a finite or countably infinite model. Since we are concerned with the satisfiability and finite satisfiability of first-order formulas in this chapter, we shall silently take all structures to be either finite or countably infinite.

The fragment $\mathcal{C}^2$ presents us with much greater challenges than any of the fragments considered so far. For a start, $\mathcal{C}^2$ lacks the finite model property:

**Example 3.** The $\mathcal{C}^2$-formula $\phi$ given by

$$\forall x \exists y s(x, y) \wedge \forall x \exists_{\leq 1} y s(y, x) \wedge \exists x \forall y \neg s(y, x)$$

is satisfiable, but not finitely satisfiable. For let $A = \mathbb{N}$ and $s^{\mathfrak{A}} = \{\langle i, i+1 \rangle \mid i \in \mathbb{N}\}$. Then $\mathfrak{A} \models \phi$. On the other hand, suppose $\mathfrak{B} \models \phi$. Since $\mathfrak{B} \models \forall x \exists y s(x, y)$,

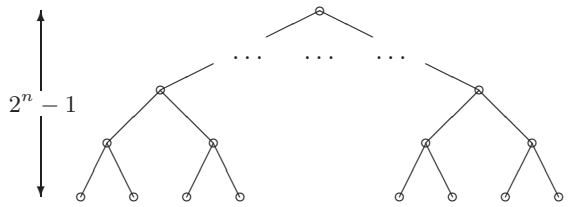


Figure 4.1: The tree $T_n$.

let $f : B \to B$ be a function such that $f \subseteq s^{\mathfrak{B}}$. Since $\mathfrak{B} \models \forall x \exists_{\leq 1} y s(y, x)$, $f$ is 1–1. Since $\mathfrak{B} \models \exists x \forall y \neg s(y, x)$, $f$ is not onto. Therefore $B$ is infinite. □

Actually, matters are even worse than Example 3 might lead one to expect. It turns out that even finitely satisfiable $\mathcal{C}^2$ formulas can require huge models.

**Example 4.** For $n \geq 0$, let $T_n$ be the complete, binary tree of depth $2^n - 1$, as depicted in Fig. 4.1. (Note: if $a$ is a node in any tree $T$, we take the *depth* of $a$, denoted $d(a)$, to be the number of edges on the unique path from $a$ to the root of $T$; and we take the depth of $T$ to be the depth of the deepest node.) Thus, $T_n$ contains $2^{2^n} - 1$ elements. Recalling the representation of a natural number $n \leq 2^n - 1$ as a string of binary digits $d_{n-1}, \ldots, d_0$ (where the zeroth digit $d_0$ is the least significant), we interpret unary predicates $X_0, \ldots, X_{n-1}$ and $X_0^*, \ldots, X_n^*$ and binary predicate $r$ over $T_n$ to form a model $\mathfrak{T}_n$, as follows. Let $X_i$ be satisfied by a node $a$ just in case the $i$th digit of $d(a)$ is 1. Likewise, for $i < n$, let $X_i^*$ be satisfied by $a$ just in case $d(a) < 2^n - 1$ and the least significant zero-digit in $d(a)$ is the $i$th digit; and let $X_n^*$ be satisfied by $a$ just in case $d(a) = 2^n - 1$. Finally, let $r$ be satisfied by the pair of nodes $\langle a, b \rangle$ just in case $b$ is a daughter of $a$.

Let $\phi_{n,0}$ be the conjunction of all formulas

$$\forall x \big( X_i^*(x) \leftrightarrow \neg X_i(x) \wedge \bigwedge_{0 \leq j < i} X_j(x) \big),$$

for $0 \leq i < n$, together with the formula

$$\forall x \big( X_n^*(x) \leftrightarrow \bigwedge_{0 \leq j < n} X_j(x) \big).$$

Thus, $\phi_{n,0}$ fixes the interpretations of the $X_i^*$ in terms of the $X_i$ in the required way. Let $\phi_{n,1}$ be the conjunction of the following formulas

$$\forall x \forall y \big( r(x, y) \rightarrow (X_i^*(x) \rightarrow (\neg X_j(y) \wedge X_i(y))) \big)$$

for $0 \leq j < i < n$, and

$$\forall x \forall y \big( r(x, y) \rightarrow (X_i^*(x) \rightarrow (X_j(x) \leftrightarrow X_j(y))) \big)$$

for $0 \leq i < j < n$. Recalling the standard algorithm for incrementing binary numerals, we see that $\phi_{n,1}$ asserts that the depth of a daughter of a node in $T$ is one greater than the depth of that node. Finally, let $\phi_{n,2}$ be the conjunction of the formulas

$$\exists x \bigwedge_{0 \leq i < n} \neg X_i(x) \qquad \forall x(\neg X_n^*(x) \to \exists_{\geq 2} y r(x,y)) \qquad \forall x \exists_{\leq 1} y r(y,x).$$

Thus, $\phi_{n,2}$ asserts that there is a node of depth zero, that each node of depth less than $2^n - 1$ has at least two daughters, and that no node has more that one mother. Let $\phi_n = \phi_{n,0} \wedge \phi_{n,1} \wedge \phi_{n,2}$. Thus, $\mathfrak{T}_n \models \phi_n$. Conversely, it is evident that every model of $\phi_n$ contains an isomorphic copy of $\mathfrak{T}_n$.

We have shown that each $\phi_n$ is finitely satisfiable, but not satisfiable over any domain containing fewer than $2^{2^n}$ elements. On the other hand, inspection of the above formulas shows that $\|\phi_n\|$ grows as a polynomial function of $n$. □

Example 3 shows that the problems Sat-$\mathcal{C}^2$ and Fin-Sat-$\mathcal{C}^2$ do not coincide. Obviously, then, we cannot hope to establish the decidability of Sat-$\mathcal{C}^2$ by describing a procedure for guessing models of bounded size. Example 4 shows that even finitely satisfiable $\mathcal{C}^2$-formulas in general require models of doubly-exponential size. Again, it follows that we cannot hope to establish that Fin-Sat-$\mathcal{C}^2$ is in NEXPTIME by describing a procedure for guessing models of exponential size. We remark that, in the course of the chapter, a matching upper bound is obtained to Example 4: any finitely satisfiable $\mathcal{C}^2$ formula $\phi$ is shown to have a model whose size is bounded by a doubly exponential function of $\|\phi\|$. (*A priori*, this is by no means obvious.)

One technique of previous chapters that does usefully apply to $\mathcal{C}^2$, however, is the existence of equisatisfiable Scott-form formulas.

**Lemma 14.** *Let $\phi$ be a $\mathcal{C}^2$-formula. We can generate, in time bounded by a polynomial function of $\|\phi\|$, a quantifier-free $\mathcal{L}^2$-formula $\alpha$, a list of positive integers $C_1, \ldots, C_m$ and a list of binary predicates $f_1, \ldots, f_m$ such that the formulas $\phi$ and*

$$\psi = \forall x \forall y(\alpha \vee x \approx y) \wedge \bigwedge_{1 \leq h \leq m} \forall x \exists_{=C_h} y(f_h(x,y) \wedge x \not\approx y) \tag{4.1}$$

*are satisfiable over the same domains containing at least $C+1$ elements, where $C = \max_h C_h$.*

*Proof.* We construct the formula $\psi$ in stages.

**Stage 1:** Let $\phi_0 = \forall x \exists y \exists x \phi$. (Thus, $\phi_0$ and $\phi$ are satisfied in exactly the same structures.) In the sequel, we take $u$, $v$ to denote the variables $x$, $y$, in either order. Suppose $\phi_0$ possesses a subformula $\theta(u) = \exists_{\leq D} v \chi$, with $\chi$ quantifier-free. Let $p$ be a new unary predicate, and $r_1$, $r_2$ new binary predicates. Define $\phi_1 := \phi_0[p(u)/\theta(u)]$ and

$$\psi_1 := \forall u \exists_{=D} v r_1(u,v) \wedge \forall u \exists_{=D+1} v r_2(u,v) \wedge$$
$$\forall u \forall v(p(u) \to (\chi \to r_1(u,v))) \wedge \forall u \forall v(\neg p(u) \to (r_2(u,v) \to \chi)).$$

We claim that $\phi_0$ and $\phi_1 \wedge \psi_1$ are satisfiable over the same domains containing at least $D+1$ elements. On the one hand, we have $\models \phi_1 \wedge \psi_1 \rightarrow \phi_0$. On the other, if $\mathfrak{A} \models \phi_0$ with $|A| \geq D+1$, expand $\mathfrak{A}$ to a structure $\mathfrak{A}'$ as follows. Set $\mathfrak{A}' \models p[a]$ if and only if $\mathfrak{A} \models \theta[a]$. For all $a \in A$, if $\mathfrak{A} \models \theta[a]$, set $\mathfrak{A}' \models r_1[a,b]$ for exactly $D$ objects $b \in A$ *including all those* $b$ *such that* $\mathfrak{A} \models \chi[a,b]$; and set $\mathfrak{A}' \models r_2[a,b]$ for exactly $D+1$ (randomly chosen) objects $b \in A$. Likewise, if $\mathfrak{A} \not\models \theta[a]$, set $\mathfrak{A}' \models r_1[a,b]$ for exactly $D$ (randomly chosen) objects $b \in A$; and set $\mathfrak{A}' \models r_2[a,b]$ for exactly $D+1$ objects $b \in A$ such that $\mathfrak{A} \models \chi[a,b]$. It is routine to check that $\mathfrak{A}' \models \phi_1 \wedge \psi_1$. Hence, $\phi_0$ and $\phi_1 \wedge \psi_1$ are satisfiable over the same domains of size $D+1$ or more, as claimed. Similarly, if $\phi_0$ contains a subformula of the forms $Qu\chi$, with $Q$ any of $\exists_{\geq D}$ $\exists_{=D}$, $\forall v$ or $\exists v$, proceed analogously, subject to the obvious adjustments. Now process $\phi_1$ in the same way, and continue until some formula $\phi_k$ is reached having the form $\forall x \chi$, with $\chi$ quantifier-free. Set

$$\psi' := \phi_k \wedge (\psi_k \wedge \psi_{k-1} \wedge \cdots \wedge \psi_1).$$

Thus, $\psi'$ and $\phi_0$ (and hence $\phi$) are satisfiable over the same domains of size at least $C'$, where $C'$ is the largest quantifier subscript occurring in any of the $\psi_i$.

**Stage 2:** By performing various trivial logical manipulations, we way take $\psi'$ to be a conjunction of formulas of the forms

$$\forall x \forall y \pi \qquad\qquad \forall x \exists_{=D} y s(x,y),$$

where $\pi$ is a quantifier-free $\mathcal{L}^2_{\approx}$-formula, $s$ is a binary predicate and $D > 0$.

Let $\psi''$ be the result of replacing any conjuncts $\forall x \exists_{=D} y s(x,y)$ of $\psi'$ by the four corresponding conjuncts

$$\begin{array}{ll} \forall x \exists_{=D-1} y(s'(x,y) \wedge x \not\approx y) & \forall x \forall y(s(x,x) \rightarrow (s(x,y) \leftrightarrow s'(x,y))) \\ \forall x \exists_{=D} y(s''(x,y) \wedge x \not\approx y) & \forall x \forall y(\neg s(x,x) \rightarrow (s(x,y) \leftrightarrow s''(r,y))), \end{array}$$

where $s'$ and $s''$ are new binary predicates. By an argument similar to that employed above, $\psi'$ and $\psi''$ are satisfiable over the same domains containing at least $D+1$ elements. And, modulo some further trivial logical re-arrangement, we may write

$$\psi'' = \forall x \forall y \alpha'(x,y) \wedge \bigwedge_{1 \leq h \leq m} \forall x \exists_{=C_h} y(f_h(x,y) \wedge x \not\approx y),$$

where $\alpha'(x,y)$ is a quantifier-free $\mathcal{L}^2_{\approx}$-formula. Thus, $\phi$ and $\psi''$ are satisfiable over the same domains containing at least $C+1$ elements, where $C = \max_h C_h$.

**Stage 3:** It remains only to reform the occurrences of $\approx$ in $\alpha'(x,y)$. We proceed exactly as in Lemma 3, obtaining a quantifier-free $\mathcal{L}^2$-formula $\alpha$ such that, restricting attention to domains containing at least 2 elements,

$$\forall x \forall y \alpha'(x,y) \equiv \forall x \forall y(\alpha(x,y) \vee x \approx y)$$

For such domains, then, $\psi''$ is logically equivalent to the formula $\psi$ given in (4.1). Moreover, it is obvious that the above computation can be effected in time bounded by a polynomial function of $\|\phi\|$. □

## 4.2 Classified signatures and featured predicates

Lemma 14 allows us to restrict attention $\mathcal{C}^2$-formulas of the form

$$\psi = \forall x \forall y(\alpha \vee x \approx y) \wedge \bigwedge_{1 \leq h \leq m} \forall x \exists_{=C_h} y(f_h(x,y) \wedge x \not\approx y),$$

where $\alpha$ is a quantifier-free $\mathcal{L}^2$-formula and the $f_1, \ldots, f_m$ are binary predicates. In the analysis of models of such formulas, the binary predicates $f_1, \ldots, f_m$ play a special role. Accordingly, we adopt the following terminology.

**Definition 4.** Let $\Sigma$ be signature of unary and binary predicates, and $f_1, \ldots, f_m$ ($m > 0$) a list of pairwise distinct binary predicates in $\Sigma$. The pair $\langle \Sigma, (f_1, \ldots, f_m) \rangle$ is called a *classified signature*, and $f_1, \ldots, f_m$ are referred to as its *featured predicates*.

The following definitions illustrate how classified signatures will be used in the ensuing argument.

**Definition 5.** Let $\mathfrak{A}$ be a structure interpreting a classified signature $\langle \Sigma, \bar{f} \rangle$, and let $a \in A$. The *degree* of $a$ in $\mathfrak{A}$ is the maximum number of elements of $A$ to which $a$ is non-reflexively related by any of the featured predicates:

$$\deg(a, \mathfrak{A}) = \max_f |\{b \in A \setminus \{a\} \mid \mathfrak{A} \models f[a,b]\}|,$$

where $f$ ranges over the featured predicates $\bar{f}$. If, for some finite $C$, $\deg(a, \mathfrak{A}) \leq C$ for all $a \in A$, then we say $\mathfrak{A}$ is of *finite degree*, and write $\deg(\mathfrak{A})$ to denote the smallest such $C$.

**Definition 6.** Let $\langle \Sigma, \bar{f} \rangle$ be a classified signature, and let $\tau$ be a 2-type over $\Sigma$. We say that $\tau$ is a *message-type* (over $\Sigma$) if $f(x,y) \in \tau$ for some featured predicate $f$. If $\tau$ is a message-type such that $\tau^{-1}$ is also a message-type, we say that $\tau$ is *invertible*. On the other hand, if $\tau$ is a 2-type such that neither $\tau$ nor $\tau^{-1}$ is a message-type, we say that $\tau$ is a *silent* 2-type.

Thus, a 2-type $\tau$ is an invertible message-type if and only if there are featured predicates $f$ and $f'$ such that $f(x,y) \in \tau$ and $f'(y,x) \in \tau$. The terminology is meant to suggest the following imagery. Let $\mathfrak{A}$ be a structure interpreting the classified signature in question. If $\mathrm{tp}^{\mathfrak{A}}[a,b]$ is a message-type $\mu$, then we may imagine that $a$ sends a message (of type $\mu$) to $b$. If $\mu$ is invertible, then $b$ replies by sending a message (of type $\mu^{-1}$) back to $a$. If $\mathrm{tp}^{\mathfrak{A}}[a,b]$ is silent, then neither element sends a message to the other.

At this point, we are ready for a technical lemma regarding structures interpreting classified signatures; this lemma will be useful in several places in the ensuing argument.

**Lemma 15.** *Let $\mathfrak{A}$ be a structure of finite degree interpreting a classified signature $\langle \Sigma, \bar{f} \rangle$ with $m$ featured predicates, and let $C = \deg(\mathfrak{A})$. Suppose that $\pi$ and $\pi'$ are 1-types over $\Sigma$ (not necessarily distinct), both realized in $\mathfrak{A}$ more than $(mC+1)^2$ times. Then there exist distinct elements $b, b'$ with $\mathrm{tp}^{\mathfrak{A}}[b] = \pi$ and $\mathrm{tp}^{\mathfrak{A}}[b'] = \pi'$, such that $\mathrm{tp}^{\mathfrak{A}}[b,b']$ is a silent 2-type.*

*Proof.* Let $B$ be a set of elements having 1-type $\pi$, with $|B| = (mC)^2 + mC + 1$, and let $B'$ be a set of elements having 1-type $\pi'$, disjoint from $B$, with $|B'| = mC + 1$. (Note that, since $|B| + |B'| = (mC + 1)^2 + 1$, and since $\pi$ and $\pi'$ are realized in $\mathfrak{A}$ at least $(mC+1)^2+1$ times, such sets can be found even if $\pi = \pi'$.) Let

$$B_0 = \{b \in B \mid \text{for some } b' \in B',\ b' \text{ sends a message to } b\}.$$

Since $\deg(\mathfrak{A}) = C$ and $|B'| = mC + 1$, we have $|B_0| \leq mC(mC + 1)$. So let $b \in B \setminus B_0$. But again, since $b$ can send a message to at most $mC$ elements of $B'$, there exists $b' \in B'$ such that $b$ does not send a message to $b'$. □

The greatest challenge when analysing the fragment $\mathcal{C}^2$ is the varied nature of the models of its formulas. It is therefore important to be able to confine attention to structures exhibiting certain characteristics which make them easier to manipulate. The notions of *chromaticity* and *differentiation* are particularly valuable in this regard.

**Definition 7.** Let $\mathfrak{A}$ be a structure interpreting a classified signature $\langle \Sigma, \bar{f} \rangle$. We say that $\mathfrak{A}$ is *chromatic* if distinct elements connected by a chain of 1 or 2 invertible message-types have distinct 1-types. That is, $\mathfrak{A}$ is chromatic just in case, for all $a, a', a'' \in A$:

1. if $a \neq a'$ and $\mathrm{tp}^{\mathfrak{A}}[a, a']$ is an invertible message-type, then $\mathrm{tp}^{\mathfrak{A}}[a] \neq \mathrm{tp}^{\mathfrak{A}}[a']$; and

2. if $a, a', a''$ are pairwise distinct and both $\mathrm{tp}^{\mathfrak{A}}[a, a']$ and $\mathrm{tp}^{\mathfrak{A}}[a', a'']$ are invertible message-types, then $\mathrm{tp}^{\mathfrak{A}}[a] \neq \mathrm{tp}^{\mathfrak{A}}[a'']$.

Chromatic structures are easy to work with because they exhibit the following properties:

**Remark 3.** *Let $\mathfrak{A}$ be a chromatic structure interpreting a classified signature $\langle \Sigma, \bar{f} \rangle$, and let $\pi'$ be a 1-type over $\Sigma$. Let $a$ be an element of $A$. Then there is at most one element $a' \in A \setminus \{a\}$ with 1-type $\pi'$ such that $a$ sends an invertible message to $a'$. Furthermore, if $\mathrm{tp}^{\mathfrak{A}}[a] = \pi'$, then there is no such element $a'$. (Fig. 4.2.)*

Structures of finite degree interpreting a classified signature $\langle \Sigma, \bar{f} \rangle$ can be rendered chromatic by means of a modest expansion:

**Lemma 16.** *Let $\langle \Sigma, \bar{f} \rangle$ be a classified signature with $m$ featured predicates, and let $\mathfrak{A}$ be a structure of finite degree interpreting $\Sigma$. Then $\mathfrak{A}$ can be expanded to a chromatic structure $\mathfrak{A}'$ by interpreting $\lceil \log((mC)^2 + 1) \rceil$ new unary predicates, where $C = \deg(\mathfrak{A})$.*

*Proof.* Consider the (undirected) graph $G$ on $A$ whose edges are the pairs of distinct elements connected by a chain of 1 or 2 invertible message-types. That is, $G = (A, E^1 \cup E^2)$, where

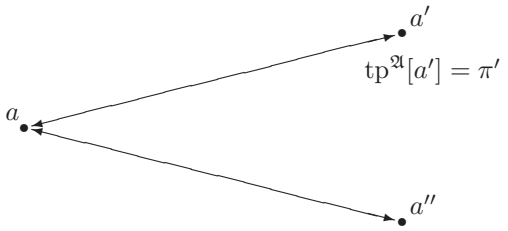


Figure 4.2: Invertible messages sent by $a$ in a chromatic structure $\mathfrak{A}$: neither $a'$ nor $a''$ can have the same 1-type as $a'$.

$$E^1 = \{(a, a') \mid a \neq a' \text{ and } \mathrm{tp}^{\mathfrak{A}}[a, a'] \text{ is an invertible message-type}\}$$
$$E^2 = \{(a, a'') \mid a \neq a'' \text{ and for some } a' \in A,\ (a, a') \text{ and } (a', a'') \text{ are both in } E^1\}.$$

Since $\mathfrak{A}$ has degree $C$, the degree of $G$ (in the normal graph-theoretic sense) is at most $(mC)^2$. Now use the standard (greedy) algorithm to colour the nodes of $G$ with $(mC)^2 + 1$ colours in such a way that no edge joins two nodes of the same colour. By interpreting the $\lceil \log((mC)^2 + 1) \rceil$ new unary predicates to encode these colours, we obtain the desired chromatic expansion of $\mathfrak{A}$. □

We turn now to the concept of differentiation:

**Definition 8.** Let $\mathfrak{A}$ be a structure interpreting a signature $\Sigma$, and let $K$ be a positive integer. We say that $\mathfrak{A}$ is *$K$-differentiated* if, for every 1-type $\pi$ over $\Sigma$, the number $u$ of elements in $A$ having 1-type $\pi$ satisfies either $u \leq 1$ or $u > K$.

Thus, in a $K$-differentiated structure, every 1-type is realized either at most once or more than $K$ times. Any structure can be made differentiated by means of a modest expansion; moreover, such an expansion preserves chromaticity:

**Lemma 17.** *Let $\mathfrak{A}$ be a chromatic structure of finite degree interpreting a classified signature $\langle \Sigma, \bar{f} \rangle$, and let $K$ be a positive integer. Then, by interpreting $\lceil \log K \rceil$ new unary predicates, $\mathfrak{A}$ can be expanded to a chromatic, $K$-differentiated structure $\mathfrak{A}'$ with $\deg(\mathfrak{A}) = \deg(\mathfrak{A}')$.*

*Proof.* For each 1-type $\pi$ realized more than once but no more than $K$ times, colour the elements having 1-type $\pi$ using $K$ different colours. For each 1-type $\pi'$ realized either once or more than $K$ times, colour the elements having 1-type $\pi'$ using a single colour. At most $K$ colours are required for this process. By interpreting the $\lceil \log K \rceil$ new unary predicates to encode these colours, we obtain the desired expansion $\mathfrak{A}'$. This process obviously preserves chromaticity and degree. □

## 4.3 Structures interpreting classified signatures

In this section, we fix an arbitrary classified signature $\langle \Sigma, (f_1, \ldots, f_m) \rangle$. Let $s$ be the total number of predicates in $\Sigma$. As in Chapter 3, we assume a *standard*

| | |
|---|---|
| $\langle \Sigma, \bar{f} \rangle$ | a classified signature |
| $s$ | the number of symbols in $\Sigma$ |
| $L$ | the number of 1-types over $\Sigma$ |
| $M^*$ | the number of invertible message-types over $\Sigma$ |
| $M$ | the number of message-types over $\Sigma$ |
| $\pi_1, \ldots, \pi_L$ | the 1-types over $\Sigma$ |
| $\mu_1, \ldots, \mu_{M^*}$ | the invertible message-types over $\Sigma$ |
| $\mu_{M^*+1}, \ldots, \mu_M$ | the non-invertible message-types over $\Sigma$ |

Table 4.1: Quick reference guide to symbols.

*enumeration* $\pi_1, \ldots, \pi_L$ of the 1-types over $\Sigma$, where $L = 2^s$. We may likewise assume some standard enumeration $\mu_1, \ldots, \mu_M$ of the *message-types* over $\Sigma$. Again, the ordering of message-types in this enumeration is essentially arbitrary: for notational convenience, however, we shall insist that the *invertible* message-types always precede the *non-invertible* message-types in this enumeration. Thus, we assume that the message-types over $\Sigma$ are standardly enumerated as

$$\mu_1, \ldots, \mu_{M^*}, \mu_{M^*+1}, \ldots, \mu_M,$$

where $\mu_1, \ldots, \mu_{M^*}$ are the invertible message-types, and $\mu_{M^*+1}, \ldots, \mu_M$ the non-invertible message-types. The above notation, which will be used throughout this section, is summarized in Table 4.1.

**Remark 4.** *The number $M$ of message-types over $\Sigma$ satisfies $M \leq m2^{4s-1}$.*

Our first task is to acquire the means to talk about 'local configurations' in structures of finite degree interpreting $\Sigma$.

**Definition 9.** Let $\mathfrak{A}$ be a structure of finite degree interpreting $\Sigma$, and let $a$ be an element of $A$. The *profile* of $a$ in $\mathfrak{A}$, denoted $\mathrm{pr}^{\mathfrak{A}}[a]$, is the $M$-tuple $\bar{v} = (v_1, \ldots, v_M)$ of natural numbers where, for all $j$ $(1 \leq j \leq M)$,

$$v_j = |\{b \in A \setminus \{a\} : \mathrm{tp}^{\mathfrak{A}}[a, b] = \mu_j\}|.$$

The tuple $\mathrm{pr}^{\mathfrak{A}}[a]$ records, for each message-type $\mu_j$ $(1 \leq j \leq M)$, how many elements $a$ sends a message of type $\mu_j$ to. Since $\mathfrak{A}$ is assumed to be of finite degree, this number must be finite. For the next definition, recall that (Notation 1), if $\tau$ is a 2-type (with variables $x$, $y$), then $\mathrm{tp}_1(\tau)$ is the 1-type consisting of those literals in $\tau$ not featuring $y$.

**Definition 10.** A *star-type over* $\Sigma$ is a pair $\sigma = \langle \pi, \bar{v} \rangle$, where $\pi$ is a 1-type over $\Sigma$ and $\bar{v} = (v_1, \ldots, v_M)$ is an $M$-tuple of natural numbers satisfying the condition that, for all $j$ $(1 \leq j \leq M)$,

$$v_j > 0 \text{ implies } \mathrm{tp}_1(\mu_j) = \pi.$$

If $\mathfrak{A}$ is a structure of finite degree interpreting $\Sigma$, and $a \in A$, then $\langle \mathrm{tp}^{\mathfrak{A}}[a], \mathrm{pr}^{\mathfrak{A}}[a] \rangle$ is evidently a star-type, which we call *the star-type of $a$ in $\mathfrak{A}$*, denoted $\mathrm{st}^{\mathfrak{A}}[a]$. We say that the star-type $\sigma$ is *realized* in $\mathfrak{A}$ if $\sigma = \mathrm{st}^{\mathfrak{A}}[a]$ for some $a \in A$.

It helps to think of $\mathrm{st}^{\mathfrak{A}}[a]$ as a description of $a$'s 'local environment' in $\mathfrak{A}$. Thus, a star-type $\sigma = \langle \pi, \bar{v} \rangle$ is a description of a possible such local environment. Importantly, certain features of structures are expressed as features of these local environments.

**Definition 11.** Let $\sigma = \langle \pi, (v_1, \ldots, v_M) \rangle$ be a star-type over $\Sigma$. We call $\sigma$ *$C$-bounded* if, for all $h$ $(1 \leq h \leq m)$,

$$\sum \{v_j \mid 1 \leq j \leq M \text{ and } f_h(x, y) \in \mu_j\} \leq C.$$

Furthermore, we call $\sigma$ *chromatic* if, for every 1-type $\pi'$, the sum

$$c = \sum \{v_j \mid 1 \leq j \leq M^* \text{ and } \mathrm{tp}_2(\mu_j) = \pi'\}$$

satisfies $c \leq 1$, and satisfies $c = 0$ if $\pi' = \pi$.

**Remark 5.** *Let $\mathfrak{A}$ be a structure of finite degree interpreting $\Sigma$ and $C$ a natural number. Then $\deg(\mathfrak{A}) \leq C$ if and only if every star-type realized in $\mathfrak{A}$ is $C$-bounded. Furthermore, $\mathfrak{A}$ is chromatic if and only if every star-type realized in $\mathfrak{A}$ is chromatic.*

The concept of $C$-bounded star-types is important, because (having fixed $\langle \Sigma, \bar{f} \rangle$) there are only finitely many of them. For the next remark, recall that $m$ is the number of featured predicates, $L$ the number of 1-types, and $M$ the number of message-types.

**Notation 2.** If $\sigma = \langle \pi, (v_1, \ldots, v_M) \rangle$ is a star-type, we denote the number $v_j$ $(1 \leq j \leq M)$ by $\sigma[j]$.

**Remark 6.** *Let $\sigma$ be a $C$-bounded star-type over $\Sigma$. Then, for all $j$ $(1 \leq j \leq M)$, $\sigma[j] \leq C$, and, in fact, $\sum_{1 \leq j \leq M} \sigma[j] \leq mC$. Further, if $C$ is a natural number, the number of $C$-bounded star-types over $\Sigma$ is bounded above by $L(C+1)^M$.*

Since $M$ can be as high as $m2^{4s-1}$, where $s$ is the number of symbols in $\Sigma$, the expression $L(C+1)^M$ represents a doubly-exponential function of $s$, which as we shall see, is too rapidly-growing for the complexity bounds we hope to achieve. Obtaining a tighter bound on the the number of different star-types realized in models of $\mathcal{C}^2$-formulas is thus one of the key issues to be addressed in this chapter; however, for the present, we put this issue to one side. We mention one further, obvious property of chromatic star-types.

**Remark 7.** *Let $\sigma$ be a chromatic star-type over $\Sigma$, and let $j$ and $j'$ be integers between 1 and $M^*$ (so that $\mu_j$ and $\mu_{j'}$ are invertible message-types). If $\mu_j^{-1} = \mu_{j'}$, then either $\sigma[j] = 0$ or $\sigma[j'] = 0$. In particular, if $\mu_j^{-1} = \mu_j$, then $\sigma[j] = 0$.*

Having dealt with characterizations of local configurations in structures interpreting $\Sigma$, we now turn to characterizations of entire structures. We begin with an auxiliary notion.

**Definition 12.** Let $\mathfrak{A}$ be a structure interpreting $\Sigma$, and let $\pi$, $\pi'$ be 1-types (not necessarily distinct). We say that $\pi$ and $\pi'$ form a *quiet pair* in $\mathfrak{A}$ if there exist distinct $a, a' \in A$ such that $\mathrm{tp}^{\mathfrak{A}}[a] = \pi$, $\mathrm{tp}^{\mathfrak{A}}[a'] = \pi'$, and $\mathrm{tp}^{\mathfrak{A}}[a, a']$ is a silent 2-type.

Informally, $\pi$ and $\pi'$ form a quiet pair just in case some element with 1-type $\pi$ neither sends a message to, nor receives a message from, some element (itself excepted) with 1-type $\pi'$.

**Definition 13.** Let $\Xi$ denote the set of silent 2-types over $\Sigma$, and let $\mathcal{I}$ be the set of unordered pairs of (not necessarily distinct) integers between 1 and $L$: that is, $\mathcal{I} = \{\{i, i'\} \mid 1 \leq i \leq i' \leq L\}$. A *frame* (over $\Sigma$) is a tuple $\mathcal{F} = (\bar{\sigma}, \mathcal{Q}, \theta)$, satisfying:

1. $\bar{\sigma} = (\sigma_1, \ldots, \sigma_N)$ is an $N$-tuple of pairwise distinct star-types for some $N > 0$;
2. $\mathcal{Q} \subseteq \mathcal{I}$; and
3. $\theta : \mathcal{Q} \to \Xi$ is a function such that, for all $\{i, i'\} \in \mathcal{Q}$ with $i \leq i'$, $\mathrm{tp}_1(\theta(\{i, i'\})) = \pi_i$ and $\mathrm{tp}_2(\theta(\{i, i'\})) = \pi_{i'}$.

The *dimension* of $\mathcal{F}$ is $N$.

Think of a frame $\mathcal{F} = (\bar{\sigma}, \mathcal{Q}, \theta)$ as a (putative) schematic description of a structure, where $\bar{\sigma}$ tells us which star-types are realized, $\mathcal{Q}$ tells us which pairs of 1-types are quiet, and $\theta$ selects, for each quiet pair of 1-types, a silent 2-type joining them. More precisely:

**Definition 14.** Let $\mathfrak{A}$ be a structure interpreting $\Sigma$, and let $\mathcal{F} = (\bar{\sigma}, \mathcal{Q}, \theta)$ be a frame over $\Sigma$. We say that $\mathcal{F}$ *describes* $\mathfrak{A}$ if the following conditions hold:

1. $\bar{\sigma}$ is a list of all and only those star-types realized in $\mathfrak{A}$;
2. if $\pi_i$ and $\pi_{i'}$ form a quiet pair in $\mathfrak{A}$, then $\{i, i'\} \in \mathcal{Q}$;
3. if $\pi_i$ and $\pi_{i'}$ form a quiet pair in $\mathfrak{A}$, then there exist distinct $a, a' \in A$ such that $\mathrm{tp}^{\mathfrak{A}}[a, a'] = \theta(\{i, i'\})$.

Any structure $\mathfrak{A}$ interpreting $\Sigma$ is evidently described by some (not necessarily unique) frame.

**Definition 15.** Let $\mathcal{F} = (\bar{\sigma}, \mathcal{Q}, \theta)$ be a frame over $\Sigma$. We call $\mathcal{F}$ *$C$-bounded* if every star-type in $\bar{\sigma}$ is $C$-bounded. Likewise, we call $\mathcal{F}$ *chromatic* if every star-type in $\bar{\sigma}$ is chromatic.

The following remark is just a re-statement of Remark 5.

**Remark 8.** *Let $\mathfrak{A}$ be a structure of finite degree interpreting $\Sigma$, and let $\mathcal{F}$ be a frame describing $\mathfrak{A}$. Then* $\deg(\mathfrak{A}) \leq C$ *if and only if $\mathcal{F}$ is $C$-bounded. Furthermore, $\mathfrak{A}$ is chromatic if and only if $\mathcal{F}$ is chromatic.*

Thus, certain interesting properties of $\mathfrak{A}$ correspond to properties of the frames which describe it.

However, while every structure interpreting $\Sigma$ is described by some frame, not every frame over $\Sigma$ describes a structure; and it is important for us to define a class of frames which do. To this end, we associate with a frame $\mathcal{F}$ a collection of numerical parameters, as follows.

**Notation 3.** Let $\mathcal{F} = (\bar{\sigma}, \mathcal{Q}, \theta)$ be a frame over $\Sigma$, where $\bar{\sigma} = (\sigma_1, \ldots, \sigma_N)$, for some $N > 0$, and recall the notation established in Table 4.1. If $\mathcal{F}$ is clear from context, for integers $i, k$ in the ranges $1 \leq i \leq L$, $1 \leq k \leq N$ write:

$$o_{ik} = \begin{cases} 1 \text{ if } \mathrm{tp}(\sigma_k) = \pi_i \\ 0 \text{ otherwise;} \end{cases}$$

$$p_{ik} = \begin{cases} 1 \text{ if, for all } j \ (1 \leq j \leq M), \ \mathrm{tp}_2(\mu_j) = \pi_i \text{ implies } \sigma_k[j] = 0 \\ 0 \text{ otherwise;} \end{cases}$$

$$r_{ik} = \sum_{j \in J} \sigma_k[j], \text{ where } J = \{j \mid M^* + 1 \leq j \leq M \text{ and } \mathrm{tp}_2(\mu_j) = \pi_i\};$$

$$s_{ik} = \sum_{j \in J} \sigma_k[j], \text{ where } J = \{j \mid 1 \leq j \leq M \text{ and } \mathrm{tp}_2(\mu_j) = \pi_i\}.$$

In addition, for integers $i, j$ in the ranges $1 \leq i \leq L$, $1 \leq j \leq M^*$, write:

$$q_{jk} = \sigma_k[j].$$

These parameters have the following intuitive interpretations.

**Remark 9.** *Let $\mathfrak{A}$ be a structure of finite degree interpreting $\Sigma$, and let $\mathcal{F}$ be a frame describing $\mathfrak{A}$. Then the symbols $o_{ik}$, $p_{ik}$, $q_{jk}$, $r_{ik}$ and $s_{ik}$ in Notation 3 have the following interpretations with respect to $\mathfrak{A}$:*

1. *$o_{ik} = 1$ just in case every element with star-type $\sigma_k$ has 1-type $\pi_i$;*
2. *$p_{ik} = 1$ just in case no element with star-type $\sigma_k$ sends a message to any element having 1-type $\pi_i$;*
3. *$q_{jk}$ counts how many messages of (invertible) type $\mu_j$ any element having star-type $\sigma_k$ sends;*
4. *$r_{ik}$ is the total number of elements having 1-type $\pi_i$ to which any element having star-type $\sigma_k$ sends a non-invertible message; and*
5. *$s_{ik}$ is the total number of elements having 1-type $\pi_i$ to which any element having star-type $\sigma_k$ sends a message.*

With this notation in hand we can characterize a class of frames whose members are guaranteed to describe *finite* structures over $\Sigma$.

**Definition 16.** Let $\mathcal{F} = (\bar{\sigma}, \mathcal{Q}, \theta)$ be a frame over $\Sigma$, where $\bar{\sigma} = (\sigma_1, \ldots, \sigma_N)$. Let $\bar{w} = (w_1, \ldots, w_N)$ be a tuple of positive integers. Using Notation 3, for all $i$ $(1 \leq i \leq L)$, all $i'$ $(1 \leq i' \leq L)$ and all $j$ $(1 \leq j \leq M^*)$, let:

$$u_i = \sum_{1 \leq k \leq N} o_{ik} w_k \qquad v_j = \sum_{1 \leq k \leq N} q_{jk} w_k \qquad x_{ii'} = \sum_{1 \leq k \leq N} o_{ik} p_{i'k} w_k.$$

Let $Z$ be a positive integer. We say that $\bar{w}$ is a *finite $Z$-solution* of $\mathcal{F}$ if the following conditions are satisfied for all $i$ $(1 \leq i \leq L)$, all $i'$ $(1 \leq i' \leq L)$, all $j$ $(1 \leq j \leq M^*)$ and all $k$ $(1 \leq k \leq N)$:

**(C1)** $v_j = v_{j'}$, where $j'$ is such that $\mu_j^{-1} = \mu_{j'}$;

**(C2)** $s_{ik} \leq u_i$;

**(C3)** $u_i \leq 1$ or $u_i > Z$;

**(C4)** if $o_{ik} = 1$, then either $u_i > 1$ or $r_{i'k} \leq x_{i'i}$;

**(C5)** if $\{i, i'\} \notin \mathcal{Q}$, then either $u_i \leq 1$ or $u_{i'} \leq 1$;

**(C6)** if $\{i, i'\} \notin \mathcal{Q}$ and $o_{ik} = 1$, then $r_{i'k} \geq x_{i'i}$.

If a finite $Z$-solution exists, $\mathcal{F}$ is said to be *finitely $Z$-solvable*.

The conditions **C1**–**C6** in Definition 16 form a system $\mathcal{E}$ of linear equations and inequalities with integer coefficients; and $\mathcal{F}$ is finitely $Z$-solvable if and only if $\mathcal{E}$ has an a solution over the natural numbers. Moreover, if $\mathcal{F}$, $Z$ and $\mathcal{E}$ are encoded in the obvious way (in particular, with integers represented as binary strings), we see that $\|\mathcal{E}\|$ is bounded by a polynomial function of $\|\mathcal{F}\| + \lfloor \log Z \rfloor$.

**Lemma 18.** *The problem of determining whether a frame $\mathcal{F}$ over a classified signature is finitely $Z$-solvable is in NP, given standard encoding of the inputs $\mathcal{F}$ and $Z$.*

*Proof.* Immediate from Corollary 3 □

**Remark 10.** *Let $\mathfrak{A}$ be a structure of finite degree interpreting $\Sigma$, and let $\mathcal{F} = (\bar{\sigma}, \mathcal{Q}, \theta)$ be a frame describing $\mathfrak{A}$. For all $k$ $(1 \leq k \leq N)$, let $w_k$ be the number of elements of $A$ having star-type $\sigma_k$ in $\mathfrak{A}$. In that case, the symbols $u_i$, $v_j$ and $x_{ii'}$ in Definition 16 have the following interpretations with respect to $\mathfrak{A}$:*

1. *$u_i$ is the number of elements $a \in A$ such that $\mathrm{tp}^{\mathfrak{A}}[a] = \pi_i$;*
2. *$v_j$ is the number of pairs $\langle a, b \rangle \in A^2$ such that $a \neq b$ and $\mathrm{tp}^{\mathfrak{A}}[a, b] = \mu_j$;*
3. *$x_{ii'}$ is the number of elements $a \in A$ such that $\mathrm{tp}^{\mathfrak{A}}[a] = \pi_i$ and $a$ does not send a message to any element having 1-type $\pi_{i'}$.*

Recall the notion of differentiation introduced in Definition 8.

**Lemma 19.** *Suppose $\mathfrak{A}$ is a finite structure interpreting $\Sigma$ (hence, of finite degree), let $C = \deg(\mathfrak{A})$, let $K \geq (mC+1)^2$, and let $\mathcal{F} = (\bar{\sigma}, \mathcal{Q}, \theta)$ be a frame describing $\mathfrak{A}$. If $\mathfrak{A}$ is $K$-differentiated, then $\mathcal{F}$ has a finite $K$-solution.*

*Proof.* Let $\bar{\sigma} = (\sigma_1, \ldots, \sigma_N)$, and let $w_k = |\{a \in A : \mathrm{st}^{\mathfrak{A}}[a] = \sigma_k\}|$ for all $k$ $(1 \leq k \leq N)$. We show that $\bar{w} = (w_1, \ldots, w_N)$ is a solution of $\mathcal{F}$. In doing so, we make free use of Remarks 9 and 10. Note that, by construction, the $w_1, \ldots, w_N$ are all positive.

**C1**: If $\mu_j^{-1} = \mu_{j'}$, then the sets $\{\langle a, b\rangle \mid a \neq b \text{ and } \mathrm{tp}^{\mathfrak{A}}[a,b] = \mu_j\}$ and $\{\langle a, b\rangle \mid a \neq b \text{ and } \mathrm{tp}^{\mathfrak{A}}[a,b] = \mu_{j'}\}$ can obviously be put in 1–1 correspondence, namely: $\langle a, b\rangle \mapsto \langle b, a\rangle$. But the cardinalities of these sets are $v_j$ and $v_{j'}$, respectively.

**C2**: Since $\mathcal{F}$ describes $\mathfrak{A}$, any element of $A$ having star-type $\sigma_k$ sends a message to exactly $s_{ik}$ elements having 1-type $\pi_i$. But $u_i$ is the number of elements of $A$ having 1-type $\pi_i$. Since $\sigma_k$ is realized in $\mathfrak{A}$, $s_{ik} \leq u_i$.

**C3**: Immediate given that $\mathfrak{A}$ is $K$-differentiated.

**C4**: If $o_{ik} = 1$ and $u_i \leq 1$, then $u_i = 1$, so that $\mathfrak{A}$ contains exactly one element with 1-type $\pi_i$; moreover, this element has star-type $\sigma_k$. Denote this element by $a$. Thus, $a$ sends a non-invertible message to exactly $r_{i'k}$ elements with 1-type $\pi_{i'}$. Clearly, none of these elements sends a message back to $a$ (since otherwise $a$'s message to it would be invertible), so that there exist at least $r_{i'k}$ elements with 1-type $\pi_{i'}$ which do not send a message to $a$. But since $a$ is the only element with 1-type $\pi_i$, there exist at least $r_{i'k}$ elements with 1-type $\pi_{i'}$ which do not send a message to any element having 1-type $\pi_i$. In other words, $r_{i'k} \leq x_{i'i}$.

**C5**: Suppose $u_i > 1$ and $u_{i'} > 1$. Now, the structure $\mathfrak{A}$ is $K$-differentiated; hence $u_i > (mC+1)^2$, and $u_{i'} > (mC+1)^2$. But $\deg(\mathfrak{A}) = C$, so that Lemma 15 implies that $\pi_i$ and $\pi_{i'}$ form a quiet pair in $\mathfrak{A}$, and hence, since $\mathcal{F}$ describes $\mathfrak{A}$, that $\{i, i'\} \in \mathcal{Q}$.

**C6**: Since $\mathcal{F}$ describes $\mathfrak{A}$, $\{i, i'\} \notin \mathcal{Q}$ implies that $\pi_i$ and $\pi_{i'}$ do not form a quiet pair in $\mathfrak{A}$. Now if $o_{ik} = 1$, there exists at least one element $a$ having 1-type $\pi_i$ and star-type $\sigma_k$. Moreover, there are $x_{i'i}$ elements having 1-type $\pi_{i'}$ which do not send a message to any element having 1-type $\pi_i$, and hence at least $x_{i'i}$ elements having 1-type $\pi_{i'}$ which do not send a message to $a$. Therefore, $a$ sends a message (in fact, a non-invertible message) to all of these elements. But since $a$ has star-type $\sigma_k$, $a$ sends a non-invertible message to exactly $r_{i'k}$ elements having 1-type $\pi_{i'}$. Thus, $r_{i'k} \geq x_{i'i}$. □

We now prove a converse of Lemma 19, for the special case of chromatic frames.

**Lemma 20.** *Let $\mathcal{F}$ be a chromatic, $C$-bounded frame, and let $K \geq 3mC$. If $\mathcal{F}$ has a finite $K$-solution, then there exists a finite structure $\mathfrak{A}$ such that $\mathcal{F}$ describes $\mathfrak{A}$.*

*Proof.* Let $\mathcal{F} = (\bar{\sigma}, \mathcal{Q}, \theta)$, let $\bar{\sigma} = (\sigma_1, \ldots, \sigma_N)$, and let $\bar{w} = (w_1, \ldots, w_N)$ be a finite $K$-solution of $\mathcal{F}$. In the sequel, we use the symbols $o_{ik}$, $p_{ik}$, $q_{jk}$, $r_{ik}$ and $s_{ik}$ (with indices in the appropriate ranges), as specified in Notation 3, and the symbols $u_i$, $v_j$ and $x_{ii'}$ (again, with indices in the appropriate ranges), as specified in Definition 16. Hence, the conditions **C1**–**C6** of Definition 16 hold.

For every $k$ ($1 \leq k \leq N$), let $A_k$ be a set of cardinality $w_k$, and let $A$ be the disjoint union of the $A_k$. Think of $A_k$ as the set of elements which 'want' to have star-type $\sigma_k$. In addition, we define for all $i$ ($1 \leq i \leq L$), all $i'$ ($1 \leq i' \leq L$) and all $j$ ($1 \leq j \leq M^*$):

$$U_i = \bigcup \{A_k \mid 1 \leq k \leq N \text{ and } o_{ik} = 1\}$$
$$X_{ii'} = \bigcup \{A_k \mid 1 \leq k \leq N \text{ and } o_{ik}p_{i'k} = 1\}$$
$$V_j = \bigcup \{A_k \mid 1 \leq k \leq N \text{ and } q_{jk} = 1\}.$$

Since $\mathcal{F}$ is chromatic, $q_{jk} \leq 1$ for all $j$ ($1 \leq j \leq M^*$) and all $k$ ($1 \leq k \leq N$). Thus, for all $i$, $i'$ and $j$ in the appropriate ranges:

$$|U_i| = u_i; \qquad |X_{ii'}| = x_{ii'}; \qquad |V_j| = v_j.$$

Think of $U_i$ as the set of elements which 'want' to have 1-type $\pi_i$, $X_{ii'}$ as the set of elements in $U_i$ which do not 'want' to send a message to any element in $U_{i'}$, and $V_j$ as the set of elements which 'want' to send an (invertible) message of type $\mu_j$ to some other element. We remark that $A_k \subseteq U_i$ if and only if $\mathrm{tp}(\sigma_k) = \pi_i$. Moreover, for all $j$ ($1 \leq j \leq M^*$), if $V_j \neq \emptyset$, there exists a unique $i$ ($1 \leq i \leq L$) such that $V_j \subseteq U_i$—namely, that $i$ such that $\mathrm{tp}_1(\mu_j) = \pi_i$. We now convert the domain $A$ into a structure $\mathfrak{A}$ in four steps.

**Step 1** (Interpreting the unary predicates and diagonals of binary predicates): For every $k$ ($1 \leq k \leq N$) and every $a \in A_k$, set $\mathrm{tp}^{\mathfrak{A}}[a] = \mathrm{tp}(\sigma_k)$. At the end of this step, we have, for every $i$ ($1 \leq i \leq L$) and every $a \in U_i$, $\mathrm{tp}^{\mathfrak{A}}[a] = \pi_i$.

**Step 2** (Fixing the invertible message-types): For every $j$ ($1 \leq j \leq M^*$), let $j'$ be such that $\mu_j^{-1} = \mu_{j'}$. By **C1**, $V_j$ and $V_{j'}$ are equinumerous. If $j' > j$, pick some 1–1 correspondence between $V_j$ and $V_{j'}$; and for every $a \in V_j$, set $\mathrm{tp}^{\mathfrak{A}}[a, a'] = \mu_j$, where $a'$ is the element of $V_{j'}$ corresponding to $a \in V_j$. This completes Step 2. We must show that these assignments are meaningful, do not clash with Step 1, and do not clash with each other. Suppose then that the assignment $\mathrm{tp}^{\mathfrak{A}}[a, a'] = \mu_j$ is made, and that $\mu_j^{-1} = \mu_{j'}$. Thus, $a \in V_j$ and $a' \in V_{j'}$. To show that the assignment is meaningful, we must prove that $a \neq a'$. For contradiction, suppose $a = a'$, and let $k$ be such that $a \in A_k$. But then $\sigma_k[j] > 0$ and $\sigma_k[j'] > 0$, which is impossible by Remark 7. To show that the assignment does not clash with Step 1, suppose $\mu_j^{-1} = \mu_{j'}$, and let $i$, $i'$ be such that $V_j \subseteq U_i$ and $V_{j'} \subseteq U_{i'}$. As observed above, $\pi_i = \mathrm{tp}_1(\mu_j)$ and $\pi_{i'} = \mathrm{tp}_1(\mu_{j'}) = \mathrm{tp}_2(\mu_j)$, which conforms to the assignments in Step 1. To show that these assignments do not clash with each other, it suffices to prove that,

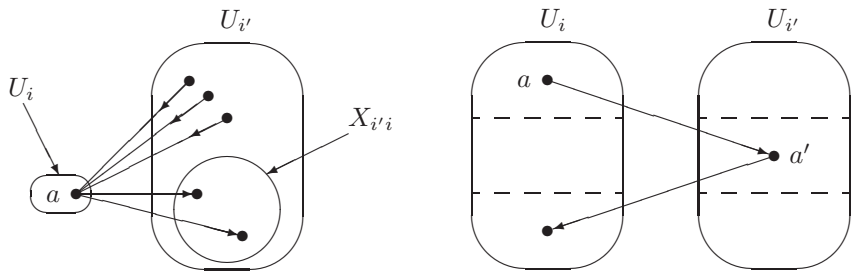


Case 2: $u_i = 1$ Case 5: $u_i, u_{i'} > K \geq 3mC$

Figure 4.3: Dealing with non-invertible messages between $U_i$ and $U_{i'}$.

if $a \in V_j \cap V_h$, $a' \in V_{j'} \cap V_{h'}$ $\mu_j^{-1} = \mu_{j'}$ and $\mu_h^{-1} = \mu_{h'}$, then $j = h$. Suppose then that the antecedent of this conditional holds; let $k$ and $k'$ be such that $a \in A_k$ and $a' \in A_{k'}$. Then $\sigma_{k'}[j'] > 0$ and $\sigma_{k'}[h'] > 0$. Since $\sigma_{k'}$ is a star-type, $\mathrm{tp}_1(\mu_{j'}) = \mathrm{tp}_1(\mu_{h'})$, whence $\mathrm{tp}_2(\mu_j) = \mathrm{tp}_2(\mu_h)$. But $\sigma_k[j] > 0$ and $\sigma_k[h] > 0$, and since $\sigma_k$ is a *chromatic* star-type, $j = h$. Note that, if $\mu_j^{-1} = \mu_j$, then $V_j = \emptyset$ by Remark 7. Thus, at the end of Step 2, for every element $a \in A$ and every $j$ $(1 \leq j \leq M^*)$, $a$ sends a (unique) message of type $\mu_j$ to some other element if and only if $a \in V_j$. That is: for all $k$ $(1 \leq k \leq N)$, all $a \in A_k$, and all $j$ $(1 \leq j \leq M^*)$, there are exactly $\sigma_k[j]$ elements $a' \in A$ such that $a \neq a'$ and $\mathrm{tp}^{\mathfrak{A}}[a, a'] = \mu_j$. We make one further observation before proceeding. Suppose that $\mathrm{tp}^{\mathfrak{A}}[a, a']$ is assigned in this step and that $a \in U_i$; we claim that $a' \notin X_{i'i}$ for any $i'$. To see this, suppose $a \in V_j \subseteq U_i$ and $a' \in A_{k'} \subseteq V_{j'}$, with $\mu_j^{-1} = \mu_{j'}$. Then $\mathrm{tp}_1(\mu_j) = \mathrm{tp}_2(\mu_{j'}) = \pi_i$. But then $\sigma_{k'}[j'] > 0$, whence $p_{ik'} = 0$, whence $a' \notin X_{i'i}$. This observation will be useful in Step 3.

**Step 3** (Fixing the non-invertible message-types): Let $i$ and $i'$ be such that $1 \leq i \leq i' \leq L$. We fix all the non-invertible messages sent, in either direction, between $U_i$ and $U_{i'}$. By **C3**, either $u_i \leq 1$ or $u_i > K$; similarly, either $u_{i'} \leq 1$ or $u_{i'} > K$. We consider five cases.

*Case 1:* $u_i = 0$. In this case, there are no elements of $U_i$ and hence no 2-type assignments to be made between elements of $U_i$ and elements of $U_{i'}$. Note that, by **C2**, $s_{ik} = 0$ for all $k$ $(1 \leq k \leq N)$, whence $\sigma_k[j] = 0$ for all $k$ $(1 \leq k \leq N)$ and for all $j$ $(1 \leq j \leq M)$ such that $\mathrm{tp}_2(\mu_j) = \pi_i$. (Intuitively, no element of $A$—and in particular of $U_{i'}$—'wants' to send a message to an element with 1-type $\pi_i$ anyway.)

*Case 2:* $u_i = 1$. The situation is illustrated in the left-hand diagram of Fig. 4.3. Let $a$ be the sole element of $U_i$, and let $k$ be such that $a \in A_k$. We deal first with

the assignment of non-invertible messages sent from $U_{i'}$ to $U_i = \{a\}$. Consider any $a' \in A_{k'} \subseteq U_{i'}$. By **C2**, $s_{ik'} \leq 1$; hence there is at most one value of $j$ in the range $1 \leq j \leq M$ such that $\mathrm{tp}_2(\mu_j) = \pi_i$ and $\sigma_{k'}[j] > 0$. Suppose then that such a $j$ exists. Again, since $s_{ik'} \leq 1$, $\sigma_{k'}[j] = 1$. If $j \leq M^*$, then this message has already been dealt with in Step 2; so we may assume $M^*+1 \leq j \leq M$. It follows from **C4** that $i \neq i'$. Hence $a \neq a'$, so that we may set $\mathrm{tp}^{\mathfrak{A}}[a', a] = \mu_j$. Since $\mathrm{tp}_1(\mu_j) = \pi_{i'}$ and $\mathrm{tp}_2(\mu_j) = \pi_i$, this assignment does not clash with Step 1. Observe also that, just as in Step 2, if this assignment is made, we have, by definition, $p_{ik'} = 0$, so that $a' \notin X_{i'i}$. By carrying out the same procedure for all $a' \in U_{i'}$, we complete the assignment of non-invertible messages sent from $U_{i'}$ to $U_i$. It remains to deal with the non-invertible messages sent from $U_i = \{a\}$ to $U_{i'}$. Remembering that $a \in A_k$, **C4** ensures the existence of a subset $R \subseteq X_{i'i}$ such that $|R| = r_{i'k}$. For each $j$ ($M^*+1 \leq j \leq M$), if $\mathrm{tp}_2(\mu_j) = \pi_{i'}$, select $\sigma_k[j]$ fresh elements $a'$ of $R$, and make the assignment $\mathrm{tp}^{\mathfrak{A}}[a, a'] = \mu_j$. (There are enough such elements by the definition of $r_{i'k}$.) These assignments clearly do not clash with those made in Step 1. Moreover, we have observed that $\mathrm{tp}^{\mathfrak{A}}[a, a']$ has previously been assigned (either in this step or in Step 2) only if $a' \notin X_{i'i}$. Thus, these assignments do not clash with those made earlier in this step or those made in Step 2.

*Case 3:* $u_{i'} = 0$ and $u_i > K$. Symmetrical to Case 1.

*Case 4:* $u_{i'} = 1$ and $u_i > K$. Symmetrical to Case 2.

*Case 5:* $u_i > K$ and $u_{i'} > K$. Since $K \geq 3mC$, partition $U_i$ into three sets $U_{i0}$, $U_{i1}$, $U_{i2}$, each containing at least $mC$ elements; and similarly for $U_{i'}$. Suppose $a \in U_i$. Then for some $h$ ($0 \leq h < 3$), $a \in U_{ih}$. Let $k$ be such that $a \in A_k$, and let $h' = h + 1 \pmod 3$. For all $j$ ($M^* + 1 \leq j \leq M$), select $\sigma_k[j]$ fresh elements $a'$ of $U_{i'h'}$ such that $\mathrm{tp}^{\mathfrak{A}}[a, a']$ was not assigned in Step 2, and set $\mathrm{tp}^{\mathfrak{A}}[a, a'] = \mu_j$. By Since $\sigma_k$ is $C$-bounded, by Remark 6, $\sum_{1 \leq j \leq M} \sigma_k[j] \leq mC$; and since $|U_{i'h'}| \geq mC$, we never run out of fresh elements to select. In this way, we deal with all messages sent from $U_i$ to $U_{i'}$; the messages sent from $U_{i'}$ to $U_i$ are dealt with symmetrically. It is obvious that these assignments do not clash with Step 1 or Step 2; and the fact that $h' = h + 1 \pmod 3$, ensures that they do not clash with each other (even if $i = i'$), as is evident from the right-hand diagram of Fig. 4.3.

Performing these assignments for all pairs $i, i'$ such that $1 \leq i \leq i' \leq L$ completes Step 3. At the end of Step 3, then, for all $k$ ($1 \leq k \leq N$), all $a \in A_k$, and all $j$ ($1 \leq j \leq M$), there are exactly $\sigma_k[j]$ elements $a' \in A$ such that $a \neq a'$ and $\mathrm{tp}^{\mathfrak{A}}[a, a'] = \mu_j$.

**Step 4** (Fixing the silent 2-types): Finally, we use the components $\mathcal{Q}$ and $\theta$ of $\mathcal{F} = (\bar{\sigma}, \mathcal{Q}, \theta)$ to deal with all the remaining 2-types in $\mathfrak{A}$. Let $a, a'$ be distinct elements of $A$ such that $\mathrm{tp}^{\mathfrak{A}}[a, a']$ has not yet been assigned. Let $i, i', k, k'$ be such that $a \in A_k \subseteq U_i$ and $a' \in A_{k'} \subseteq U_{i'}$, and assume, without loss of generality, that $i \leq i'$. We claim that $\{i, i'\} \in \mathcal{Q}$. For suppose otherwise. By **C5**, we have either $u_i = 1$ or $u_{i'} = 1$. Assume the former. If $p_{ik'} = 0$,

then there is some $j'$ ($1 \leq j' \leq M$) such that $\sigma_{k'}[j'] > 0$ and $\mathrm{tp}_2(\mu_{j'}) = \pi_i$, whence—bearing in mind that $a$ is the unique element of $U_i$—$\mathrm{tp}^{\mathfrak{A}}[a, a']$ will certainly have been assigned in Step 2 (if $\mu_{j'}$ is an invertible message-type) or in Step 3 Case 2 (if $\mu_{j'}$ is a non-invertible message-type), contradicting the fact that $\mathrm{tp}^{\mathfrak{A}}[a, a']$ is unassigned. Thus, $p_{ik'} = 1$, and hence $o_{i'k'}p_{ik'} = 1$. That is: $a' \in X_{i'i}$. But $|X_{i'i}| = x_{i'i}$. And by **C6**, $x_{i'i} \leq r_{i'k}$. Yet in Step 3 (Case 2), $r_{i'k}$ elements of $X_{i'i}$ were chosen to receive messages from $a$. Hence $a'$ must be among these elements, again contradicting the fact that $\mathrm{tp}^{\mathfrak{A}}[a, a']$ is unassigned. The case where $u_{i'} \leq 1$ proceeds symmetrically. Thus, we have established that, if $\mathrm{tp}^{\mathfrak{A}}[a, a']$ has not yet been assigned, then $\{i, i'\} \in \mathcal{Q}$, so that we can make the assignment $\mathrm{tp}^{\mathfrak{A}}[a, a'] = \theta(\{i, i'\})$. Since $\mathrm{tp}_1(\theta(\{i, i'\})) = \pi_i$ and $\mathrm{tp}_2(\theta(\{i, i'\})) = \pi_{i'}$, there is no clash with Step 1. Evidently, we can proceed in this way until all remaining 2-types have been assigned. Moreover, since each $\theta(\{i, i'\})$ is silent, this step does not spoil the work of Steps 2–3: we still have that, for all $k$ ($1 \leq k \leq N$), all $a \in A_k$, and all $j$ ($1 \leq j \leq M$), there are exactly $\sigma_k[j]$ elements $a' \in A$ such that $a \neq a'$ and $\mathrm{tp}^{\mathfrak{A}}[a, a'] = \mu_j$.

This completes the construction of $\mathfrak{A}$. It is easy to see that $\mathcal{F}$ describes $\mathfrak{A}$. □

We need to generalize the above results to allow us to deal with infinite structures. In fact, there is little further work to do.

**Notation 4.** Let $\mathbb{N}^*$ denote the set $\mathbb{N} \cup \{\aleph_0\}$. We extend the ordering $>$ and the arithmetic operations $+$ and $\cdot$ from $\mathbb{N}$ to $\mathbb{N}^*$ in the obvious way. Specifically, we define $\aleph_0 > n$ for all $n \in \mathbb{N}$; we define $\aleph_0 + \aleph_0 = \aleph_0 \cdot \aleph_0 = \aleph_0$ and $0 \cdot \aleph_0 = \aleph_0 \cdot 0 = 0$; we define $n + \aleph_0 = \aleph_0 + n = \aleph_0$ for all $n \in \mathbb{N}$; and we define $n \cdot \aleph_0 = \aleph_0 \cdot n = \aleph_0$ for all $n \in \mathbb{N}$ such that $n > 0$. Under this extension, $>$ remains a total order, and $+$, $\cdot$ remain associative and commutative.

A system of linear equalities and inequalities defining an integer programming problem can of course be re-interpreted so that solutions are sought not over $\mathbb{N}$ but over $\mathbb{N}^*$. (We always assume that the coefficients occurring in such problems are in $\mathbb{N}$.) As an example, the single inequality $x_1 \geq x_1 + 1$ has no solutions over $\mathbb{N}$, but it does have a solution over $\mathbb{N}^*$, namely, $x_1 = \aleph_0$.

**Lemma 21.** *Let $\mathcal{E}$ be a finite set of linear inequalities of the form*

$$a_0 + a_1x_1 + \cdots + a_nx_n \leq b_0 + b_1x_1 + \cdots + b_nx_n$$

*in variables $x_1, \ldots, x_n$. Here, all coefficients are assumed to be in $\mathbb{N}$. If $\mathcal{E}$ has a solution over $\mathbb{N}^*$, then $\mathcal{E}$ has a solution over $\mathbb{N}^*$ such that all finite values are bounded by some fixed exponential function of $\|\mathcal{E}\|$.*

*Proof.* Suppose that $\mathcal{E}$ has a solution over $\mathbb{N}^*$. Re-order the variables if necessary so that this solution has the form $\bar{w}\bar{\aleph}_0$, with $\bar{w} = w_1, \ldots, w_k \in \mathbb{N}^k$ for some $k$ ($0 \leq k \leq n$) and $\bar{\aleph}_0$ an $(n-k)$-tuple of $\aleph_0$s. Say that an inequality in $\mathcal{E}$ *does not involve* the variable $x_i$ if the corresponding coefficients $a_i$ and $b_i$ are both zero. Let $\mathcal{E}'$ be the set of inequalities in $\mathcal{E}$ involving none of the $x_{k+1}, \ldots, x_n$.

Thus, $\mathcal{E}'$, considered as a problem in variables $x_1, \ldots, x_k$, has a solution $\bar{w}$ over $\mathbb{N}$, whence, by Theorem 4, it has a solution $\bar{w}'$ with values bounded by some fixed exponential function of $\|\mathcal{E}'\|$ (and hence of $\|\mathcal{E}\|$). But it is easy to see that $\bar{w}'\bar{\aleph}_0$ is a solution of $\mathcal{E}$. □

Lemmas 15, 16 and 17 apply to both finite and infinite structures. Definition 16 requires modification, however.

**Definition 17.** Let $\mathcal{F} = (\bar{\sigma}, \mathcal{Q}, \theta)$ and $Z$ be as in Definition 16. Let $\bar{w} = (w_1, \ldots, w_N)$ be a tuple of non-zero elements of $\mathbb{N}^*$. We say that $\bar{w}$ is a *$Z$-solution* of $\mathcal{F}$ if $\bar{w}$ satisfies the conditions of Definition 16, but with the arithmetic interpreted over $\mathbb{N}^*$ as specified in Notation 4.

We then note that the reasoning of Lemmas 19 and 20 works unproblematically for countably infinite models.

**Lemma 22.** *Suppose $\mathfrak{A}$ is a structure of finite degree interpreting $\Sigma$, let $C = \deg(\mathfrak{A})$, let $K \geq (mC+1)^2$, and let $\mathcal{F} = (\bar{\sigma}, \mathcal{Q}, \theta)$ be a frame describing $\mathfrak{A}$. If $\mathfrak{A}$ is $K$-differentiated, then $\mathcal{F}$ has a $K$-solution.*

**Lemma 23.** *Let $\mathcal{F}$ be a chromatic, $C$-bounded frame, and let $K \geq 3mC$. If $\mathcal{F}$ has a $K$-solution, then there exists a structure $\mathfrak{A}$ of finite degree interpreting $\Sigma$ such that $\mathcal{F}$ describes $\mathfrak{A}$.*

The proofs are exactly the same as in the finite case. Note that the variables $u_i$, $v_j$ and $x_{ii'}$ as well as the $w_k$ may now take the value $\aleph_0$; by contrast, the coefficients $o_{ik}$, $p_{ik}$, $q_{jk}$, $r_{ik}$ and $s_{ik}$ remain finite. Remark 10 generalizes unproblematically to countably infinite structures, so that the quantities $u_i$, $v_j$ and $x_{ii'}$ (implicitly) mentioned in Definition 16 continue to have their familiar interpretations. The proofs of Lemmas 22 and 23 then proceed exactly as for Lemmas 19 and 20.

Finally, we note that Lemma 18 also generalizes to the infinite case.

**Lemma 24.** *The problem of determining whether a frame $\mathcal{F}$ over a classified signature is $Z$-solvable is in NP, given standard encoding of the inputs $\mathcal{F}$ and $Z$.*

*Proof.* Immediate from Lemma 21. □

Lemmas 19 and 20 put us in a position to characterize exactly those frames $\mathcal{F}$ which describe finite, chromatic, $(mC+1)^2$-differentiated structures over $\Sigma$. Likewise, Lemmas 22 and 23 put us in a position to characterize exactly those frames $\mathcal{F}$ which describe finite-degree (but not necessarily finite), chromatic, $(mC+1)^2$-differentiated structures over $\Sigma$. The significance of these characterizations lies in the facts that: (i) the number of $C$-bounded frames over $\Sigma$ is finite; and (ii) if a frame does describe a structure, then that frame alone contains enough information to determine the truth of a suitable Scott-form $\mathcal{C}^2$ formula in that structure, as we shall presently see.

For the next definition, recall that a 1-type $\pi$ is simply a finite collection of formulas, so that $\bigwedge \pi$ denotes the conjunction of those formulas; similarly for 2-types. Recall further that $f_1, \ldots, f_m$ are the featured predicates of $\Sigma$.

**Definition 18.** Let $\psi$ be any formula over the signature $\Sigma$ having the form

$$\psi = \forall x \forall y(\alpha \vee x \approx y) \wedge \bigwedge_{1 \leq h \leq m} \forall x \exists_{=C_h} y(f_h(x,y) \wedge x \not\approx y),$$

where $\alpha$ is a quantifier-free $\mathcal{L}^2$-formula, $m$ is a positive integer, and the $C_h$ are positive integers. Let $\mathcal{F} = (\bar{\sigma}, \mathcal{Q}, \theta)$ be a frame over $\Sigma$, where $\bar{\sigma} = (\sigma_1, \ldots, \sigma_N)$. We write $\mathcal{F} \models \psi$ if the following conditions are satisfied:

1. for all $k$ $(1 \leq k \leq N)$ and all $j$ $(1 \leq j \leq M)$, if $\sigma_k[j] > 0$ then $\models \bigwedge \mu_j \rightarrow \alpha(x,y) \wedge \alpha(y,x)$;
2. for all $\{i, i'\} \in \mathcal{Q}$, $\models \bigwedge \theta(\{i, i'\}) \rightarrow \alpha(x,y) \wedge \alpha(y,x)$;
3. for all $k$ $(1 \leq k \leq N)$ and all $h$ $(1 \leq h \leq m)$, the sum of all the $\sigma_k[j]$ $(1 \leq j \leq M)$ such that $f_h(x,y) \in \mu_j$ equals $C_h$.

**Lemma 25.** *Let $\mathfrak{A}$ be a structure of finite degree interpreting $\Sigma$, and let $\psi$ be a formula of the form given in Definition 18. If $\mathcal{F}$ is a frame describing $\mathfrak{A}$, and $\mathcal{F} \models \psi$, then $\mathfrak{A} \models \psi$. Conversely, if $\mathfrak{A} \models \psi$, then there exists a frame $\mathcal{F}$ describing $\mathfrak{A}$, such that $\mathcal{F} \models \psi$.*

*Proof.* The first statement of the lemma is almost immediate. For the second statement, suppose $\mathfrak{A} \models \psi$. Let $\bar{\sigma}$ be a list of exactly the star-types realized in $\mathfrak{A}$, and let $\mathcal{Q}$ be the set of *exactly* those sets $\{i, i'\}$ such that $\pi_i$ and $\pi_{i'}$ form a quiet pair in $\mathfrak{A}$. For all $\{i, i'\} \in \mathcal{Q}$ with $i \leq i'$, pick $a$, $a'$ with $\mathrm{tp}^{\mathfrak{A}}[a] = \pi_i$ and $\mathrm{tp}^{\mathfrak{A}}[a'] = \pi_{i'}$ such that $\mathrm{tp}^{\mathfrak{A}}[a, a']$ is a silent 2-type; and set $\theta(\{i, i'\}) = \mathrm{tp}^{\mathfrak{A}}[a, a']$. Then $\mathcal{F} = \langle \bar{\sigma}, \mathcal{Q}, \theta \rangle$ describes $\mathfrak{A}$. (Note that $\mathcal{Q}$ contains only those pairs $\{i, i'\}$ which it must contain in order for $\mathcal{F}$ to describe $\mathfrak{A}$.) It is then easy to see that $\mathcal{F} \models \psi$. □

We now have all the essential ingredients for proving that Fin-Sat-$\mathcal{C}^2$ and Sat-$\mathcal{C}^2$ are decidable. By Lemma 14 we can restrict attention to formulas $\psi$ in Scott normal form. Consider the case of Fin-Sat-$\mathcal{C}^2$. By the discussion of Section 4.2 we can restrict attention to finding a finite, chromatic $(mC+1)^2$-differentiated model of $\psi$. And by Lemmas 19 and 20, we can reduce the task of determining the existence of such a model to that of finding a chromatic, $C$-bounded, finitely $(mC+1)^2$-solvable frame $\mathcal{F}$ such that $\mathcal{F} \models \psi$. But we may enumerate all $C$-bounded frames $\mathcal{F}$, and check for finite $(mC+1)^2$-solvability by Lemma 18. The case of Sat-$\mathcal{C}^2$ proceeds analogously.

The above argument even provides us with an upper complexity bound: if the size of the input formula is $n$, then we may work with a classified signature $\Sigma$ whose size is bounded by a polynomial function of $n$. Moreover, we may confine attention to $C$-bounded frames, where $C$ is bounded by an exponential function of $n$. The size of any $C$-bounded frame $\mathcal{F}$ over $\Sigma$ is therefore bounded by a doubly exponential function of $n$. Such a frame may therefore be guessed and checked for (finite) $(mC+1)^2$-solvability in non-deterministic doubly exponential time. Thus, Fin-Sat-$\mathcal{C}^2$ and Sat-$\mathcal{C}^2$ are both in 2-NEXPTIME. In the next section, we see how to reduce this upper bound to NEXPTIME.

## 4.4 Approximating structures

Throughout this section, we assume the classified signature $\langle \Sigma, (f_1, \ldots, f_m) \rangle$. Thus, "structure" means "structure interpreting $\Sigma$", "1-type" means "1-type over $\Sigma$, and so on. We also continue to employ the symbols $s$, $L$, $M^*$, $M$, $\pi_1, \ldots, \pi_L$ and $\mu_1, \ldots, \mu_M$ as summarized in Table 4.1. The first step is to manufacture some tools for manipulating structures.

**Notation 5.** Let $\mathfrak{A}$ be a structure, $\pi$ a 1-type, and $\Pi$ a set of 1-types. When $\mathfrak{A}$ is clear from context, denote by $A_\pi$ the set $\{a \in A | \mathrm{tp}^{\mathfrak{A}}[a] = \pi\}$, and denote by $A_\Pi$ the set $\{a \in A | \mathrm{tp}^{\mathfrak{A}}[a] \in \Pi\}$. In addition, denote by $\Pi^c$ the set of all and only those 1-types not contained in $\Pi$.

**Remark 11.** *For any structure $\mathfrak{A}$ and any set of 1-types $\Pi$, $A_{\Pi^c} = A \setminus A_\Pi$.*

We need to elaborate the notion of *profiles*, introduced in Definition 9. Recall, in this regard, that $\mu_1, \ldots, \mu_M$ are the message-types (invertible and non-invertible).

**Definition 19.** Let $\mathfrak{A}$ be a structure of finite degree, let $a$ be an element of $A$, and let $\Pi$ be any set of 1-types. The $\Pi$-*profile* of $a$ in $\mathfrak{A}$, denoted $\mathrm{pr}^{\mathfrak{A}}_\Pi[a]$, is the $M$-tuple $\bar{v} = (v_1, \ldots, v_M)$ of natural numbers where, for all $j$ $(1 \leq j \leq M)$,

$$v_j = |\{b \in A_\Pi : b \neq a \text{ and } \mathrm{tp}^{\mathfrak{A}}[a, b] = \mu_j\}|.$$

From Definition 9, we see that $\mathrm{pr}^{\mathfrak{A}}[a]$ is the tuple $\mathrm{pr}^{\mathfrak{A}}_\Pi[a]$ in the case where $\Pi$ is the set of all 1-types.

**Remark 12.** *For any set of 1-types $\Pi$, the tuple $\mathrm{pr}^{\mathfrak{A}}_\Pi[a]$ is obtained from the tuple $\mathrm{pr}^{\mathfrak{A}}[a]$ by simply zeroing its $j$th coordinate whenever $\mathrm{tp}_2(\mu_j)$ is not a member of $\Pi$. Hence, if $\Pi' \supseteq \Pi$, then $\mathrm{pr}^{\mathfrak{A}}_{\Pi'}[a]$ always determines $\mathrm{pr}^{\mathfrak{A}}_\Pi[a]$.*

For the next definition, recall that $f_1, \ldots, f_m$ are the counting predicates.

**Definition 20.** Let $\mathfrak{A}$ be a structure of finite degree, let $a$ be an element of $A$, and let $\Pi$ be any set of 1-types. The $\Pi$-*count* of $a$ in $\mathfrak{A}$, denoted $\mathrm{ct}^{\mathfrak{A}}_\Pi[a]$, is the $m$-tuple $\bar{u} = (u_1, \ldots, u_m)$ of natural numbers where, for all $h$ $(1 \leq h \leq m)$,

$$u_h = |\{b \in A_\Pi : b \neq a \text{ and } \mathfrak{A} \models f_h[a, b]\}|.$$

If $\Pi$ is the set of all 1-types, we call $\mathrm{ct}^{\mathfrak{A}}_\Pi[a]$ simply the *count* of $a$ in $\mathfrak{A}$, and denote it $\mathrm{ct}^{\mathfrak{A}}[a]$.

The tuple $\mathrm{ct}^{\mathfrak{A}}_\Pi[a]$ records, for each counting predicate $f_h$ $(1 \leq h \leq m)$, the number of elements with 1-type in $\Pi$ to which $a$ is non-reflexively related by $f_h$. In particular, $\mathrm{ct}^{\mathfrak{A}}[a]$ records, for each counting predicate $f_h$ $(1 \leq h \leq m)$, the number of elements to which $a$ is non-reflexively related by $f_h$. Notice that $\mathrm{pr}^{\mathfrak{A}}_\Pi[a]$ is an $M$-tuple, while $\mathrm{ct}^{\mathfrak{A}}_\Pi[a]$ is an $m$-tuple.

**Remark 13.** *For any set of 1-types* $\Pi$, $\mathrm{pr}^{\mathfrak{A}}_{\Pi}[a]$ *determines* $\mathrm{ct}^{\mathfrak{A}}_{\Pi}[a]$. *Specifically, if* $\mathrm{pr}^{\mathfrak{A}}_{\Pi}[a] = (v_1, \ldots, v_M)$ *and* $\mathrm{ct}^{\mathfrak{A}}_{\Pi}[a] = (u_1, \ldots, u_m)$ *we have, for all* $h$ $(1 \leq h \leq m)$

$$u_h = \sum\{v_j \mid 1 \leq j \leq M \text{ and } f_h(x,y) \in \mu_j\}.$$

Of course, $\mathrm{ct}^{\mathfrak{A}}_{\Pi}[a]$ does not in general determine $\mathrm{pr}^{\mathfrak{A}}_{\Pi}[a]$, because it contains strictly less information. Think of the tuple $\mathrm{ct}^{\mathfrak{A}}_{\Pi}[a]$ as providing a 'statistical summary' of the tuple $\mathrm{pr}^{\mathfrak{A}}_{\Pi}[a]$.

**Definition 21.** Let $\mathfrak{A}$ and $\mathfrak{A}'$ be chromatic structures of finite degree over some common domain $A$, let $\Pi$ be a set of 1-types, and let $B$ be a subset of $A$. We call $\mathfrak{A}'$ a $(\Pi, B)$-*approximation to* $\mathfrak{A}$ if every 2-type realized in $\mathfrak{A}'$ is also realized in $\mathfrak{A}$, and, for all $a \in A$:

1. $\mathrm{tp}^{\mathfrak{A}'}[a] = \mathrm{tp}^{\mathfrak{A}}[a]$;
2. $\mathrm{pr}^{\mathfrak{A}'}_{\Pi^c}[a] = \mathrm{pr}^{\mathfrak{A}}_{\Pi^c}[a]$;
3. $a \in A \setminus B$ implies $\mathrm{pr}^{\mathfrak{A}'}[a] = \mathrm{pr}^{\mathfrak{A}}[a]$;
4. $a \in B$ implies $\mathrm{ct}^{\mathfrak{A}'}_{\Pi}[a] = \mathrm{ct}^{\mathfrak{A}}_{\Pi}[a]$.

Very roughly, a $(\Pi, B)$-approximation to $\mathfrak{A}$ is a surgically modified version of $\mathfrak{A}$ in which only the $\Pi$-profiles of elements of $B$ have been interfered with. In particular: all elements of $A$ retain their old 1-types and their old $\Pi^c$-profiles; all elements of $A \setminus B$ retain their old profiles; and all elements of $B$ retain their old $\Pi$-counts. In addition, chromaticity is preserved, and no new 2-types (or 1-types) are introduced. We remark that, in Condition 4 of Definition 21, the restriction that $a \in B$ is in fact logically redundant, since if $a \notin B$, Condition 3 certainly entails $\mathrm{ct}^{\mathfrak{A}'}_{\Pi}[a] = \mathrm{ct}^{\mathfrak{A}}_{\Pi}[a]$.

**Remark 14.** *Let* $\mathfrak{A}$, $\mathfrak{A}'$ *and* $\mathfrak{A}''$ *be chromatic structures of finite degree over some common domain* $A$, *let* $\Pi$, $\Pi'$ *be sets of 1-types and let* $B$, $B'$ *be subsets of* $A$. *Then* $\mathfrak{A}$ *is a* $(\Pi, B)$*-approximation to itself. Furthermore, if* $\mathfrak{A}'$ *is a* $(\Pi, B)$*-approximation to* $\mathfrak{A}$, $\Pi \subseteq \Pi'$, *and* $B \subseteq B'$, *then* $\mathfrak{A}'$ *is also a* $(\Pi', B')$*-approximation to* $\mathfrak{A}$. *Finally, if* $\mathfrak{A}'$ *is a* $(\Pi, B)$*-approximation to* $\mathfrak{A}$, *and* $\mathfrak{A}''$ *is a* $(\Pi, B)$*-approximation to* $\mathfrak{A}'$, *then* $\mathfrak{A}''$ *is a* $(\Pi, B)$*-approximation to* $\mathfrak{A}$.

A crucial fact about $(\Pi, B)$-approximations is that they maintain satisfaction of certain Scott-form formulas.

**Lemma 26.** *Let* $\mathfrak{A}$ *be a chromatic structure of finite degree over some domain* $A$, *let* $\Pi$ *be a set of 1-types, let* $B$ *be a subset of* $A$, *and let* $\mathfrak{A}'$ *be a* $(\Pi, B)$*-approximation to* $\mathfrak{A}$. *Let* $\alpha$ *be a quantifier-free* $\mathcal{L}^2$*-formula and* $C_1, \ldots, C_m$ *be positive integers. Consider the formula*

$$\psi = \forall x \forall y(\alpha \vee x \approx y) \wedge \bigwedge_{1 \leq h \leq m} \forall x \exists_{=C_h} y(f_h(x,y) \wedge x \not\approx y).$$

*If* $\mathfrak{A} \models \psi$, *then* $\mathfrak{A}' \models \psi$.

*Proof.* By Remark 14, we may assume without loss of generality that $\Pi$ is the set of all 1-types and $B = A$. Since every 2-type realized in $\mathfrak{A}'$ is also realized in $\mathfrak{A}$, $\mathfrak{A}' \models \forall x \forall y(\alpha \vee x \approx y)$. And since $\mathrm{ct}^{\mathfrak{A}'}[a] = \mathrm{ct}^{\mathfrak{A}}[a]$ for every $a \in A$, $\mathfrak{A}' \models \bigwedge_{1 \leq h \leq m} \forall x \exists_{=C_h} y(f_h(x,y) \wedge x \not\approx y)$. □

Our strategy now is to show that, given a structure $\mathfrak{A}$ of finite degree, a set of 1-types $\Pi$ and a subset $B$ of $A$, we can find a $(\Pi, B)$-approximation to $\mathfrak{A}$ in which 'few' star-types are realized.

For the next definition, recall that $\mu_1, \ldots, \mu_{M^*}$ are the *invertible* message-types. Thus, for any structure $\mathfrak{A}$ of finite degree, and any 1-type $\pi$, the first $M^*$ coordinates of any $\{\pi\}$-profile $\mathrm{pr}^{\mathfrak{A}}_{\{\pi\}}[a]$ tell us, for each invertible message-type $\mu_j$ $(1 \leq j \leq M^*)$, how many elements in $A_\pi$ the element $a$ sends a message of type $\mu_j$ to. Notice that, if $\mathfrak{A}$ is chromatic, then, by Remark 3, the first $M^*$ coordinates of $\mathrm{pr}^{\mathfrak{A}}_{\{\pi\}}[a]$ are either all zero, or else are all zero except for a single occurrence of unity.

**Definition 22.** Let $\mathfrak{A}$ be a chromatic structure of finite degree over domain $A$, let $\Pi$ be a set of 1-types, let $\pi$ be a 1-type, and let $B$ be a subset of $A$. We say that $B$ is a $\Pi$-*group* if every element of $B$ has the same 1-type and every element of $B$ has the same $\Pi$-count. We say that $B$ is a $\pi$-*patch* if $B$ is a $\{\pi\}$-group and, for all $a, b \in B$, the tuples $\mathrm{pr}^{\mathfrak{A}}_{\{\pi\}}[a]$ and $\mathrm{pr}^{\mathfrak{A}}_{\{\pi\}}[b]$ agree in each of their first $M^*$ coordinates.

We now demonstrate that, if $B$ is a $\pi$-patch in a structure satisfying certain conditions, then there exists a $(\{\pi\}, B)$-approximation to that structure in which every element of $B$ has the same $\{\pi\}$-profiles.

**Lemma 27.** *Let $\mathfrak{A}$ be a chromatic, $(mC+1)^2$-differentiated structure, of degree at most $C$, over a domain $A$; let $\pi$ be a 1-type; and let $B \subseteq A$ be a $\pi$-patch in $\mathfrak{A}$. Then there exists a structure $\mathfrak{A}'$ such that $\mathfrak{A}'$ is a $(\{\pi\}, B)$-approximation to $\mathfrak{A}$ in which the elements of $B$ all have the same $\{\pi\}$-profile.*

*Proof.* Since a $\pi$-patch is by definition a $\{\pi\}$-group, there is some 1-type $\pi^*$ such that $B \subseteq A_{\pi^*}$. If $|B| \leq 1$, then $\mathfrak{A}' = \mathfrak{A}$ satisfies the conditions of the lemma. Hence, we may assume that $|B| > 1$. Similarly, if no element of $B$ sends a message to any element of $A_\pi$, then $\mathrm{pr}^{\mathfrak{A}}_{\{\pi\}}[b]$ is the 'zero-vector' (i.e. the tuple with all entries 0) for all $b \in B$, and $\mathfrak{A}' = \mathfrak{A}$ again satisfies the conditions of the Lemma. Hence, we may assume that some element of $B$ sends a message to some element of $A_\pi$, and, in particular, that $A_\pi$ is non-empty. Now suppose for the moment that $A_\pi$ is a singleton $\{a\}$. Thus, $\pi \neq \pi^*$, so that $B$ and $A_\pi$ are disjoint (Fig. 4.4). We claim that no element of $B$ can send an invertible message to $a$. For, if any does, then all do (remember that $B$ is a $\pi$-patch, so the profiles of all elements agree in their first $M^*$ coordinates), in which case $a$ sends invertible messages to more than one element having type $\pi^*$, contradicting the chromaticity of $\mathfrak{A}$. Thus, we may fix some element $b_0 \in B$ which sends a non-invertible message to $a$. Since $B$ is a $\{\pi\}$-group, it follows that every element of $B$ sends a message—and hence a non-invertible message—to $a$ (though the

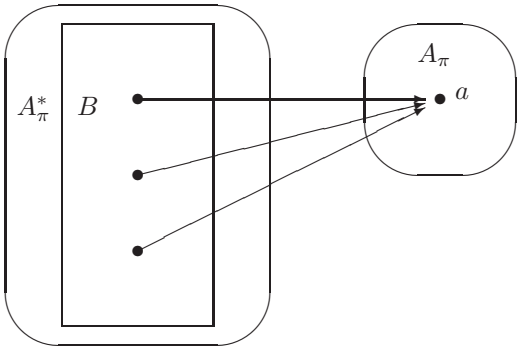


Figure 4.4: A $\pi$-patch $B$ with $|A_\pi| = 1$ and $|B| > 1$.

messages sent by different elements of $B$ need not all have the same type). In particular, $a$ does not send a message to any element of $B$. Now define the structure $\mathfrak{A}'$ to be exactly like $\mathfrak{A}$ except that, for every $b \in B$, we set

$$\mathrm{tp}^{\mathfrak{A}'}[b, a] = \mathrm{tp}^{\mathfrak{A}}[b_0, a]. \tag{4.2}$$

Clearly, $\mathfrak{A}'$ is chromatic and realizes no 2-types not realized in $\mathfrak{A}$. Clearly also, $\mathrm{tp}^{\mathfrak{A}'}[a] = \mathrm{tp}^{\mathfrak{A}}[a]$; $\mathrm{pr}^{\mathfrak{A}'}_{\Pi^c}[a] = \mathrm{pr}^{\mathfrak{A}}_{\Pi^c}[a]$, and $a \in A \setminus B$ implies $\mathrm{pr}^{\mathfrak{A}'}[a] = \mathrm{pr}^{\mathfrak{A}}[a]$, since no messages sent by any elements in $A \setminus B$ have been interfered with. Finally, by the construction in (4.2), and bearing in mind that $A_\pi = \{a\}$, we have $\mathrm{ct}^{\mathfrak{A}'}_{\{\pi\}}[a] = \mathrm{ct}^{\mathfrak{A}}_{\{\pi\}}[b]$. But since $B$ is a $\{\pi\}$-group, $\mathrm{ct}^{\mathfrak{A}}_{\{\pi\}}[b] = \mathrm{ct}^{\mathfrak{A}}_{\{\pi\}}[a]$, whence $\mathrm{ct}^{\mathfrak{A}'}_{\{\pi\}}[a] = \mathrm{ct}^{\mathfrak{A}}_{\{\pi\}}[a]$. Thus, we have shown that $\mathfrak{A}'$ is a $(\{\pi\}, B)$-approximation to $\mathfrak{A}$. Again, by the construction (4.2), every element of $B$ has the same $\{\pi\}$-profile in $\mathfrak{A}'$.

Thus, we may assume henceforth that $B$ and $A_\pi$ both contain more than one element. Since $\mathfrak{A}$ is $(mC+1)^2$-differentiated, $A_\pi$ and $A_{\pi^*} \supseteq B$ both contain more than $(mC + 1)^2$ elements. By Lemma 15, then, let $\tau$ be a silent 2-type such that, for some $a \in A_{\pi^*}$ and some $a' \in A_\pi$, $\mathrm{tp}^{\mathfrak{A}}[a, a'] = \tau$. We will use $\tau$ below as a 'filler' 2-type, for specifying 2-types when we do not want to disturb the profiles of elements in a structure we are manipulating.

For $a \in B$, let

$$A_a = \{a' \in A_\pi \mid a \neq a' \text{ and } \mathrm{tp}^{\mathfrak{A}}[a, a'] \text{ is a non-invertible message-type}\};$$

and for $a \notin B$, let $A_a = \emptyset$. Notice, incidentally, that $a' \in A_a$ implies $a \notin A_{a'}$. Choose $b \in B$ for which $|A_b|$ is smallest, and fix $b$. Enumerate $A_b$ as $b_1, b_2, \ldots$. For any $a \in B$ not equal to $b$, let $\hat{A}_a$ be a subset of $A_a$ having the same number of elements as $A_b$, and enumerate $\hat{A}_a$ as $a_1, a_2, \ldots$. The situation is depicted schematically in Fig. 4.5.

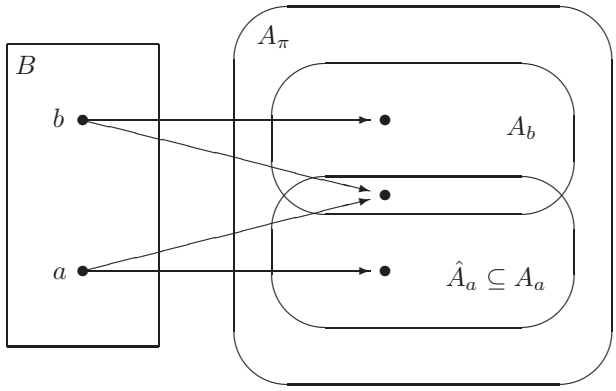


Figure 4.5: A $\pi$-patch $B$ with $|A_\pi|$ and $|B|$ both greater than 1. There is no assumption that $B$ and $A_\pi$ are disjoint.

We now define the structure $\mathfrak{A}'$ by assigning 2-types as follows. For all $a \in B$ such that $a \neq b$, set

$$\mathrm{tp}^{\mathfrak{A}'}[a, a_i] = \mathrm{tp}^{\mathfrak{A}}[b, b_i], \tag{4.3}$$

where $i$ ranges over the enumeration of $\hat{A}_a$, and set

$$\mathrm{tp}^{\mathfrak{A}'}[a, a'] = \tau, \tag{4.4}$$

where $a'$ is any element of $A_a \setminus \hat{A}_a$. (Remember: $\tau$ is our 'filler' 2-type.) In addition, for all distinct $a$, $a'$ such that $a' \notin A_a$ and $a \notin A_{a'}$, set

$$\mathrm{tp}^{\mathfrak{A}'}[a, a'] = \mathrm{tp}^{\mathfrak{A}}[a, a']. \tag{4.5}$$

Since $a' \in A_a$ implies $a \notin A_{a'}$, none of these assignments overwrites any other. And since $B \subseteq A_{\pi^*}$, the 1-type assignments implicit in (4.3)–(4.5) never clash: indeed, we have $\mathrm{tp}^{\mathfrak{A}'}[a] = \mathrm{tp}^{\mathfrak{A}}[a]$ for all $a \in A$. Furthermore, the transformation from $\mathfrak{A}$ to $\mathfrak{A}'$ does not affect invertible message-types. That is: for distinct $a$, $a'$, $\mathrm{tp}^{\mathfrak{A}}[a, a']$ is an invertible message-type if and only if $\mathrm{tp}^{\mathfrak{A}'}[a, a']$ is an invertible message-type; and moreover, if $\mathrm{tp}^{\mathfrak{A}}[a, a']$ is an invertible message-type, then $\mathrm{tp}^{\mathfrak{A}'}[a, a'] = \mathrm{tp}^{\mathfrak{A}}[a, a']$.

We now verify that $\mathfrak{A}'$ is a $(\{\pi\}, B)$-approximation to $\mathfrak{A}$. From the remarks of the previous paragraph and the fact that $\mathfrak{A}$ is chromatic, we have that $\mathfrak{A}'$ is also chromatic. In addition, it is immediate from (4.3)–(4.5) that every 2-type realized in $\mathfrak{A}'$ is also realized in $\mathfrak{A}$. Now let $a$ be any element of $B$. Since $B$ is a $\pi$-patch, the tuples $\mathrm{pr}^{\mathfrak{A}}_{\{\pi\}}[a]$ and $\mathrm{pr}^{\mathfrak{A}}_{\{\pi\}}[b]$ by definition agree in their

first $M^*$ coordinates (corresponding to the invertible message-types). Hence, since we have just shown that the transformation from $\mathfrak{A}$ to $\mathfrak{A}'$ does not affect invertible message-types, the tuples $\mathrm{pr}^{\mathfrak{A}'}_{\{\pi\}}[a]$ and $\mathrm{pr}^{\mathfrak{A}}_{\{\pi\}}[b]$ also agree in their first $M^*$ coordinates. Furthermore, the assignments (4.3)–(4.5) guarantee that $\mathrm{pr}^{\mathfrak{A}'}_{\{\pi\}}[a]$ and $\mathrm{pr}^{\mathfrak{A}}_{\{\pi\}}[b]$ also agree in the remaining coordinates $M^*+1,\ldots,M$ (corresponding to the non-invertible message-types). Hence,

$$\mathrm{pr}^{\mathfrak{A}'}_{\{\pi\}}[a] = \mathrm{pr}^{\mathfrak{A}}_{\{\pi\}}[b] \qquad \text{for all } a \in B. \tag{4.6}$$

It is now a simple matter to check the numbered conditions in Definition 21. Let $a$ be an arbitrary element of $A$.

1. We have already established that $\mathrm{tp}^{\mathfrak{A}'}[a] = \mathrm{tp}^{\mathfrak{A}}[a]$.
2. Let $a'$ be any element of $A_{\{\pi\}^c}$ with $a \neq a'$. Then certainly $a' \notin A_a \subseteq A_\pi$, so that $\mathrm{tp}^{\mathfrak{A}'}[a,a']$ can be different from $\mathrm{tp}^{\mathfrak{A}}[a,a']$ only if $a \in A_{a'}$. But if $a \in A_{a'}$, then neither $\mathrm{tp}^{\mathfrak{A}}[a,a']$ nor $\mathrm{tp}^{\mathfrak{A}'}[a,a']$ can be a message-type. Hence, $\mathrm{pr}^{\mathfrak{A}'}_{\{\pi\}^c}[a] = \mathrm{pr}^{\mathfrak{A}}_{\{\pi\}^c}[a]$.
3. Suppose $a \in A \setminus B$, and let $a'$ be any element of $A$ with $a \neq a'$. The argument now proceeds much as for the previous condition: certainly, $a' \notin A_a = \emptyset$, so $\mathrm{tp}^{\mathfrak{A}'}[a,a']$ can be different from $\mathrm{tp}^{\mathfrak{A}}[a,a']$ only if $a \in A_{a'}$. But if $a \in A_{a'}$, then neither $\mathrm{tp}^{\mathfrak{A}}[a,a']$ nor $\mathrm{tp}^{\mathfrak{A}'}[a,a']$ can be a message-type. Hence $\mathrm{pr}^{\mathfrak{A}'}[a] = \mathrm{pr}^{\mathfrak{A}}[a]$.
4. Suppose $a \in B$. Equation (4.6) yields $\mathrm{ct}^{\mathfrak{A}'}_{\{\pi\}}[a] = \mathrm{ct}^{\mathfrak{A}}_{\{\pi\}}[b]$. And since $B$ is a $\{\pi\}$-group, $\mathrm{ct}^{\mathfrak{A}}_{\{\pi\}}[b] = \mathrm{ct}^{\mathfrak{A}}_{\{\pi\}}[a]$. It follows that $\mathrm{ct}^{\mathfrak{A}'}_{\{\pi\}}[a] = \mathrm{ct}^{\mathfrak{A}}_{\{\pi\}}[a]$.

Finally, it is immediate from Equation (4.6) that all elements of $B$ have the same $\{\pi\}$-profile in $\mathfrak{A}'$. □

We now demonstrate that, if $B$ is a $\Pi$-group in a $(mC+1)^2$-differentiated, chromatic structure $\mathfrak{A}$ with $\deg(\mathfrak{A}) = C$, then, by taking a $(\Pi, B)$-approximation to that model, we can bound the number of $\Pi$-profiles realized by the elements of $B$. This demonstration will occupy Lemmas 28–30. Our strategy is first to partition $\Pi$ into roughly equal sets $\Pi'$ and $\Pi''$. We then recursively bound the number of $\Pi'$- and $\Pi''$-profiles realized by elements of $B$, and finally align these $\Pi'$- and $\Pi''$-profiles so as to bound the number of $\Pi$-profiles that result.

**Lemma 28.** *Let $\mathfrak{A}$ be a chromatic structure of finite degree over domain $A$, let $\Pi$ be a set of 1-types, let $B \subseteq A$ be a $\Pi$-group, and let $\omega$ be a permutation of $B$. Then there exists a structure $\mathfrak{A}'$ such that $\mathfrak{A}'$ is a $(\Pi, B)$-approximation to $\mathfrak{A}$, and for all $b \in B$, $\mathrm{pr}^{\mathfrak{A}'}_{\Pi}[\omega(b)] = \mathrm{pr}^{\mathfrak{A}}_{\Pi}[b]$.*

*Proof.* First, extend $\omega$ to the whole of $A$ by setting $\omega(a) = a$ for $a \in A \setminus B$. Next, for all $b \in A$, define:

$$\omega_{b\Pi}(a) = \begin{cases} \omega(a) \text{ if } b \in A_\Pi \\ a \text{ otherwise;} \end{cases}$$

Thus, $\omega_{b\Pi}$ is a permutation of $A$ (which may be the identity). Since $(\omega_{b\Pi})^{-1}$ and $(\omega^{-1})_{b\Pi}$ are the same permutation, we may unambiguously write $\omega_{b\Pi}^{-1}$. Clearly, $\omega$ fixes $B$ setwise and $A \setminus B$ pointwise; so, therefore, does $\omega_{b\Pi}^{-1}$. Moreover, since $B$ is a $\Pi$-group, every element of $B$ by definition has the same 1-type, and this 1-type is either a member of $\Pi$ or it is not. Hence, either $B \subseteq A_\Pi$ or $B \subseteq A_{\Pi^c} = A \setminus A_\Pi$. Thus, $\omega$ fixes both $A_\Pi$ and $A_{\Pi^c}$ setwise, and so therefore does $\omega_{b\Pi}^{-1}$.

Define the structure $\mathfrak{A}'$ over domain $A$ by setting, for all distinct $a, a' \in A$:

$$\mathrm{tp}^{\mathfrak{A}'}[a, a'] = \mathrm{tp}^{\mathfrak{A}}[\omega_{a'\Pi}^{-1}(a), \omega_{a\Pi}^{-1}(a')]. \tag{4.7}$$

To show that $\mathfrak{A}'$ is well-defined, we must show first that the elements $\omega_{a'\Pi}^{-1}(a)$ and $\omega_{a\Pi}^{-1}(a')$ in each instance of (4.7) are distinct, and second, that the 1-type assignments implicit in the different instances of (4.7) do not clash. Suppose, then that $a \neq a'$; we prove that $\omega_{a'\Pi}^{-1}(a) \neq \omega_{a\Pi}^{-1}(a')$. Since the permutations $\omega_{a\Pi}^{-1}$ and $\omega_{a'\Pi}^{-1}$ fix $B$ setwise and $A \setminus B$ pointwise, we may assume that $a, a' \in B$. We have already noted that either $B \subseteq A_\Pi$ or $B \subseteq A \setminus A_\Pi$. If $B \subseteq A_\Pi$, then $\omega_{a'\Pi}^{-1}(a) = \omega^{-1}(a)$ and $\omega_{a\Pi}^{-1}(a') = \omega^{-1}(a')$; if, on the other hand, $B \subseteq A \setminus A_\Pi$, then $\omega_{a'\Pi}^{-1}(a) = a$ and $\omega_{a\Pi}^{-1}(a') = a'$. Either way, $\omega_{a'\Pi}^{-1}(a) \neq \omega_{a\Pi}^{-1}(a')$. Next, we prove that the 1-type assignments in (4.7) never clash. Since all elements of $B$ have the same 1-type, and since $\omega$ is the identity outside $B$, we have, for all $a, a'$, $\mathrm{tp}^{\mathfrak{A}}[\omega_{a'\Pi}^{-1}(a)] = \mathrm{tp}^{\mathfrak{A}}[a]$; thus, $\mathrm{tp}^{\mathfrak{A}}[\omega_{a'\Pi}^{-1}(a)]$ does not depend on $a'$. Hence, the 1-type assignments implicit in (4.7) cannot clash, and $\mathfrak{A}'$ is indeed well-defined. In fact, this argument establishes that $\mathrm{tp}^{\mathfrak{A}'}[a] = \mathrm{tp}^{\mathfrak{A}}[a]$ for all $a \in A$.

We first check the numbered conditions of Definition 21 in turn. Let $a$ be an arbitrary element of $A$.

1. We have just established that $\mathrm{tp}^{\mathfrak{A}'}[a] = \mathrm{tp}^{\mathfrak{A}}[a]$.

2. For all $b \in A_{\Pi^c}$, $\omega_{b\Pi}^{-1}(a) = a$; in particular, if $a \in A_{\Pi^c}$, then $\omega_{a\Pi}^{-1}(a) = a$. Therefore, $\omega_{a\Pi}^{-1}$ is always a permutation of $A_{\Pi^c} \setminus \{a\}$, and moreover, for all $b \in A_{\Pi^c} \setminus \{a\}$, $\mathrm{tp}^{\mathfrak{A}'}[a, b] = \mathrm{tp}^{\mathfrak{A}}[a, \omega_{a\Pi}^{-1}(b)]$. Thus, the list of 2-types $\mathrm{tp}^{\mathfrak{A}'}[a, b]$ obtained as $b$ ranges over $A_{\Pi^c} \setminus \{a\}$ is (in some order) the list of 2-types $\mathrm{tp}^{\mathfrak{A}}[a, b']$ obtained as $b'$ ranges over $A_{\Pi^c} \setminus \{a\}$. It follows that $\mathrm{pr}_{\Pi^c}^{\mathfrak{A}'}[a] = \mathrm{pr}_{\Pi^c}^{\mathfrak{A}}[a]$.

3. Suppose $a \in A \setminus B$. Then, for all $b \in A$, $\omega_{b\Pi}^{-1}(a) = a$; in particular, $\omega_{a\Pi}^{-1}(a) = a$. Therefore, $\omega_{a\Pi}^{-1}$ is a permutation of $A \setminus \{a\}$, and moreover, for all $b \in A \setminus \{a\}$, $\mathrm{tp}^{\mathfrak{A}'}[a, b] = \mathrm{tp}^{\mathfrak{A}}[a, \omega_{a\Pi}^{-1}(b)]$. Thus, the list of 2-types $\mathrm{tp}^{\mathfrak{A}'}[a, b]$ obtained as $b$ ranges over $A \setminus \{a\}$ is (in some order) the list of 2-types $\mathrm{tp}^{\mathfrak{A}}[a, b']$ obtained as $b'$ ranges over $A \setminus \{a\}$. It follows that $\mathrm{pr}^{\mathfrak{A}'}[a] = \mathrm{pr}^{\mathfrak{A}}[a]$.

4. For all $b \in A_\Pi$, $\omega_{b\Pi}^{-1}(a) = \omega^{-1}(a)$; in particular, if $a \in A_\Pi$, then $\omega_{a\Pi}^{-1}(a) = \omega^{-1}(a)$. Therefore, $\omega_{a\Pi}^{-1}$ is a bijection from the set $A_\Pi \setminus \{a\}$ to the set $A_\Pi \setminus \{\omega^{-1}(a)\}$, and moreover, for all $b \in A_\Pi \setminus \{a\}$, $\mathrm{tp}^{\mathfrak{A}'}[a, b] =$

$\mathrm{tp}^{\mathfrak{A}}[\omega^{-1}(a), \omega_{a\Pi}^{-1}(b)]$. Thus, the list of 2-types $\mathrm{tp}^{\mathfrak{A}'}[a,b]$ obtained as $b$ ranges over $A_\Pi \setminus \{a\}$ is (in some order) the list of 2-types $\mathrm{tp}^{\mathfrak{A}}[\omega^{-1}(a), b']$ obtained as $b'$ ranges over $A_\Pi \setminus \{\omega^{-1}(a)\}$. It follows that

$$\mathrm{pr}_\Pi^{\mathfrak{A}'}[a] = \mathrm{pr}_\Pi^{\mathfrak{A}}[\omega^{-1}(a)]. \tag{4.8}$$

Certainly, then, we have $\mathrm{ct}_\Pi^{\mathfrak{A}'}[a] = \mathrm{ct}_\Pi^{\mathfrak{A}}[\omega^{-1}(a)]$. But since $B$ is a $\Pi$-group in $\mathfrak{A}$, $a \in B$ implies $\mathrm{ct}_\Pi^{\mathfrak{A}}[\omega^{-1}(a)] = \mathrm{ct}_\Pi^{\mathfrak{A}}[a]$, whence $\mathrm{ct}_\Pi^{\mathfrak{A}'}[a] = \mathrm{ct}_\Pi^{\mathfrak{A}}[a]$.

We have thus established that, for all $a \in A$, $\mathrm{pr}_\Pi^{\mathfrak{A}'}[a] = \mathrm{pr}_\Pi^{\mathfrak{A}}[\omega^{-1}(a)]$ and $\mathrm{pr}_{\Pi^c}^{\mathfrak{A}'}[a] = \mathrm{pr}_{\Pi^c}^{\mathfrak{A}}[a]$. Since $\mathfrak{A}$ is chromatic, it follows easily that $\mathfrak{A}'$ is chromatic. Moreover, all 2-types realized in $\mathfrak{A}'$ are realized in $\mathfrak{A}$. Hence, $\mathfrak{A}'$ is a $(\Pi, B)$-approximation to $\mathfrak{A}$. Finally, it follows from Equation (4.8) that, for all $b \in B$, $\mathrm{pr}_\Pi^{\mathfrak{A}'}[\omega(b)] = \mathrm{pr}_\Pi^{\mathfrak{A}}[b]$. □

Suppose $\mathfrak{A}$, $\Pi$ and $B$ are as in Lemma 28. That lemma then assures us that, as long as we are content to work with $(\Pi, B)$-approximations, we can permute the $\Pi$-profiles of the elements in $B$ at will! The following lemma exploits this facility.

**Lemma 29.** *Let $\mathfrak{A}$ be a chromatic structure of finite degree over domain $A$; let $\Pi'$, $\Pi''$ be disjoint, non-empty sets of 1-types; and let $\Pi = \Pi' \cup \Pi''$. Suppose the non-empty set $B \subseteq A$ is a $\Pi''$-group. Let the number of different $\Pi'$-profiles realized in $\mathfrak{A}$ by the elements of $B$ be $J'$; and let the number of different $\Pi''$-profiles realized in $\mathfrak{A}$ by the elements of $B$ be $J''$. Then there exists a $(\Pi, B)$-approximation $\mathfrak{A}'$ to $\mathfrak{A}$ in which at most $J' + J'' - 1$ different $\Pi$-profiles are realized by the elements of $B$.*

*Proof.* For perspicuity, we assume first that $B$ is finite. Enumerate $B$ as $b_1, \ldots, b_I$. Let the various $\Pi'$-profiles realized by at least one element of $B$ be $\mathbf{v}'_1, \ldots, \mathbf{v}'_{J'}$; and let the various $\Pi''$-profiles realized by at least one element of $B$ be $\mathbf{v}''_1, \ldots, \mathbf{v}''_{J''}$. By re-numbering the $b_1, \ldots, b_I$ if necessary, we may assume without loss of generality that the $\Pi'$-profiles of the $b_1, \ldots, b_I$ fall into consecutive blocks as depicted in the middle column in Fig. 4.6. More precisely, we have integers $0 = I_1 < I_2 < \cdots < I_{J'+1} = I$ such that, for all $j$ $(1 \leq j \leq J')$, $\mathrm{pr}_{\Pi'}^{\mathfrak{A}'}[b_i] = \mathbf{v}'_j$ for $i$ in the range $[I_j + 1, I_{j+1}]$. Since $B$ is a $\Pi''$-group, by Lemma 28, we can obtain a structure $\mathfrak{A}'$ such that $\mathfrak{A}'$ is a $(\Pi'', B)$-approximation to $\mathfrak{A}$ in which the elements of $B$ have $\Pi''$-profiles likewise arranged in consecutive blocks. Since $\mathfrak{A}'$ is a $(\Pi'', B)$-approximation to $\mathfrak{A}$ and the sets $\Pi'$ and $\Pi''$ are disjoint, the $\Pi'$-profiles of the elements of $B$ will be unaffected by the transformation from $\mathfrak{A}$ to $\mathfrak{A}'$: a typical alignment of $\Pi'$-profiles and $\Pi''$-profiles in $\mathfrak{A}'$ is shown in Fig. 4.6. By inspection, at most $J' + J'' - 1$ $\Pi$-profiles are realized in $\mathfrak{A}'$ by the elements of $B$.

The same argument applies in the case where $B$ is infinite, with only minor modifications. Of course, the number of different $\Pi'$- and $\Pi''$-profiles realized by the elements of $B$ can still only be finite, but some of the resulting blocks of $B$ may contain infinitely many entries. The matching up of these blocks so that at most $J' + J'' - 1$ $\Pi$-profiles result is routine. □

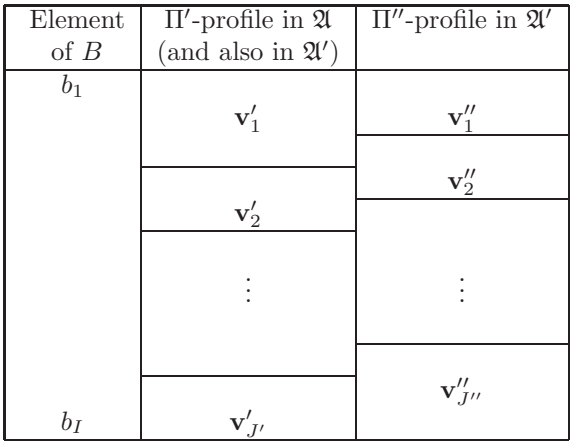


Figure 4.6: Arrangement of $\Pi'$-profiles and $\Pi''$-profiles in $B$.

Recall that $m$ is the number of featured predicates and $M^*$ the number of invertible message-types.

**Lemma 30.** *Fix an $(mC+1)^2$-differentiated, chromatic structure $\mathfrak{A}$, of finite degree. Let $C = \deg(\mathfrak{A})$, let $l \geq 0$, let $\Pi$ be a non-empty set of 1-types such that $|\Pi| \leq 2^l$, and let $B \subseteq A$ be a $\Pi$-group. Then there is a structure $\mathfrak{A}'$ such that $\mathfrak{A}'$ is a $(\Pi, B)$-approximation to $\mathfrak{A}$ and the number of different $\Pi$-profiles realized in $\mathfrak{A}'$ by the elements of $B$ is at most $2^l(M^*+1)(C+1)^{lm}$.*

*Proof.* By induction on $l$. To aid readability, let $K_l$ stand for $2^l(M^*+1)(C+1)^{lm}$. If $l = 0$, $\Pi$ is a singleton, so write $\Pi = \{\pi\}$. Decompose $B$ into maximal $\pi$-patches $B_1, \ldots, B_H$. Since $\mathfrak{A}$ is chromatic, Remark 3 guarantees that the first $M^*$ coordinates of any tuple $\mathrm{pr}^{\mathfrak{A}}_{\{\pi\}}[a]$ are either all zero, or else are all zero except for a single occurrence of unity. Therefore, $H \leq M^*+1$. Now let $\mathfrak{A}_0 = \mathfrak{A}$, and for all $h$ $(1 \leq h \leq H)$, apply Lemma 27 to obtain a structure $\mathfrak{A}_h$ such that $\mathfrak{A}_h$ is a $(\{\pi\}, B_h)$-approximation to $\mathfrak{A}_{h-1}$ in which the elements of $B_h$ all have the same $\{\pi\}$-profile. By Remark 14, $\mathfrak{A}_H$ is a $(\{\pi\}, B)$-approximation to $\mathfrak{A}$. And because the $B_h$ are pairwise disjoint, $1 \leq h < h' \leq H$ implies $\mathrm{pr}^{\mathfrak{A}_{h'}}[a] = \mathrm{pr}^{\mathfrak{A}_h}[a]$ for all $a \in B_h$. Hence, the total number of $\{\pi\}$-profiles realized by elements of $B$ in $\mathfrak{A}_H$ is at most $H \leq M^*+1 = K_0$. Thus, setting $\mathfrak{A}' = \mathfrak{A}_H$ establishes the case $l = 0$.

Now suppose $l > 0$. We may assume $\Pi$ is not a singleton, since otherwise, we can employ the argument of the case $l = 0$; so let $\Pi$ be partitioned into non-empty sets $\Pi'$ and $\Pi''$ each of cardinality at most $2^{l-1}$. Also, partition $B$ into maximal $\Pi'$-groups $B_1, \ldots, B_H$ (say). Since $B$ is a $\Pi$-group, the $B_1, \ldots, B_H$ will also be $\Pi''$-groups. Moreover, since $\deg(\mathfrak{A}) = C$, the $\Pi$-count of any element in $\mathfrak{A}$ is one of at most $(C+1)^m$ different tuples; and since $B$ is a $\Pi$-group, all elements of $B$ must have the same 1-type, whence $H \leq (C+1)^m$.

Again, let $\mathfrak{A}_0 = \mathfrak{A}$, and consider the set $B_1$. By inductive hypothesis, let $\mathfrak{A}'_1$ be a $(\Pi', B_1)$-approximation to $\mathfrak{A}_0$ in which at most $K_{l-1}$ $\Pi'$-profiles are realized by the elements of $B_1$. Again, by inductive hypothesis, let $\mathfrak{A}''_1$ be a $(\Pi'', B_1)$-approximation to $\mathfrak{A}'_1$ in which at most $K_{l-1}$ $\Pi''$-profiles are realized by the elements of $B_1$. Thus, in the structure $\mathfrak{A}''_1$, $B_1$ is a $\Pi'$-group realizing at most $K_{l-1}$ different $\Pi'$-profiles, and also a $\Pi''$-group realizing at most $K_{l-1}$ different $\Pi''$-profiles. By Lemma 29, let $\mathfrak{A}_1$ be a $(\Pi, B_1)$-approximation to $\mathfrak{A}''_1$ in which the elements of $B_1$ realize at most $2K_{l-1} - 1 < 2K_{l-1}$ different $\Pi$-profiles. By Remark 14, $\mathfrak{A}_1$ is a $(\Pi, B_1)$-approximation to $\mathfrak{A}_0$. Treating the sets $B_2, \ldots, B_H$ in the same way, we obtain structures $\mathfrak{A}_h$ $(1 \leq h \leq H)$ such that, for each $h$ in this range, $\mathfrak{A}_h$ is a $(\Pi, B_h)$-approximation to $\mathfrak{A}_{h-1}$ in which at most $2K_{l-1}$ different $\Pi$-profiles are realized by the elements of $B_h$. By Remark 14, $\mathfrak{A}_H$ is a $(\Pi, B)$-approximation to $\mathfrak{A}$. And because the $B_h$ are pairwise disjoint, $1 \leq h < h' \leq H$ implies that $\mathrm{pr}^{\mathfrak{A}_{h'}}[a] = \mathrm{pr}^{\mathfrak{A}_h}[a]$ for all $a \in B_h$. Hence, the total number of $\Pi$-profiles realized by elements of $B$ in $\mathfrak{A}_H$ is at most $2HK_{l-1} \leq 2(C+1)^m K_{l-1} = K_l$. Thus, setting $\mathfrak{A}' = \mathfrak{A}_H$ completes the induction. □

## 4.5 Complexity of $\mathcal{C}^2$

Finally, we may bring all of the above results together.

**Theorem 7.** *The problems Sat-$\mathcal{C}^2$ and Fin-Sat-$\mathcal{C}^2$ are both NEXPTIME-complete.*

*Proof.* The NEXPTIME-hardness of Fin-Sat-$\mathcal{C}^2$ and of Sat-$\mathcal{C}^2$ follows from Lemma 5 and the fact that, by Lemma 4, $\mathcal{L}^2$ has the finite model property. Hence we need only show that Sat-$\mathcal{C}^2$ and Fin-Sat-$\mathcal{C}^2$ are in NEXPTIME.

Let the $\mathcal{C}^2$-formula $\phi$ be given. By Lemma 14, we may produce, in polynomial time, a Scott-form formula

$$\psi := \forall x \forall y(\alpha \vee x \approx y) \wedge \bigwedge_{1 \leq h \leq m} \forall x \exists_{=C_h} y(f_h(x,y) \wedge x \not\approx y), \tag{4.9}$$

such that $\psi$ is satisfiable over any domain of size at least $C+1$ if and only if $\phi$ is, where $C = \max_{1 \leq h \leq m} C_h$. Since $\psi$ certainly has no model of size $C$ or less, then, $\phi$ is (finitely) satisfiable if and only if either (i) $\phi$ has a model of size $C$ or less, or (ii) $\psi$ is (finitely) satisfiable. Moreover, since $C$ is bounded by a singly exponential function of $\|\phi\|$, Case (i) can obviously be checked by a non-deterministic Turing machine running in time bounded by a singly exponential function of $\|\phi\|$. Therefore, we need only deal with Case (ii). Note that in the formula (4.9), $\alpha$ is a quantifier-free $\mathcal{L}^2$ formula, $1 \leq m \leq \|\phi\|$, the $C_1, \ldots, C_m$ are all positive and bounded by a singly exponential function of $\|\phi\|$, and the $f_1, \ldots, f_m$ are pairwise distinct.

Let $\Sigma$ be the signature of $\psi$ together with $d$ additional unary predicates, where

$$d = \lceil \log(mC+1)^2 \rceil + \lceil \log((mC)^2 + 1) \rceil,$$

and let $\Sigma$ be treated as a classified signature, by taking the featured predicates to be $f_1, \ldots, f_m$. Thus, $\Sigma$ satisfies the conditions set out at the beginning of Sections 4.3 and 4.4, so we may employ all the results derived there. Let $n = \max(\|\psi\|, |\Sigma|)$. Thus, $n$ is bounded by a polynomial function of $\|\phi\|$. We shall establish the following claim.

**Claim.** The formula $\psi$ is (finitely) satisfiable if and only if there exists a chromatic, $C$-bounded, (finitely) $3mC$-soluble frame $\mathcal{F}$, with dimension at most $2^{n^3+6n}$, such that $\mathcal{F} \models \psi$.

Assuming the truth of this claim, we may establish the satisfiability of $\psi$ using the following nondeterministic procedure.

1. Guess a chromatic, $C$-bounded, frame $\mathcal{F}$, with dimension $N$ for some $N \leq 2^{n^3+6n}$.
2. Check that $\mathcal{F} \models \psi$ (Definition 18).
3. Determine whether $\mathcal{F}$ has a (finite) $3mC$-solution (Definitions 16 and 17).

The number of symbols needed to encode $\mathcal{F}$ (under any sensible encoding) is bounded by a singly exponential function of $n$. So, therefore, is the time required to check that $\mathcal{F}$ is chromatic and that $\mathcal{F} \models \psi$. That the existence of a (finite) $3mC$-solution of $\mathcal{F}$ can be checked in nondeterministic time bounded by an exponential function of $n$ follows from Lemma 24 (or Lemma 18).

It remains only to establish the above claim. Suppose first that $\psi$ has a (finite) model $\mathfrak{A}$ of finite degree interpreting $\Sigma$. Since $\Sigma$ contains $d$ unary predicates not occurring in $\psi$, Lemmas 16 and 17 guarantee that, by re-interpreting these predicates if necessary, we may take $\mathfrak{A}$ to be chromatic and $(mC+1)^2$-differentiated. Letting $\Pi$ be the set of all 1-types over $\Sigma$, and $B$ be any of the sets $A_\pi$ (with $\pi$ a 1-type), Lemma 30 yields a structure $\mathfrak{A}'$ such that $\mathfrak{A}'$ is a $(\Pi, B)$-approximation to $\mathfrak{A}$ realizing at most $2^n(M^*+1)(C+1)^{nm}$ different profiles. Repeating this process for all 1-types realized in $\mathfrak{A}$, we obtain a structure $\mathfrak{A}^*$ such that $\mathfrak{A}^*$ is a $(\Pi, A)$-approximation to $\mathfrak{A}$ realizing at most $2^{2n}(M^*+1)(C+1)^{nm}$ different star-types. Simple calculation shows that $(C+1) \leq 2^n$ and $(M^*+1) \leq 2^{4n}$, so that the number of star-types realized in $\mathfrak{A}^*$ is bounded by $2^{n^3+6n}$. By Lemma 26, $\mathfrak{A}^* \models \psi$. By Lemma 25, let $\mathcal{F}$ be a frame describing $\mathfrak{A}^*$ such that $\mathcal{F} \models \psi$. Thus, the dimension of $\mathcal{F}$ (= the number of star-types realized in $\mathfrak{A}^*$) is at most $2^{n^3+6n}$. By Remark 8, $\mathcal{F}$ is $C$-bounded and chromatic. By Lemma 22 (Lemma 19), $\mathcal{F}$ is (finitely) $(mC+1)^2$-soluble and hence (finitely) $3mC$-soluble.

Conversely, suppose there exists a $C$-bounded, chromatic, (finitely) $3mC$-soluble frame $\mathcal{F}$ over $\Sigma$, such that $\mathcal{F} \models \psi$. Then by Lemma 23 (Lemma 20), there exists some (finite) structure $\mathfrak{A}$ such that $\mathcal{F}$ describes $\mathfrak{A}$, and by Lemma 25, $\mathfrak{A} \models \psi$. This establishes the above claim, and hence the theorem. □

The proof of Theorem 7 actually yields a little more information. Recall the sequence $\{\phi_n\}$ of finitely satisfiable $\mathcal{C}^2$-formulas presented in Example 4. We

showed there that, while $\|\phi_n\|$ grows as a polynomial function of $n$, the size of the smallest satisfying models grows as a doubly exponential function of $n$. We can now see that this is, essentially, as bad as it gets:

**Corollary 4.** *Let $\phi$ be a formula of $\mathcal{C}^2$. If $\phi$ is finitely satisfiable, then it is satisfiable in a structure of size bounded by a doubly exponential function of $\|\phi\|$.*

*Proof.* The structure built in Lemma 20 from $\mathcal{F}$ and its solution $\bar{w}$ has domain of cardinality $w_1 + \cdots + w_N$. □

Notice that the complexity result of Theorem 7 is better than one might naïvely expect on the basis of the small model property of Corollary 4. Thus, while $\mathcal{C}^2$ is expressive enough to force very large models, it is not expressive enough simultaneously to say a great deal about them: in particular, no $\mathcal{C}^2$-formula can prevent the kind of structure manipulation discussed in Section 4.4.

## 4.6 Bibliographic notes

The decidability of Sat-$\mathcal{C}^2$ and Fin-Sat-$\mathcal{C}^2$ was proved, though with no complexity bounds, in a very terse paper, by Grädel, Otto and Rosen [13] (who are also responsible for Example 4). A complexity bound of (in effect) 2-NEXPTIME for Sat-$\mathcal{C}^2$ was reported by Pacholski, Szwast and Tendera [25, 26]. This was finally reduced to NEXPTIME (and extended to Fin-Sat-$\mathcal{C}^2$) by Pratt-Hartmann [30].

# Chapter 5

# Modal and Guarded Logics with Counting

## 5.1 Modal logic

Fix a countably infinite set $\Pi$. The language of *modal logic* is defined to be the smallest set of expressions, $\mathcal{M}$, satisfying the following conditions:

1. $\Pi \subseteq \mathcal{M}$;
2. if $\phi$ and $\psi$ are in $\mathcal{M}$, then so are $\neg\phi$, $\phi \wedge \psi$, $\phi \vee \psi$, $\phi \rightarrow \psi$ and $\phi \leftrightarrow \psi$;
3. if $\phi$ is in $\mathcal{M}$, then so are $\Diamond\phi$ and $\Box\phi$.

We refer to expressions in this set as $\mathcal{M}$-*formulas* (or simply *formulas*, if clear from context).

Let $\Sigma$ be the relational signature with unary predicates $\Pi$ and single binary predicate $r$, and let $\mathfrak{A}$ be a $\Sigma$-structure with domain $W$. We define the *satisfaction* relation for $\mathcal{M}$-formulas inductively as follows:

1. $\mathfrak{A} \models_w p$ if and only if $w \in p^{\mathfrak{A}}$;
2. $\mathfrak{A} \models_w \neg\phi$ if and only if $\mathfrak{A} \not\models_w \phi$, and similarly for $\wedge$, $\vee$, $\rightarrow$, $\leftrightarrow$;
3. $\mathfrak{A} \models_w \Diamond\phi$ if and only if there exists $v \in W$ such that $\langle w, v\rangle \in r^{\mathfrak{A}}$ and $\mathfrak{A} \models_v \phi$;
4. $\mathfrak{A} \models_w \Box\phi$ if and only if, for all $v \in W$ such that $\langle w, v\rangle \in r^{\mathfrak{A}}$, $\mathfrak{A} \models_v \phi$.

The notion of satisfaction extends to sets of $\mathcal{M}$-formulas $\Phi$ as expected: $\mathfrak{A} \models_w \Phi$ if $\mathfrak{A} \models_w \phi$ for all $\phi \in \Phi$.

The study of modal logic was original motivated by philosophical investigations into the concepts of necessity and possibility. In this connection, we are invited to think of the elements of $W$ as *worlds*, the elements of $\Pi$ as *proposition*

| | |
|---|---|
| reflexive frames | $\forall x r(x,x)$ |
| serial frames | $\forall x \exists y r(x,y)$ |
| symmetric frames | $\forall x \forall y (r(x,y) \rightarrow r(y,x))$ |
| transitive frames | $\forall x \forall y \forall z (r(x,y) \wedge r(y,z) \rightarrow r(x,z))$ |
| Euclidean frames | $\forall x \forall y \forall z (r(x,y) \wedge r(x,z) \rightarrow r(y,z))$. |

Table 5.1: Frame classes considered in these notes.

*letters*, and the relation $r^{\mathfrak{A}} \subseteq W \times W$ as a relation of *accessibility* on possible worlds. If $\mathfrak{A} \models_w \phi$, we say that $\phi$ is *true at* the world $w$. Thus, $\Diamond\phi$ may be read: "$\phi$ is true at some world accessible from the current world" (or: "Possibly, $\phi$"). Likewise, $\Box\phi$ may be read: "$\phi$ is true at all worlds accessible from the current world" (or: "Necessarily, $\phi$"). The Boolean connectives are of course understood as in propositional logic. In these notes, we shall not concern ourselves with the philosophical value of this formal reconstruction.

By a *frame*, we mean an $\{r\}$-structure—in other words, a non-empty (possibly infinite) digraph. If $\mathfrak{A}$ is a $\Sigma$-structure, then its $\{r\}$-reduct is a frame $\mathfrak{F}$: we say that $\mathfrak{A}$ is a structure *over* $\mathfrak{F}$. Further, we call the mapping $V : \Pi \rightarrow \mathbb{P}(W)$ given by $p \mapsto p^{\mathfrak{A}}$ the *valuation* of $\mathfrak{A}$ (*on* $W$). We write $\mathfrak{A} = (W, R, V)$ to indicate that $\mathfrak{A}$ is a $\Sigma$-structure over the frame $(W, R)$ with valuation $V$. Obviously, this determines $\mathfrak{A}$ completely. Henceforth, when discussing modal logic (and its variants), the term "structure", with no signature qualification, will always mean "$\Sigma$-structure". Let $\phi$ be an $\mathcal{M}$-formula. We say that $\phi$ is *satisfiable over* a frame $\mathfrak{F}$ if there exists a structure $\mathfrak{A}$ over $\mathfrak{F}$ and a world $w$ of $\mathfrak{A}$ such that $\mathfrak{A} \models_w \phi$. Further, $\phi$ is *satisfiable over* a class of frames $\mathcal{K}$ if it is satisfiable over some frame in $\mathcal{K}$. We denote by $\mathcal{M}_{\mathcal{K}}$-*Sat* the problem of determining whether a given $\mathcal{M}$-formula is satisfiable over $\mathcal{K}$.

Any first-order sentence $\alpha$ over the signature $\{r\}$ defines a class of frames $\{\mathfrak{F} : \mathfrak{F} \models \alpha\}$. The most common frame classes are those listed, together with their respective defining first-order sentences, Table 5.1. We denote the class of reflexive frames by Rfl, and, similarly, the classes of serial, symmetric, transitive and Euclidean frames by Ser, Sym, Tr and Eucl, respectively. Note that, in the presence of symmetry, transitivity is equivalent to the Euclidean property. A structure over a reflexive frame will simply be called a *reflexive* structure, and similarly for the other frame properties. The first question we ask is: what is the complexity of the satisfiability problem for $\mathcal{M}$ over any frame-class characterized by a conjunction of (zero or more of) these properties?

The following notation will be useful. Let $\mathcal{F}$ be a subset (possibly empty) of the set of frame classes $\{\mathrm{Rfl}, \mathrm{Ser}, \mathrm{Sym}, \mathrm{Tr}, \mathrm{Eucl}\}$. Then $\cap\mathcal{F}$ simply denotes the class of frames having all the properties listed in $\mathcal{F}$. For instance, $\bigcap\{\mathrm{Rfl}, \mathrm{Tr}\}$ denotes the class of reflexive, transitive frames, $\bigcap\{\mathrm{Ser}, \mathrm{Tr}, \mathrm{Eucl}\}$ denotes the class of serial, transitive, Euclidean frames, and so on. As a limiting case, $\bigcap\emptyset$ denotes the class of all frames.

The following two theorems are well-known, and may be proved using techniques found in any modern text on modal logic (e.g. [3], Ch. 6).

**Theorem 8.** *Let* $\mathcal{F} \subseteq \{\mathrm{Rfl}, \mathrm{Ser}, \mathrm{Sym}, \mathrm{Tr}, \mathrm{Eucl}\}$*, with* $\mathrm{Eucl} \in \mathcal{F}$ *or* $\{\mathrm{Sym}, \mathrm{Tr}\} \subseteq \mathcal{F}$*. Then the satisfiability problem for modal logic over* $\bigcap \mathcal{F}$ *is NPTime-complete.*

**Theorem 9.** *If* $\mathcal{F} \subseteq \{\mathrm{Rfl}, \mathrm{Ser}, \mathrm{Tr}\}$*, then the satisfiability problem for modal logic over* $\bigcap \mathcal{F}$ *is PSpace-complete. Also, if* $\mathcal{F} \subseteq \{\mathrm{Rfl}, \mathrm{Ser}, \mathrm{Sym}\}$*, then the satisfiability problem for modal logic over* $\bigcap \mathcal{F}$ *is PSpace-complete.*

## 5.2 Graded modal logic

The language of *graded modal logic* is defined to be the smallest set of expressions, $\mathcal{GM}$, satisfying the following conditions:

1. $\Pi \subseteq \mathcal{GM}$;
2. if $\phi$ and $\psi$ are in $\mathcal{GM}$, then so are $\neg\phi$, $\phi \wedge \psi$, $\phi \vee \psi$, $\phi \rightarrow \psi$ and $\phi \leftrightarrow \psi$;
3. if $\phi$ is in $\mathcal{GM}$, then so are $\Diamond_{\leq C}\phi$ and $\Diamond_{\geq C}\phi$, for any bit-string $C$.

We refer to expressions in this set as $\mathcal{GM}$*-formulas* (or simply *formulas*, if clear from context).

We define the *satisfaction* relation for $\mathcal{GM}$-formulas inductively as follows:

1. $\mathfrak{A} \models_w p$ if and only if $w \in p^{\mathfrak{A}}$;
2. $\mathfrak{A} \models_w \neg\phi$ if and only if $\mathfrak{A} \not\models_w \phi$, and similarly for $\wedge$, $\vee$, $\rightarrow$, $\leftrightarrow$;
3. $\mathfrak{A} \models_w \Diamond_{\geq C}\phi$ if and only if there exist at least $C$ worlds $v \in W$ such that $\langle w, v\rangle \in r^{\mathfrak{A}}$ and $\mathfrak{A} \models_v \phi$;
4. $\mathfrak{A} \models_w \Diamond_{\leq C}\phi$ if and only if there exist at most $C$ worlds $v \in W$ such that $\langle w, v\rangle \in r^{\mathfrak{A}}$ and $\mathfrak{A} \models_v \phi$.

The notion of satisfaction extends to sets of $\mathcal{GM}$-formulas $\Phi$ as expected. Again, for a given structure $\mathfrak{A} = (W, R, V)$, we think of the elements of $\Pi$ as proposition letters to which $V$ assigns truth-values relative to the worlds in $W$. Thus, graded modal logic is the formal language obtained by decorating the $\Diamond$-operator of ordinary modal logic with subscripts expressing cardinality constraints. Specifically, for $C \geq 0$, the formula $\Diamond_{\leq C}\phi$ may be glossed: "$\phi$ is true at no more than $C$ accessible worlds," and the formula $\Diamond_{\geq C}\phi$ may be glossed: "$\phi$ is true at no fewer than $C$ accessible worlds." From a philosophical point of view, one could be forgiven for thinking this sort of modal accountancy poorly motivated. However, if we consider $\mathcal{M}$ to be, in effect, a fragment of first-order logic, the language $\mathcal{GM}$ is, in the context of the present enquiry, a very natural generalization. Furthermore, the results we obtain are non-trivial and not at all what one might naïvely expect.

We denote by $\mathcal{GM}_{\mathcal{K}}$-*Sat* the problem of determining whether a given $\mathcal{GM}$-formula is satisfiable over $\mathcal{K}$. In this section, we ask the following. Let $\mathcal{F}$ be a subset (possibly empty) of the set of frame classes $\{\mathrm{Rfl}, \mathrm{Ser}, \mathrm{Sym}, \mathrm{Tr}, \mathrm{Eucl}\}$. We ask: what is the complexity of $\mathcal{GM}_{\cap\mathcal{F}}$-Sat?

The results are encapsulated in the following three theorems.

**Theorem 10.** *Let* $\mathcal{F} \subseteq \{\mathrm{Rfl}, \mathrm{Ser}, \mathrm{Sym}, \mathrm{Tr}, \mathrm{Eucl}\}$*, with* $\mathrm{Eucl} \in \mathcal{F}$ *or* $\{\mathrm{Sym}, \mathrm{Tr}\} \subseteq \mathcal{F}$*. Then the satisfiability problem for graded modal logic over* $\bigcap \mathcal{F}$ *is NPTime-complete.*

**Theorem 11.** *Let* $\mathcal{F} \subseteq \{\mathrm{Rfl}, \mathrm{Ser}, \mathrm{Sym}\}$*. Then the satisfiability problem for graded modal logic over* $\bigcap \mathcal{F}$ *is PSpace-complete.*

**Theorem 12.** *Let* $\mathcal{F} \subseteq \{\mathrm{Rfl}, \mathrm{Ser}, \mathrm{Tr}\}$*, with* $\mathrm{Tr} \in \mathcal{F}$*. Then the satisfiability problem for graded modal logic over* $\bigcap \mathcal{F}$ *is NExpTime-complete. It remains NExpTime-hard, even when all numerical subscripts in modal operators are at most* 1*.*

Thus, adding counting to the operators of modal logic produces a somewhat more complicated complexity-theoretic landscape. The proofs of these theorems are omitted from these notes.

## 5.3 The guarded fragment

The relational semantics for the language $\mathcal{M}$ means that we can regard it as a fragment of first-order logic in which quantification is restricted to the patterns $\forall v(r(u, v) \rightarrow \psi(v))$ and $\exists v(r(u, v) \wedge \psi(v))$. This idea can be generalized.

We again fix some purely relational signature $\Sigma$, but this time with predicates of any arity, and we again denote by $\mathcal{L}$ the language of first-order logic over $\Sigma$. The *guarded fragment* of first-order logic, $\mathcal{G}$, is defined to be the smallest set of $\mathcal{L}$-formulas satisfying the following conditions:

1. every atomic formula is in $\mathcal{G}$;
2. $\mathcal{G}$ is closed under Boolean connectives;
3. if $\psi \in \mathcal{G}$ with $\mathrm{Vars}(\psi) = \bar{x}, \bar{y}$, and $p$ is a predicate of the same arity as $\bar{x}, \bar{y}$, then $\forall \bar{x}(p(\bar{x}, \bar{y}) \rightarrow \psi) \in \mathcal{G}$ and $\exists \bar{x}(p(\bar{x}, \bar{y}) \wedge \psi) \in \mathcal{G}$;
4. if $\psi \in \mathcal{G}$, and $\psi$ has at most one free variable, then $\forall x \psi \in \mathcal{G}$ and $\exists x \psi \in \mathcal{G}$.

For $k > 1$, we denote by $\mathcal{G}^k_{\approx}$ the subset of $\mathcal{G}$ involving at most the variables $x_1, \ldots x_k$. Thus, $\mathcal{G}^k = \mathcal{G} \cap \mathcal{L}^k_{\approx}$, for all $k > 0$. $\mathcal{L}^k_{\approx}$ not involving the equality predicate. We call $\mathcal{G}^k$ the *k-variable guarded fragment.* Thus, $\mathcal{M}$ is a proper subset of the 2-variable guarded fragment $\mathcal{G}^2$.

We have three important theorems concerning $\mathcal{G}$:

**Theorem 13.** *The problem* $\mathcal{G}$*-Sat is 2-*ExpTime*-complete.*

**Theorem 14.** *For every* $k > 0$*, the problem* $\mathcal{G}^k$*-Sat is* ExpTime*-complete.*

**Theorem 15.** *The guarded fragment has the finite model property.*

It follows from Theorem 15 that the problems $\mathcal{G}$-Sat and $\mathcal{G}$-Fin-Sat coincide, as do $\mathcal{G}^k$-Sat and $\mathcal{G}^k$-Fin-Sat, for all $k > 0$.

The proofs of Theorems 13–15 are omitted from these notes.

## 5.4 The guarded fragment with counting

In this section, we consider the effect of adding counting quantifiers to the guarded fragment.

In the sequel, we restrict consideration to a purely relational signature of 0-ary, unary and binary predicates. If $r$ is any binary predicate (including $\approx$), we call an atomic formula having either of the forms $r(x, y)$ or $r(y, x)$ a *guard-atom*. The two-variable guarded fragment with counting quantifiers, $\mathcal{GC}^2$, can then be defined as the smallest set of formulas satisfying the following conditions:

1. $\mathcal{GC}^2$ contains all atomic formulas and is closed under Boolean combinations;
2. if $\phi$ is a formula of $\mathcal{GC}^2$ with at most one free variable, and $u$ is a variable (i.e. either $x$ or $y$), then the formulas $\forall u\phi$ and $\exists u\phi$ are in $\mathcal{GC}^2$;
3. if $\phi$ is a formula of $\mathcal{GC}^2$, $\gamma$ a guard-atom, $u$ a variable, and $Q$ any of the quantifiers $\exists$, $\exists_{\leq C}$, $\exists_{\geq C}$, $\exists_{=C}$ (for $C > 0$), then the formulas $\forall u(\gamma \rightarrow \phi)$, $Qu(\gamma \wedge \phi)$ and $Qu\gamma$ are in $\mathcal{GC}^2$.

According to the above syntax, the non-counting quantifiers $\exists$ and $\forall$ may apply without restriction to formulas with at most one free variable; however, they may apply to formulas with two free variables only in the presence of a guard-atom. By contrast, the counting quantifiers $\exists_{\leq C}$, $\exists_{\geq C}$, $\exists_{=C}$ may only ever apply in the presence of a guard atom (which by definition has two free variables). Note in particular that the formula $\exists_{=1} xp(x)$ is not in $\mathcal{GC}^2$. In fact, the next lemma shows that no formula of $\mathcal{GC}^2$ can force a predicate $p$ to be uniquely instantiated in its models.

**Lemma 31.** *Let $\phi$ be a formula of $\mathcal{GC}^2$ with signature $\Sigma$ (so that $\Sigma$ has no individual constants), $\mathfrak{A}$ a structure interpreting $\Sigma$, and $I$ a nonempty set. For $i \in I$, let $\mathfrak{A}_i$ be a copy of $\mathfrak{A}$, with the domains $A_i$ pairwise disjoint. If $\phi$ is satisfied in $\mathfrak{A}$, then it is satisfied in the structure $\mathfrak{A}'$ with domain $A' = \bigcup_{i \in I} A_i$ and interpretations $\sigma^{\mathfrak{A}'} = \bigcup_{i \in I} \sigma^{\mathfrak{A}_i}$ for every $\sigma \in \Sigma$.*

*Proof.* If $\theta : \{x, y\} \rightarrow A$ is any variable assignment, and $i \in I$, let $\theta_i$ be the variable assignment which maps $x$ and $y$ to the corresponding elements in $A_i \subseteq A'$. A routine structural induction on $\phi$ shows that $\mathfrak{A} \models_\theta \phi$ if and only if, for some (= for all) $i \in I$, $\mathfrak{A}' \models_{\theta_i} \phi$. □

It follows immediately that, if a formula of $\mathcal{GC}^2$ has a finite model, then it has arbitrarily large finite models, and indeed infinite models.

**Lemma 32.** *Let $\phi$ be a $\mathcal{GC}^2$-formula. We can compute, in time bounded by a polynomial function of $\|\phi\|$, a formula*

$$\psi = \forall x\alpha \wedge \bigwedge_{1 \leq h \leq l} \forall x \forall y(e_h(x, y) \rightarrow (\beta_h \vee x \approx y)) \wedge \bigwedge_{1 \leq i \leq m} \forall x \exists_{=C_i} y(f_i(x, y) \wedge x \not\approx y) \tag{5.1}$$

*such that:* $(i)$ $\alpha$ *is a quantifier-free formula not involving* $\approx$ *with* $x$ *as its only variable;* $(ii)$ $l$ *and* $m$ *are positive integers;* $(iii)$ *for all* $h$ $(1 \leq h \leq l)$, $e_h$ *is a binary predicate other than* $\approx$, *and* $\beta_h$ *is a quantifier-free formula not involving* $\approx$ *with* $x$ *and* $y$ *as its only variables;* $(iv)$ *for all* $i$ $(1 \leq i \leq m)$, $C_i$ *is a positive integer, and* $f_i$ *is a binary predicate other than* $\approx$*;* $(v)$ $\phi$ *is satisfiable if and only if* $\psi$ *is satisfiable;* $(vi)$ $\phi$ *is finitely satisfiable if and only if* $\psi$ *is finitely satisfiable.*

*Proof.* We proceed exactly as for Lemma 14. Note that, if $\phi$ is finitely satisfiable, then it has models over arbitrarily large finite domains. Therefore, we may without loss of generality assume that the size of any model of $\phi$ of interest is greater than any quantifier subscript mentioned in $\phi$.

The only difficulty is to show that the conjuncts of the form $\forall u \forall v \chi$ generated in Stage 1 of the procedure described there can be made guarded. To see that this is so, consider the treatment of a subformula $\theta(u) = \exists_{\leq D} v \chi$ of $\phi_0$. Since $\phi_0$ is guarded, $\chi$ must be of the form $\gamma \wedge \eta$, where $\gamma$ is an atomic formula $g(u, v)$ or $g(v, u)$. Again, we define $\phi_1 := \phi_0[p(u)/\theta(u)]$, where $p$ is a new unary predicate, and

$$\begin{aligned}\psi_1 := \forall u \exists_{=D} v r_1(u, v) \wedge \forall u \exists_{=D+1} v r_2(u, v) \wedge \\ \forall u \forall v (p(u) \rightarrow (\theta \rightarrow r_1(u, v))) \wedge \\ \forall u \forall v (\neg p(u) \rightarrow (r_2(u, v) \rightarrow \theta)),\end{aligned}$$

arguing, just as in Lemma 14 that $\psi_0$ and $\phi_1 \wedge \psi_1$ are satisfiable over sufficiently large domains. Now consider in more detail $\psi_1$. Its latter two conjuncts are, as they stand, not guarded. However, noting that $\theta$ is $(\gamma \wedge \eta)$, we see that these conjuncts are in fact logically equivalent to the guarded formula

$$\forall u \forall v (\gamma \rightarrow ((p(u) \wedge \eta) \rightarrow r_1(u, v))) \wedge \forall u \forall v (\neg r_2(u, v) \rightarrow (p(u) \rightarrow (\gamma \wedge \eta))).$$

The other cases are dealt with similarly, and we obtain the desired formula $\psi$. □

To show that the (finite) satisfiability problem for $\mathcal{GC}^2$ is in ExpTime, it therefore suffices to consider only formulas $\phi$ of the form (5.1). Furthermore, we may assume without loss of generality that no 0-ary predicates (proposition letters) occur in $\phi$, since we can consider each of the (at most $2^{\|\phi\|}$) truth-value assignments to the 0-ary predicates of $\phi$ in turn, replacing each 0-ary predicate with $\top$ or $\bot$ according to its truth-value in the considered assignment.

Accordingly, fix $\phi$ to be some formula of the form (5.1) over a signature of unary and binary predicates. Set $C = \max_{1 \leq i \leq m} C_i$, and let $\Sigma$ be the signature of $\phi$ together with $\log((mC)^2 + 1)$ (rounded up) new unary predicates. Thus, $|\Sigma|$ is bounded by a polynomial (actually, linear) function of $\|\phi\|$. Since $\Sigma$ is the only signature we shall be concerned with in the sequel, we generally suppress reference to it. Thus, 'predicate' henceforth means 'predicate in $\Sigma \cup \{\approx\}$', 'structure' henceforth means 'structure interpreting $\Sigma$', and so on. We keep the

meanings of the symbols

$$\Sigma, \phi, \alpha, l, e_1, \ldots, e_l, \beta_1, \ldots, \beta_l, m, C_1, \ldots, C_m, C, f_1, \ldots, f_m$$

fixed throughout this paper. The predicates $f_1, \ldots, f_m$ will play a key role in the ensuing argument; we refer to them as the *counting predicates*. There is no restriction on these predicates' occurring in other parts of $\phi$.

Again, we can regard the signature $\Sigma$ as a classified signature ...

Let the 1-types (over $\Sigma$) be enumerated in some order as the sequence

$$\Pi = \pi_0, \ldots, \pi_{P-1}.$$

Evidently, $P$ is a power of 2, so $p = \log P$ is an integer. (Actually, $p = |\Sigma|$.) Now let $s$ be any bit string ($0 \leq |s| \leq p$), and denote the string of length 0 by $\epsilon$. We inductively define the sub-sequence $\Pi_s$ of $\Pi$ by setting $\Pi_\epsilon$ to be the whole of $\Pi$, and setting $\Pi_{s0}$ and $\Pi_{s1}$ to be the left and right halves of $\Pi_s$, respectively. Formally:

$$\Pi_\epsilon \quad = \quad \pi_0, \ldots, \pi_{P-1},$$

and if $\Pi_s = \pi_j, \ldots, \pi_{k-1}$, with $|s| < p$,

$$\begin{aligned} \Pi_{s0} &= \pi_j, \ldots, \pi_{\frac{k+j}{2}-1} \\ \Pi_{s1} &= \pi_{\frac{k+j}{2}}, \ldots, \pi_{k-1}. \end{aligned}$$

Thus, if $|s| = p$, then $\Pi_s$ is a one-element sequence $\pi_j$, where $j$ is the integer ($0 \leq j < P$) encoded by the bit-string $s$ in the usual way. To avoid clumsy circumlocutions, we occasionally equivocate between bit-strings and the integers they encode, thus, for instance, writing $\pi_s$ instead of $\pi_j$ in this case. But we will only ever write $\pi_s$ if $|s| = p$. In addition, we occasionally for convenience treat sequences as if they were sets, writing, for instance, $\pi \in \Pi_s$.

We now use the sequences $\Pi_s$ to define sets of invertible message-types indexed by bit-strings as follows. Let $\Lambda$ be the set of all invertible message-types. If $\pi$ is any 1-type, and $s$ is any bit-string such that $|s| \leq p$, let

$$\Lambda_{\pi,s} = \{\lambda \in \Lambda \mid \mathrm{tp}_1(\lambda) = \pi \text{ and } \mathrm{tp}_2(\lambda) \in \Pi_s\}.$$

Thus, the $\Lambda_{\pi,s}$ are sets of invertible message-types identified purely by their terminal 1-types. Except in very special cases, these sets will contain more than one member, even when $|s| = p$. However, for chromatic structures, we have the following important fact.

**Lemma 33.** *Let $\mathfrak{A}$ be a chromatic structure, $a \in A$, $\pi = \mathrm{tp}^{\mathfrak{A}}[a]$, and $s$ a bit-string with $|s| = p$. Then there can be at most one element $b \in A \setminus \{a\}$ such that $\mathrm{tp}^{\mathfrak{A}}[a, b] \in \Lambda_{\pi,s}$.*

*Proof.* Any two such elements would be connected by a chain of two invertible message-types, and would both have 1-type $\pi_s$. □

Finally, we use bit strings to index sequences of 2-types that are not invertible message-types. Again, fix any 1-type $\pi$, and consider the set of non-invertible message-types $\mu$ such that $\mathrm{tp}_1(\mu) = \pi$. Let these be enumerated in some way as a sequence

$$\mu_{\pi,0}, \ldots, \mu_{\pi,R-1}.$$

Furthermore, consider the set of silent 2-types $\mu$ such that $\mathrm{tp}_1(\mu) = \pi$. Let these be enumerated in some way as a sequence

$$\mu_{\pi,R}, \ldots, \mu_{\pi,Q-1}.$$

Thus, the sequence

$$M_\pi = \mu_{\pi,0}, \ldots, \mu_{\pi,Q-1}$$

is an enumeration of precisely those 2-types $\tau$ such that $\mathrm{tp}_1(\tau) = \pi$ and $\tau^{-1}$ is not a message-type. Evidently, $R$ and $Q$ are independent of the choice of $\pi$; moreover, $Q$ is a power of 2, so $q = \log Q$ is an integer. (We remark that $R$ is not a power of 2.) Let $t$ be any bit string ($0 \leq |t| \leq q$). We inductively define the sub-sequence $M_{\pi,t}$ of $M_\pi$ by setting $M_{\pi,\epsilon}$ to be the whole of $M_\pi$, and setting $M_{\pi,t0}$ and $M_{\pi,t1}$ to be the left and right halves of $M_{\pi,t}$, respectively. Formally:

$$M_{\pi,\epsilon} \quad = \quad \mu_{\pi,0}, \ldots, \mu_{\pi,Q-1},$$

and if $M_{\pi,t} = \mu_{\pi,j}, \ldots, \mu_{\pi,k-1}$, with $|t| < q$,

$$\begin{aligned} M_{\pi,t0} &= \mu_{\pi,j}, \ldots, \mu_{\pi,\frac{k+j}{2}-1} \\ M_{\pi,t1} &= \mu_{\pi,\frac{k+j}{2}}, \ldots, \mu_{\pi,k-1}. \end{aligned}$$

Thus, if $|t| = q$, then $M_{\pi,t}$ is a one-element sequence $\mu_{\pi,j}$, where $j$ is the integer ($0 \leq j < Q$) encoded by the bit-string $t$ in the usual way. Again we may for convenience write $\mu_{\pi,t}$ instead of $\mu_{\pi,j}$ in this case, but here too we only ever write $\mu_{\pi,t}$ if $|t| = q$.

## 5.5 Spectra and tallies

The approach taken here involves identifying various configurational properties of elements in finite structures. These we now proceed to define. We continue to use the symbols introduced in Section 5.4 with their advertised meanings. In particular, $f_1, \ldots, f_m$ are the counting predicates occurring in the formula $\phi$ given in (5.1), and the integers $C_1, \ldots, C_m$ are the corresponding numerical quantifier subscripts.

Let $\mathfrak{A}$ be a finite structure, $a \in A$, and $\pi = \mathrm{tp}^{\mathfrak{A}}[a]$, and suppose $\mathfrak{A} \models \phi$. Evidently, for all $i$ ($1 \leq i \leq m$), there are exactly $C_i$ elements $b \in A \setminus \{a\}$ such that $\mathfrak{A} \models f_i[a, b]$:

$$C_i = |\{b \in A \setminus \{a\} : \mathfrak{A} \models f_i[a, b] \ \}|.$$

Now, for any bit-string $s$ ($0 \leq |s| \leq p$), define the *s-spectrum* of $a$ in $\mathfrak{A}$, denoted $\mathrm{sp}_s^{\mathfrak{A}}[a]$, to be the $m$-element vector whose $i$th component ($1 \leq i \leq m$) is the number of elements $b \in A \setminus \{a\}$ such that $\mathfrak{A} \models f_i[a,b]$ and $\mathrm{tp}^{\mathfrak{A}}[a,b] \in \Lambda_{\pi,s}$:

$$\left(\mathrm{sp}_s^{\mathfrak{A}}[a]\right)_i = |\{b \in A \setminus \{a\} : \mathfrak{A} \models f_i[a,b] \text{ and } \mathrm{tp}^{\mathfrak{A}}[a,b] \in \Lambda_{\pi,s}\}|.$$

Similarly, for any bit-string $t$ ($0 \leq |t| \leq q$), define the *t-tally* of $a$ in $\mathfrak{A}$, denoted $\mathrm{tl}_t^{\mathfrak{A}}[a]$, to be the $m$-element vector whose $i$th component is the number of elements $b \in A \setminus \{a\}$ such that $\mathfrak{A} \models f_i[a,b]$ and $\mathrm{tp}^{\mathfrak{A}}[a,b] \in M_{\pi,t}$:

$$\left(\mathrm{tl}_t^{\mathfrak{A}}[a]\right)_i = |\{b \in A \setminus \{a\} : \mathfrak{A} \models f_i[a,b] \text{ and } \mathrm{tp}^{\mathfrak{A}}[a,b] \in M_{\pi,t}\}|.$$

Henceforth, by *vector*, we shall always mean "$m$-dimensional vector over $\mathbb{N}$", with indices running from 1 to $m$. We denote the vector $(C_1, \ldots, C_m)$ by $\mathbf{C}$ and the $m$-dimensional zero vector $(0, \ldots, 0)$ by $\mathbf{0}$. If $\mathbf{u}$ and $\mathbf{v}$ are vectors, we write $\mathbf{u} \leq \mathbf{v}$ if every component of $\mathbf{u}$ is less than or equal to the corresponding component of $\mathbf{v}$; we write $\mathbf{u} < \mathbf{v}$ if $\mathbf{u} \leq \mathbf{v}$ and $\mathbf{u} \neq \mathbf{v}$. Similarly for $\geq$ and $>$. Recalling further that $C = \max_{1 \leq i \leq m} C_i$, the number of vectors $\mathbf{u}$ such that $\mathbf{u} \leq \mathbf{C}$ is evidently bounded by $(C+1)^m$, and hence by an exponential function of $\|\phi\|$. Moreover, the $s$-spectrum and $t$-tally of $a$ is always a vector $\leq \mathbf{C}$. Lastly, given any 2-type $\tau$, we write $\mathbf{C}_\tau$ for the vector whose $i$th component is given by:

$$\left(\mathbf{C}_\tau\right)_i = \begin{cases} 1 & \text{if } f_i(x,y) \in \tau, \\ 0 & \text{otherwise.} \end{cases} \tag{5.2}$$

To better understand this apparatus, let $\mathfrak{A}$, $a$, $\pi$ be as above, and consider first the case where $s$ and $t$ are the empty string $\epsilon$. Since $\Lambda_{\pi,\epsilon}$ is the set of invertible message-types $\lambda$ such that $\mathrm{tp}_1(\lambda) = \pi$, $\mathrm{sp}_\epsilon^{\mathfrak{A}}[a]$ is simply the vector whose $i$th component records the number of elements $b$ to which $a$ sends a message of *invertible* type containing the atom $f_i(x,y)$. Likewise, $\mathrm{tl}_\epsilon^{\mathfrak{A}}[a]$ is the vector whose $i$th component records the number of elements $b$ to which $a$ sends a message of *non-invertible* type containing the atom $f_i(x,y)$. If $0 < |s| \leq p$, then $\mathrm{sp}_s^{\mathfrak{A}}[a]$ is obtained in the same way as $\mathrm{sp}_\epsilon^{\mathfrak{A}}[a]$, except that we discount all messages whose type is not a member of $\Lambda_{\pi,s}$. Likewise, if $0 < |t| \leq q$, then $\mathrm{tl}_t^{\mathfrak{A}}[a]$ is obtained in the same way as $\mathrm{tl}_\epsilon^{\mathfrak{A}}[a]$, except that we discount all messages whose type is not a member of $M_{\pi,t}$.

**Lemma 34.** *Suppose $\mathfrak{A}$ is a model of $\phi$. Let $a \in A$, $\pi = \mathrm{tp}^{\mathfrak{A}}[a]$, and $s$, $t$ be bit-strings such that $|s| < p$ and $|t| < q$. Then*

$$\mathrm{sp}_\epsilon^{\mathfrak{A}}[a] + \mathrm{tl}_\epsilon^{\mathfrak{A}}[a] = \mathbf{C} \tag{5.3}$$

$$\mathrm{sp}_{s0}^{\mathfrak{A}}[a] + \mathrm{sp}_{s1}^{\mathfrak{A}}[a] = \mathrm{sp}_s^{\mathfrak{A}}[a] \tag{5.4}$$

$$\mathrm{tl}_{t0}^{\mathfrak{A}}[a] + \mathrm{tl}_{t1}^{\mathfrak{A}}[a] = \mathrm{tl}_t^{\mathfrak{A}}[a]. \tag{5.5}$$

*Proof.* Immediate. □

Thus, while $\Lambda_{\pi,s}$ is the *union* of the sets $\Lambda_{\pi,s0}$ and $\Lambda_{\pi,s1}$, $\mathrm{sp}^{\mathfrak{A}}_s[a]$ is the *vector sum* of the spectra $\mathrm{sp}^{\mathfrak{A}}_{s0}[a]$ and $\mathrm{sp}^{\mathfrak{A}}_{s1}[a]$; similarly for tallies.

As the strings $s$ and $t$ get longer, the sets $\Lambda_{\pi,s}$ and the sequences $M_{\pi,t}$ get smaller, and so the vectors $\mathrm{sp}^{\mathfrak{A}}_s[a]$ and $\mathrm{tl}^{\mathfrak{A}}_t[a]$ become, as it were, more selective in the information they record. It is therefore instructive to consider what happens to $s$-spectra and $t$-tallies when $s$ and $t$ have maximal length. The latter case is the easier to understand, and so we consider it first. If $|t| = q$, then $M_{\pi,t}$ by construction contains just one 2-type, $\mu = \mu_{\pi,t}$, which is either a non-invertible message-type, or else a silent 2-type. If $\mu$ is a non-invertible message-type, then, since $\mathfrak{A} \models \phi$, there can be only finitely many elements $b \in A \setminus \{a\}$ such that $\mathrm{tp}^{\mathfrak{A}}[a, b] = \mu$. Let this number be $n$. Evidently:

$$\mathrm{tl}^{\mathfrak{A}}_t[a] = n\mathbf{C}_\mu. \tag{5.6}$$

On the other hand, if $\mu$ is a silent 2-type, then $\mathrm{tl}^{\mathfrak{A}}_t[a] = \mathbf{C}_\mu = \mathbf{0}$.

Turning now to $s$-spectra with $|s| = p$, recall that $\Lambda_{\pi,s}$ in general has more than one element. However, Lemma 33 tells us that, in a *chromatic* model, no element may send more than one message whose type is in $\Lambda_{\pi,s}$. This enables us to establish the following simple result, which will be useful in the sequel:

**Lemma 35.** *Suppose that $\mathfrak{A}$ is a chromatic model of $\phi$. Let $a \in A$, $\pi$ be a 1-type, and $s$ be a bit-string with $|s| = p$. If $\mathrm{tp}^{\mathfrak{A}}[a] = \pi$ and $\mathrm{sp}^{\mathfrak{A}}_s[a] > \mathbf{0}$, then there exists $\lambda \in \Lambda_{\pi,s}$ with $\mathrm{sp}^{\mathfrak{A}}_s[a] = \mathbf{C}_\lambda$ such that $a$ sends a message of type $\lambda$ to some $b \in A \setminus \{a\}$. Conversely, if there exists $\lambda \in \Lambda_{\pi,s}$ such that $a$ sends a message of type $\lambda$ to some $b \in A \setminus \{a\}$, then $\mathrm{tp}^{\mathfrak{A}}[a] = \pi$ and $\mathrm{sp}^{\mathfrak{A}}_s[a] = \mathbf{C}_\lambda$.*

*Proof.* Suppose $\mathrm{tp}^{\mathfrak{A}}[a] = \pi$ and $\mathrm{sp}^{\mathfrak{A}}_s[a] > \mathbf{0}$. Then there exists $b \in A \setminus \{a\}$ such that $\mathrm{tp}^{\mathfrak{A}}[a, b] \in \Lambda_{\pi,s}$. By Lemma 33, this $b$ is unique, so, letting $\lambda = \mathrm{tp}^{\mathfrak{A}}[a, b]$, we have $\mathrm{sp}^{\mathfrak{A}}_s[a] = \mathbf{C}_\lambda$ as required. Conversely, suppose $a$ sends a message of type $\lambda \in \Lambda_{\pi,s}$ to some element $b \in A \setminus \{a\}$. Certainly, then, $\mathrm{tp}^{\mathfrak{A}}[a] = \pi$, and again, by Lemma 33, this $b$ is the only element in $A \setminus \{a\}$ to which $a$ sends a message having any type in $\Lambda_{\pi,s}$; it follows that $\mathrm{sp}^{\mathfrak{A}}_s[a] = \mathbf{C}_\lambda$ as required. □

## 5.6 Transformation into a constraint satisfaction problem

In the sequel, we take $\pi$ to vary over the set of 1-types, $\lambda$ to vary over the set of invertible message-types, $s$ to vary over the set of bit-strings of length at most $p$, $t$ to vary over the set of bit-strings of length at most $q$, and $\mathbf{u}$, $\mathbf{v}$ and $\mathbf{w}$ to vary over the set of vectors $\leq \mathbf{C}$. (Similarly for their primed counterparts $\pi'$, $\lambda'$, $s'$, $t'$, $\mathbf{u}'$, $\mathbf{v}'$ and $\mathbf{w}'$.) We refer to these sets as the *standard ranges* of the respective letters; they are summarized in Table 5.2. Occasionally, additional restrictions will be imposed on these ranges.

Now let $V$ be the set whose elements are the following (distinct) symbols, where the indices $\lambda$, $\pi$, $s$, $t$, $\mathbf{u}$, $\mathbf{v}$, $\mathbf{w}$ vary over their standard ranges:

$$x_\lambda, \qquad y_{\pi,s,\mathbf{u}}, \qquad z_{\pi,t,\mathbf{u}},$$
$$\hat{y}_{\pi,s,\mathbf{v},\mathbf{w}} \text{ whenever } |s| < p, \qquad \hat{z}_{\pi,t,\mathbf{v},\mathbf{w}} \text{ whenever } |t| < q.$$

| Symbol | Standard range |
|---|---|
| $\pi$, $\pi'$ | all 1-types |
| $\lambda$, $\lambda'$ | all invertible message-types |
| $s$, $s'$ | all bit-strings of length at most $p$ |
| $t$, $t'$ | all bit-strings of length at most $q$ |
| $\mathbf{u}$, $\mathbf{v}$, $\mathbf{w}$, $\mathbf{u}'$, $\mathbf{v}'$, $\mathbf{w}'$ | all vectors $\leq \mathbf{C}$ |

Table 5.2: Standard ranges of symbols used as indices

The symbols $\hat{y}_{\pi,s,\mathbf{v},\mathbf{w}}$ and $\hat{z}_{\pi,t,\mathbf{v},\mathbf{w}}$ are not defined when $|s| = p$ and $|t| = q$. The cardinality of $V$ is evidently bounded by an exponential function of $\|\phi\|$. We impose some arbitrary order on $V$, and refer to its elements as *variables*. If $U$ is a non-empty set of variables, enumerated, in the imposed order, as $\{u_1, \ldots, u_k\}$, let $\sum U$ denote the term $u_1 + \cdots + u_k$; if $U$ is the empty set, let $\sum U$ denote the (constant) term 0. In the sequel, we take a *constraint* to be an equation or inequality involving arithmetical terms over $V$, or a conditional statement formed from two such inequalities. A *solution* of a set of constraints over some numerical domain $\mathbb{D}$ is simply an assignment $\theta : V \to \mathbb{D}$ under which all the constraints in question evaluate (in the obvious way) to true. Using this apparatus, we proceed to construct, given the formula $\phi$ in (5.1), a set $\mathcal{E}$ of constraints. We prove below that $\mathcal{E}$ has a solution over $\mathbb{N}$ if and only if $\phi$ is finitely satisfiable.

To motivate this construction, suppose, for the moment, that $\mathfrak{A}$ is a finite model of $\phi$. We may then define the assignment $\theta : V \to \mathbb{N}$, which we may think of as the 'intended' assignment relative to $\mathfrak{A}$, as follows. If $\pi$ is a 1-type, write $A_\pi$ to denote the set of elements of $A$ having 1-type $\pi$ in $\mathfrak{A}$:

$$A_\pi = \{a \in A \mid \mathrm{tp}^{\mathfrak{A}}[a] = \pi\}.$$

Now, for any $\lambda$ in its standard range, let $\theta(x_\lambda)$ be the number of elements of $A$ sending any message of type $\lambda$ to some other element. For any $\pi$, $s$, $t$, $\mathbf{u}$ in their standard ranges, let $\theta(y_{\pi,s,\mathbf{u}})$ be the number of elements of $A_\pi$ having $s$-spectrum $\mathbf{u}$, and let $\theta(z_{\pi,t,\mathbf{u}})$ be the number of elements of of $A_\pi$ having $t$-tally $\mathbf{u}$. Finally, for any $\pi$, $s$, $t$, $\mathbf{v}$, $\mathbf{w}$ in their standard ranges, with $|s| < p$ and $|t| < q$, let $\theta(\hat{y}_{\pi,s,\mathbf{v},\mathbf{w}})$ be the number of elements of $A_\pi$ having $s0$-spectrum $\mathbf{v}$ and $s1$-spectrum $\mathbf{w}$, and let $\theta(\hat{z}_{\pi,t,\mathbf{v},\mathbf{w}})$ be the number of elements of $A_\pi$ having $t0$-tally $\mathbf{v}$ and $t1$-tally $\mathbf{w}$. This assignment is summarized in Table 5.3.

We construct $\mathcal{E}$ in three stages. First, let $\mathcal{E}_1$ be the following set of constraints involving the variables $V$, where $\pi$, $\mathbf{u}$, $\mathbf{v}$, $\mathbf{w}$ again vary over their standard

| Variable | Value in $\mathbb{N}$ under $\theta$ |
|---|---|
| $x_\lambda$ | $\lvert\{a \in A :$ there exists $b \in A \setminus \{a\}$ such that $\mathrm{tp}^{\mathfrak{A}}[a,b] = \lambda\}\rvert$ |
| $y_{\pi,s,\mathbf{u}}$ | $\lvert\{a \in A_\pi : \mathrm{sp}^{\mathfrak{A}}_s[a] = \mathbf{u}\}\rvert$ |
| $z_{\pi,t,\mathbf{u}}$ | $\lvert\{a \in A_\pi : \mathrm{tl}^{\mathfrak{A}}_t[a] = \mathbf{u}\}\rvert$ |
| $\hat{y}_{\pi,s,\mathbf{v},\mathbf{w}}$ | $\lvert\{a \in A_\pi : \mathrm{sp}^{\mathfrak{A}}_{s0}[a] = \mathbf{v}$ and $\mathrm{sp}^{\mathfrak{A}}_{s1}[a] = \mathbf{w}\}\rvert$, whenever $\lvert s\rvert < p$ |
| $\hat{z}_{\pi,t,\mathbf{v},\mathbf{w}}$ | $\lvert\{a \in A_\pi : \mathrm{tl}^{\mathfrak{A}}_{t0}[a] = \mathbf{v}$ and $\mathrm{tl}^{\mathfrak{A}}_{t1}[a] = \mathbf{w}\}\rvert$, whenever $\lvert t\rvert < q$ |

Table 5.3: The assignment $\theta : V \to \mathbb{N}$, assuming that $\mathfrak{A}$ is a finite model of $\phi$.

ranges, and $s$, $t$ vary over bit-strings such that $|s| < p$ and $|t| < q$:

$$z_{\pi,\epsilon,\mathbf{u}} = y_{\pi,\epsilon,\mathbf{C}-\mathbf{u}} \tag{5.7}$$

$$y_{\pi,s,\mathbf{u}} = \sum\{\hat{y}_{\pi,s,\mathbf{v}',\mathbf{w}'} \mid \mathbf{v}' + \mathbf{w}' = \mathbf{u}\} \tag{5.8}$$

$$z_{\pi,t,\mathbf{u}} = \sum\{\hat{z}_{\pi,t,\mathbf{v}',\mathbf{w}'} \mid \mathbf{v}' + \mathbf{w}' = \mathbf{u}\} \tag{5.9}$$

$$y_{\pi,s0,\mathbf{v}} = \sum\{\hat{y}_{\pi,s,\mathbf{v},\mathbf{w}'} \mid \mathbf{v} + \mathbf{w}' \leq \mathbf{C}\} \tag{5.10}$$

$$y_{\pi,s1,\mathbf{w}} = \sum\{\hat{y}_{\pi,s,\mathbf{v}',\mathbf{w}} \mid \mathbf{v}' + \mathbf{w} \leq \mathbf{C}\} \tag{5.11}$$

$$z_{\pi,t0,\mathbf{v}} = \sum\{\hat{z}_{\pi,t,\mathbf{v},\mathbf{w}'} \mid \mathbf{v} + \mathbf{w}' \leq \mathbf{C}\} \tag{5.12}$$

$$z_{\pi,t1,\mathbf{w}} = \sum\{\hat{z}_{\pi,t,\mathbf{v}',\mathbf{w}} \mid \mathbf{v}' + \mathbf{w} \leq \mathbf{C}\} \tag{5.13}$$

$$1 \leq \sum\{y_{\pi',\epsilon,\mathbf{u}'} \mid \pi' \text{ a 1-type}, \mathbf{u}' \leq \mathbf{C}\}. \tag{5.14}$$

**Lemma 36.** *Suppose $\mathfrak{A}$ is a finite model of $\phi$, and let the variables in $V$ take the values specified, relative to $\mathfrak{A}$, in Table 5.3. Then the constraints in $\mathcal{E}_1$ are all satisfied.*

*Proof.* The constraints (5.7)–(5.9) follow easily from the respective equations in Lemma 34. The constraints (5.10)–(5.13) are immediate. In the (single) constraint (5.14), the sum on the right-hand side evaluates to the cardinality of $A$, since every $a \in A$ has a unique 1-type and a unique $\epsilon$-spectrum. But $A$ is nonempty by definition. □

To define the next set of constraints, we require some additional terminology. Let $\tau$ be any 2-type. Since $\tau$ is a finite set of formulas with free variables $x$ and $y$, we may write $\bigwedge\tau$ to denote their conjunction. Referring again to the formula (5.1), we say that $\tau$ is *forbidden*, if the formula

$$\alpha(x) \wedge \alpha(y) \wedge \bigwedge_{1 \leq h \leq l} \Big( (e_h(x,y) \to \beta_h(x,y)) \wedge (e_h(y,x) \to \beta_h(y,x)) \Big) \wedge \bigwedge\tau \tag{5.15}$$

is unsatisfiable. Thus, if $\mathfrak{A} \models \phi$ and $a$, $b$ are distinct elements of $A$, then $\mathrm{tp}^{\mathfrak{A}}[a, b]$ cannot be forbidden. Since (5.15) is purely Boolean, we can evidently identify the forbidden 2-types in time bounded by an exponential function of $\|\phi\|$.

Now let $\mathcal{E}_2$ consist of the following constraints, where $\lambda$, $\pi$ vary over their standard ranges, $s$, $t$ vary over bit-strings such that $|s| = p$, $|t| = q$, and $\mathbf{u}$ varies over vectors such that $\mathbf{0} < \mathbf{u} \leq \mathbf{C}$:

$$y_{\pi,s,\mathbf{u}} = \sum\{x_{\lambda'} \mid \lambda' \in \Lambda_{\pi,s} \text{ and } \mathbf{C}_{\lambda'} = \mathbf{u}\} \tag{5.16}$$

$$x_{(\lambda^{-1})} = x_\lambda \tag{5.17}$$

$$x_\lambda = 0 \quad \text{whenever } \mathrm{tp}_1(\lambda) = \mathrm{tp}_2(\lambda) \tag{5.18}$$

$$x_\lambda = 0 \quad \text{whenever } \lambda \text{ is forbidden} \tag{5.19}$$

$$z_{\pi,t,\mathbf{u}} = 0 \quad \text{whenever } \mu_{\pi,t} \text{ is forbidden.} \tag{5.20}$$

$$z_{\pi,t,\mathbf{u}} = 0 \quad \text{whenever } \mathbf{u} \text{ is not a scalar multiple of } \mathbf{C}_\mu, \text{ for } \mu = \mu_{\pi,t} \tag{5.21}$$

**Lemma 37.** *Suppose $\mathfrak{A}$ is a finite, chromatic model of $\phi$, and let the variables in $V$ take the values specified, relative to $\mathfrak{A}$, in Table 5.3. Then the constraints in $\mathcal{E}_2$ are all satisfied.*

*Proof.* To see that the constraints (5.16) hold, fix $\pi$, $s$, $\mathbf{u}$ (with $|s| = p$ and $\mathbf{u} > \mathbf{0}$), and write

$$A_{\pi,s,\mathbf{u}} = \{a \in A \mid \mathrm{tp}^{\mathfrak{A}}[a] = \pi \text{ and } \mathrm{sp}^{\mathfrak{A}}_s[a] = \mathbf{u}\}.$$

In addition, for any invertible message-type $\lambda$, write

$$A_\lambda = \{a \in A \mid \text{ there exists } b \in A \setminus \{a\} \text{ such that } \mathrm{tp}^{\mathfrak{A}}[a, b] = \lambda\}.$$

From Table 5.3, we have $|A_{\pi,s,\mathbf{u}}| = \theta(y_{\pi,s,\mathbf{u}})$ and $|A_\lambda| = \theta(x_\lambda)$. Moreover, by Lemma 33, the sets $A_{\lambda'}$, for $\lambda'$ ranging over the elements of $\Lambda_{\pi,s}$, are pairwise disjoint. But Lemma 35 just states that

$$A_{\pi,s,\mathbf{u}} = \bigcup\{A_{\lambda'} \mid \lambda' \in \Lambda_{\pi,s} \text{ and } \mathbf{C}_{\lambda'} = \mathbf{u}\},$$

whence the relevant instance of the constraints (5.16) follows. To see that the constraints (5.17) hold, observe that, since $\mathfrak{A}$ is chromatic, $\theta(x_\lambda)$ is actually the total number of messages of (invertible) type $\lambda$ sent by elements of $\mathfrak{A}$, and similarly for $\lambda^{-1}$. But then $\theta(x_\lambda)$ and $\theta(x_{(\lambda^{-1})})$ are obviously equal. The constraints (5.18) are immediate given that $\mathfrak{A}$ is chromatic. The constraints (5.19) and (5.20) are immediate given that $\mathfrak{A} \models \phi$. Lastly, Equation (5.6) states that, for $|t| = q$, no element with 1-type $\pi$ can have a $t$-tally which is not a scalar multiple of $\mathbf{C}_\mu$, where $\mu$ is the sole element in $M_{\pi,t}$. The constraints (5.21) then follow. □

Let $\mathcal{E}_3$ consist of the following constraints, where $\pi$ varies over all 1-types, $t$ varies over bit-strings such that $|t| = q$, and $\mathbf{u}$ varies over vectors such that $\mathbf{0} < \mathbf{u} \leq \mathbf{C}$:

$$z_{\pi,t,\mathbf{u}} > 0 \quad \Rightarrow \quad \sum\{y_{\pi',\epsilon,\mathbf{u}'} \mid \pi' = \mathrm{tp}_2(\mu_{\pi,t}) \text{ and } \mathbf{u}' \leq \mathbf{C}\} > 0. \tag{5.22}$$

**Lemma 38.** *Suppose $\mathfrak{A}$ is a finite model of $\phi$, and let the variables in $V$ take the values specified, relative to $\mathfrak{A}$, in Table 5.3. Then the constraints in $\mathcal{E}_3$ are all satisfied.*

*Proof.* Fix $\pi$, $t$, $\mathbf{u}$ (with $|t| = q$ and $\mathbf{u} > \mathbf{0}$), so that $\mu_{\pi,t}$ is the sole element of $M_{\pi,t}$. If $\theta(z_{\pi,t,\mathbf{u}}) > 0$, some element has $t$-tally $\mathbf{u}$; and since $\mathbf{u} > \mathbf{0}$, $\mu_{\pi,t}$ must be a non-invertible message-type (i.e. not a silent 2-type), and moreover at least one message of that type must be sent in $\mathfrak{A}$. Therefore, $\mathfrak{A}$ contains at least one element whose 1-type is $\mathrm{tp}_2(\mu_{\pi,t})$. But the number of elements in $A$ whose 1-type is $\mathrm{tp}_2(\mu_{\pi,t})$ is $\sum\{\theta(y_{\pi',\epsilon,\mathbf{u}'}) \mid \pi' = \mathrm{tp}_2(\mu_{\pi,t}) \text{ and } \mathbf{u}' \leq \mathbf{C}\}$. □

Let $\mathcal{E} = \mathcal{E}_1 \cup \mathcal{E}_2 \cup \mathcal{E}_3$. Measuring the size $\|\mathcal{E}\|$ of $\mathcal{E}$ in the usual way, it is evident that $\|\mathcal{E}\|$ is bounded above by an exponential function of $\|\phi\|$. From the preceding lemmas:

**Lemma 39.** *Let $\phi$ and $\mathcal{E}$ be as above. If $\phi$ is finitely satisfiable, then $\mathcal{E}$ has a solution over $\mathbb{N}$.*

*Proof.* Suppose $\phi$ is finitely satisfiable. By Lemma 16, let $\mathfrak{A}$ be a finite, chromatic model of $\phi$. Assign to the variables in $V$ the values given in Table 5.3. Then apply Lemmas 36–38. □

When $\phi$ is finitely satisfiable, we may think of the variables $V$ as corresponding to configurational properties of elements in its finite, chromatic models. If the values of these variables are taken to record how often these configurational properties are realized in some such model, as prescribed in Table 5.3, then the constraints $\mathcal{E}$ will be satisfied.

## 5.7 Main result

We proceed to establish a converse of Lemma 39: given a solution of $\mathcal{E}$ over $\mathbb{N}$, there exists a finite model $\mathfrak{A}$ of $\phi$ such that that solution records how often certain configurational properties are realized in $\mathfrak{A}$, as specified in Table 5.3.

To reduce notational clutter, we use the variable names $x_\lambda$, $y_{\pi,s,\mathbf{u}}$, $z_{\pi,t,\mathbf{u}}$, $\hat{y}_{\pi,s,\mathbf{v},\mathbf{w}}$, $\hat{z}_{\pi,t,\mathbf{v},\mathbf{w}}$ to stand for the corresponding natural numbers in some solution of $\mathcal{E}$ (and similarly for terms involving these variables). Fix some 1-type $\pi$, and let $A_\pi$ be a set with exactly

$$\sum\{y_{\pi,\epsilon,\mathbf{u}'} \mid \mathbf{u}' \leq \mathbf{C}\}$$

elements (possibly zero). Think of $A_\pi$ as a set of elements which 'want' to have 1-type $\pi$ in some yet-to-be-built finite model $\mathfrak{A}$ of $\phi$.

Our first step will be to define, for any $s$, $t$ in their standard ranges, vector-valued functions $\mathbf{f}_{\pi,s}$ and $\mathbf{g}_{\pi,t}$ on $A_\pi$. For $a \in A_\pi$, think of $\mathbf{f}_{\pi,s}(a)$ as the $s$-spectrum which $a$ wants to have in $\mathfrak{A}$ (when $\mathfrak{A}$ is eventually built); and think of $\mathbf{g}_{\pi,t}(a)$ as the the $t$-tally which $a$ wants to have in $\mathfrak{A}$. Formally, the definitions of these functions simply depend on our solution of $\mathcal{E}$; informally, however, it helps

to keep in mind Table 5.3 when understanding the construction. In particular, if we want $y_{\pi,s,\mathbf{u}}$ to represent the number of elements having 1-type $\pi$ and $s$-spectrum $\mathbf{u}$ in $\mathfrak{A}$, we will need to ensure that exactly this number of elements $a \in A_\pi$ satisfy $\mathbf{f}_{\pi,s}(a) = \mathbf{u}$; and similarly for $t$-tallies. That is, we will need to ensure that, for all $\mathbf{u} \leq \mathbf{C}$,

$$|\mathbf{f}_{\pi,s}^{-1}(\mathbf{u})| = y_{\pi,s,\mathbf{u}} \tag{5.23}$$

$$|\mathbf{g}_{\pi,t}^{-1}(\mathbf{u})| = z_{\pi,t,\mathbf{u}}. \tag{5.24}$$

Furthermore, if $|s| < p$ and $|t| < q$, then, recalling Lemma 34, we will also need to ensure that, for all $a \in A_\pi$,

$$\mathbf{f}_{\pi,\epsilon}(a) + \mathbf{g}_{\pi,\epsilon}(a) = \mathbf{C} \tag{5.25}$$

$$\mathbf{f}_{\pi,s0}(a) + \mathbf{f}_{\pi,s1}(a) = \mathbf{f}_{\pi,s}(a) \tag{5.26}$$

$$\mathbf{g}_{\pi,t0}(a) + \mathbf{g}_{\pi,t1}(a) = \mathbf{g}_{\pi,t}(a). \tag{5.27}$$

The following rather technical lemma simply guarantees that these requirements can be satisfied.

**Lemma 40.** *Suppose $x_\lambda$, $y_{\pi,s,\mathbf{u}}$, $z_{\pi,t,\mathbf{u}}$, $\hat{y}_{\pi,s,\mathbf{v},\mathbf{w}}$, $\hat{z}_{\pi,t,\mathbf{v},\mathbf{w}}$ (with indices having the appropriate ranges) are natural numbers satisfying the constraints $\mathcal{E}$ given above. Fix any 1-type $\pi$, and let $A_\pi$ be a set of cardinality $\sum\{y_{\pi,\epsilon,\mathbf{u}'} \mid \mathbf{u}' \leq \mathbf{C}\}$. Then there exists a system of functions on $A_\pi$*

$$\mathbf{f}_{\pi,s} : A_\pi \to \{\mathbf{u} \mid \mathbf{u} \leq \mathbf{C}\} \qquad \mathbf{g}_{\pi,t} : A_\pi \to \{\mathbf{u} \mid \mathbf{u} \leq \mathbf{C}\},$$

*where the indices $s$ and $t$ vary over their standard ranges, satisfying the following conditions: (i) Equations* (5.23) *and* (5.24) *hold for all vectors $\mathbf{u} \leq \mathbf{C}$; (ii) if $|s| < p$ and $|t| < q$, then Equations* (5.25)–(5.27) *hold for all $a \in A_\pi$.*

*Proof.* Decompose the set $A_\pi$ into pairwise disjoint (possibly empty) sets $A_\mathbf{u}$ such that $|A_\mathbf{u}| = y_{\pi,\epsilon,\mathbf{u}}$, where the index $\mathbf{u}$ varies over all vectors $\leq \mathbf{C}$. This is possible by the cardinality of $A_\pi$. For all $\mathbf{u} \leq \mathbf{C}$, and all $a \in A_\mathbf{u}$, set

$$\mathbf{f}_{\pi,\epsilon}(a) = \mathbf{u} \qquad \mathbf{g}_{\pi,\epsilon}(a) = \mathbf{C} - \mathbf{u}.$$

This assignment evidently satisfies (5.23) for $s = \epsilon$; and by the constraints (5.7), it also satisfies (5.24) for $t = \epsilon$. Moreover, it is immediate that, for all $a \in A_\pi$, Equation (5.25) holds as required.

We now construct the functions $\mathbf{f}_{\pi,s}$, where $0 < |s| \leq p$, by induction on $s$. Assume that, for some $s$ ($0 \leq |s| < p$), $\mathbf{f}_{\pi,s}$ has been defined and satisfies (5.23). For every vector $\mathbf{u} \leq \mathbf{C}$, decompose $\mathbf{f}_{\pi,s}^{-1}(\mathbf{u})$ into pairwise disjoint (possibly empty) sets $A_{\mathbf{v},\mathbf{w}}$ such that $|A_{\mathbf{v},\mathbf{w}}| = \hat{y}_{\pi,s,\mathbf{v},\mathbf{w}}$, where the indices $\mathbf{v}$, $\mathbf{w}$ vary over all vectors satisfying $\mathbf{v} + \mathbf{w} = \mathbf{u}$. This is possible by the constraints (5.8) together with the assumption that $\mathbf{f}_{\pi,s}$ satisfies (5.23). Having thus decomposed the sets $\mathbf{f}_{\pi,s}^{-1}(\mathbf{u})$ (for all $\mathbf{u} \leq \mathbf{C}$), we see that, for any $a \in A_\pi$, there is precisely one (ordered) pair of vectors $\mathbf{v}$, $\mathbf{w}$ such that $a \in A_{\mathbf{v},\mathbf{w}}$; hence we may set

$$\mathbf{f}_{\pi,s0}(a) = \mathbf{v} \qquad \mathbf{f}_{\pi,s1}(a) = \mathbf{w}.$$

This defines the functions $\mathbf{f}_{\pi,s0}$ and $\mathbf{f}_{\pi,s1}$. It is immediate that, for all $a \in A_\pi$, Equation (5.26) holds as required.

To see that $\mathbf{f}_{\pi,s0}$ and $\mathbf{f}_{\pi,s1}$ both satisfy Equation (5.23), note that $\mathbf{f}_{\pi,s0}(a) = \mathbf{v}$ if and only if, for some vector $\mathbf{w}'$ such that $\mathbf{v} + \mathbf{w}' \leq \mathbf{C}$, $a \in A_{\mathbf{v},\mathbf{w}'}$. Similarly, $\mathbf{f}_{\pi,s1}(a) = \mathbf{w}$ if and only if, for some vector $\mathbf{v}'$ such that $\mathbf{v}' + \mathbf{w} \leq \mathbf{C}$, $a \in A_{\mathbf{v}',\mathbf{w}}$. That is,

$$\begin{aligned} \mathbf{f}^{-1}_{\pi,s0}(\mathbf{v}) &= \bigcup\{A_{\mathbf{v},\mathbf{w}'} \mid \mathbf{v} + \mathbf{w}' \leq \mathbf{C}\} \\ \mathbf{f}^{-1}_{\pi,s1}(\mathbf{w}) &= \bigcup\{A_{\mathbf{v}',\mathbf{w}} \mid \mathbf{v}' + \mathbf{w} \leq \mathbf{C}\}, \end{aligned}$$

with the collections of sets on the respective right-hand sides being pairwise disjoint. By the constraints (5.10)–(5.11), together with the fact that $|A_{\mathbf{v},\mathbf{w}}| = \hat{y}_{\pi,s,\mathbf{v},\mathbf{w}}$ for all $\mathbf{v}, \mathbf{w}$, we have:

$$\begin{aligned} |\mathbf{f}^{-1}_{\pi,s0}(\mathbf{v})| &= y_{\pi,s0,\mathbf{v}} \\ |\mathbf{f}^{-1}_{\pi,s1}(\mathbf{w})| &= y_{\pi,s1,\mathbf{w}}, \end{aligned}$$

which establishes (5.23) for the functions $\mathbf{f}_{\pi,s0}$ and $\mathbf{f}_{\pi,s1}$. This completes the induction. The construction of the functions $\mathbf{g}_{\pi,t}$ proceeds completely analogously, using the constraints (5.9), (5.12) and (5.13). □

**Lemma 41.** *Let the functions $\mathbf{f}_{\pi,s}$ and $\mathbf{g}_{\pi,t}$ be constructed as in Lemma 40. Then, for all $a \in A_\pi$, we have*

$$\sum\{\mathbf{f}_{\pi,s'}(a) : |s'| = p\} + \sum\{\mathbf{g}_{\pi,t'}(a) : |t'| = q\} = \mathbf{C}.$$

*Proof.* We prove the stronger result that, for all $a \in A_\pi$, $j$ $(0 \leq j \leq p)$ and $k$ $(0 \leq k \leq q)$,

$$\sum\{\mathbf{f}_{\pi,s'}(a) : |s'| = j\} + \sum\{\mathbf{g}_{\pi,t'}(a) : |t'| = k\} = \mathbf{C}, \tag{5.28}$$

using a double induction on $j$ and $k$. If $j = k = 0$, then the left-hand side of (5.28) is simply $\mathbf{f}_{\pi,\epsilon}(a) + \mathbf{g}_{\pi,\epsilon}(a)$, which is equal to $\mathbf{C}$ by (5.25). Suppose now that the result holds for the pair $j$, $k$, with $j < p$. Then

$$\begin{aligned} & \sum\{\mathbf{f}_{\pi,s'}(a) : |s'| = (j+1)\} + \sum\{\mathbf{g}_{\pi,t'}(a) : |t'| = k\} \\ = \; & \sum\{\mathbf{f}_{\pi,s'0}(a) + \mathbf{f}_{\pi,s'1}(a) : |s'| = j\} + \sum\{\mathbf{g}_{\pi,t'}(a) : |t'| = k\} \\ = \; & \sum\{\mathbf{f}_{\pi,s'}(a) : |s'| = j\} + \sum\{\mathbf{g}_{\pi,t'}(a) : |t'| = k\} \qquad \text{by (5.26)} \\ = \; & \mathbf{C} \qquad \text{by inductive hypothesis.} \end{aligned}$$

This establishes the result for the pair $j + 1, k$. An analogous argument using (5.27) applies when $k < m$, completing the induction. □

Before we come to the promised converse of Lemma 39, we remark on the (exponentially many) choices made during the construction of the various functions $\mathbf{f}_{\pi,s}$ and $\mathbf{g}_{\pi,t}$ in the proof of Lemma 40—specifically, in the decomposition of certain sets into collections of subsets. A given solution of $\mathcal{E}$ ensures that a system of functions $\mathbf{f}_{\pi,s}$ and $\mathbf{g}_{\pi,t}$ exists, subject to the given conditions; but it by no means determines them.

**Lemma 42.** *Let $\mathcal{E}$ be as above and $k$ a positive integer. If $\mathcal{E}$ has a solution over $\mathbb{N}$, then it has a solution over $\mathbb{N}$ in which all positive values are at least $k$.*

*Proof.* Suppose $\mathcal{E}$ has a solution $\theta : V \to \mathbb{N}$. Now define $\theta' : V \to \mathbb{N}$ by $\theta'(v) = k\theta(v)$. By inspection, $\theta'$ is a solution of $\mathcal{E}$. □

Model-theoretically, Lemma 42 is simply a reflection of Lemma 31.

**Lemma 43.** *Let $\phi$ and $\mathcal{E}$ be as above. If $\mathcal{E}$ has a solution over $\mathbb{N}$, then $\phi$ is finitely satisfiable.*

*Proof.* By Lemma 42, we may assume that $\mathcal{E}$ has a solution in which all positive values are greater than or equal to $3mC$. Again, we use the variable names $x_\lambda$, $y_{\pi,s,\mathbf{u}}$, $z_{\pi,t,\mathbf{u}}$, $\hat{y}_{\pi,s,\mathbf{v},\mathbf{w}}$, $\hat{z}_{\pi,t,\mathbf{v},\mathbf{w}}$ to stand for the corresponding values in this solution. Our task is to construct a model $\mathfrak{A}$ of $\phi$.

For each 1-type $\pi$, let $A_\pi$ be a set of cardinality $\sum\{y_{\pi,\epsilon,\mathbf{u}} \mid \mathbf{u} \leq \mathbf{C}\}$, with the $A_\pi$ pairwise disjoint; and let $A = \bigcup\{A_\pi \mid \pi \text{ a 1-type}\}$. Think of $A_\pi$ as the set of elements of $A$ which 'want' to have 1-type $\pi$. By the constraint (5.14), $A \neq \emptyset$. For every 1-type $\pi$, let the functions $\mathbf{f}_{\pi,s}$ and $\mathbf{g}_{\pi,t}$ on $A_\pi$ be constructed as in Lemma 40; we are interested only in those $\mathbf{f}_{\pi,s}$ and $\mathbf{g}_{\pi,t}$ where $|s| = p$, and $|t| = q$. Think of $\mathbf{f}_{\pi,s}(a)$ as the $s$-spectrum which $a$ wants to have, and of $\mathbf{g}_{\pi,t}(a)$ as the $t$-tally which $a$ wants to have. Finally, consider any set $\mathbf{f}^{-1}_{\pi,s}(\mathbf{u})$, where $\mathbf{0} < \mathbf{u} \leq \mathbf{C}$ and $|s| = p$. Using the constraints (5.16) and Equation (5.23), we can decompose $\mathbf{f}^{-1}_{\pi,s}(\mathbf{u})$ into pairwise disjoint (possibly empty) sets $A_\lambda$ with $|A_\lambda| = x_\lambda$, where $\lambda$ varies over the set of invertible message-types such that $\lambda \in \Lambda_{\pi,s}$ and $\mathbf{C}_\lambda = \mathbf{u}$. It follows that, if $a \in A_\lambda$, with $\lambda \in \Lambda_{\pi,s}$, then $\mathbf{C}_\lambda = \mathbf{f}_{\pi,s}(a)$. Think of $A_\lambda$ as the set of elements of $A_\pi$ which want to send a single message of (invertible) type $\lambda$.

Before proceeding, we pause to consider the construction just described in respect of any of the sets $A_\pi$. Fixing, for the moment, some bit-string $s$ with $|s| = p$, we see that $A_\pi$ is decomposed into the pairwise disjoint sets $\mathbf{f}^{-1}_{\pi,s}(\mathbf{u})$ (as $\mathbf{u}$ varies over vectors $\leq \mathbf{C}$), and that each of the sets $\mathbf{f}^{-1}_{\pi,s}(\mathbf{u})$, where $\mathbf{0} < \mathbf{u} \leq \mathbf{C}$, is further decomposed into the pairwise disjoint subsets $A_\lambda$ (as $\lambda$ varies over the elements of $\Lambda_{\pi,s}$ such that $\mathbf{C}_\lambda = \mathbf{u}$). Note that the set $\mathbf{f}^{-1}_{\pi,s}(\mathbf{0})$ is not subject to this further stage of decomposition. This process is performed for *every* bit string $s$ with $|s| = p$, so that different values of $s$ lead to independent—and possibly overlapping—decompositions, as illustrated in Fig. 5.1. Likewise, for every bit-string $t$ with $|t| = q$, $A_\pi$ is decomposed into the pairwise disjoint sets $\mathbf{g}^{-1}_{\pi,t}(\mathbf{u})$ (as $\mathbf{u}$ varies over vectors $\leq \mathbf{C}$). Again, decompositions corresponding to different values of $t$ should be thought of as independent of each other.

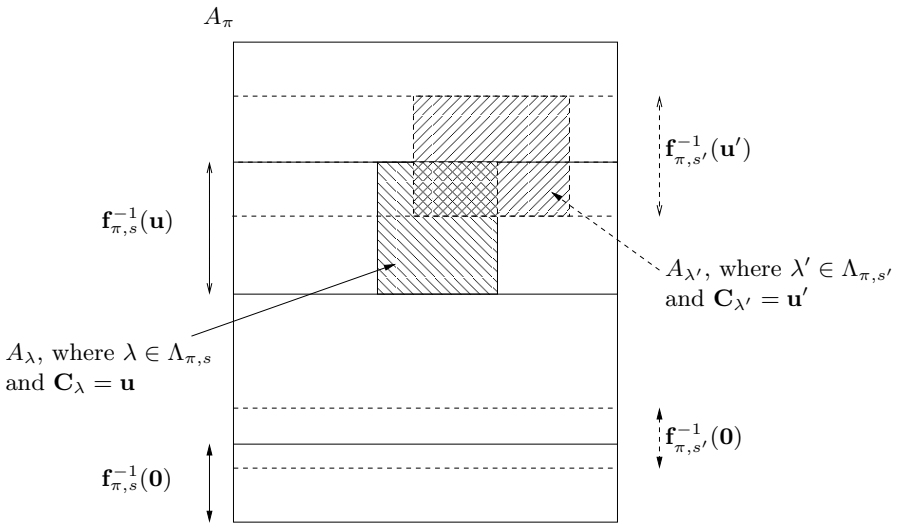


Figure 5.1: The decompositions of $A_\pi$ for the strings $s$ and $s'$.

We now proceed to construct, for every $a \in A$, a 'mosaic piece'; these pieces will be assembled into the desired structure $\mathfrak{A}$. Formally, a mosaic piece is a finite multiset of message-types (the reader is asked to excuse the mixed metaphor); informally, we may think of a mosaic piece as a finite collection of 'messages' sent by $a$, each of which is labelled with some (invertible or non-invertible) message-type (Fig. 5.2). Recall from Section 5.4 that, if $\pi$ is any 1-type, then $\mu_{\pi,0}, \ldots, \mu_{\pi,R-1}$ is an enumeration of the non-invertible message-types $\mu$ such that $\mathrm{tp}_1(\mu) = \pi$. Fix $a \in A$, and let $\pi$ be the unique 1-type such that $a \in A_\pi$. The messages in the mosaic piece corresponding to $a$ shall be as follows. (i) For every bit-string $s$ such that $|s| = p$, if $\mathbf{f}_{\pi,s}(a) > \mathbf{0}$, let $\lambda_{a,s}$ be the invertible message-type $\lambda \in \Lambda_{\pi,s}$ such that $a \in A_\lambda$ (hence $\mathbf{C}_\lambda = \mathbf{f}_{\pi,s}(a)$), and let the mosaic piece corresponding to $a$ contain a single message labelled $\lambda_{a,s}$. Note that, if $\mathbf{f}_{\pi,s}(a) > \mathbf{0}$, then $\lambda_{a,s}$ exists and is unique by the construction of the sets $A_\lambda$. (ii) For every bit string $t$ such that $|t| = q$, if $\mu = \mu_{\pi,t}$ is a non-invertible message-type and $\mathbf{C}_\mu > \mathbf{0}$, let $n_{a,t}$ be the unique natural number $n$ such that $\mathbf{g}_{\pi,t}(a) = n\mathbf{C}_\mu$, and let the mosaic piece corresponding to $a$ contain $n_{a,t}$ distinct messages labelled $\mu_{\pi,t}$. Note that, if $\mathbf{g}_{\pi,t}(a) = \mathbf{0}$, then $n_{a,t} = 0$; on the other hand, if $\mathbf{g}_{\pi,t}(a) > \mathbf{0}$, then $n_{a,t}$ exists by the constraints (5.21) and Equation (5.24). The resulting mosaic piece is depicted in Fig. 5.2, where, for readability, we have replaced any bit-strings by the integers they conventionally denote.

For all $a \in A$ and all $i$ $(1 \leq i \leq m)$, let $C_{a,i}$ $(1 \leq i \leq m)$ be the number of messages in the mosaic piece for $a$ (as just constructed) having any label $\nu$ for which $f_i(x, y) \in \nu$, and furthermore let $\mathbf{C}_a$ be the vector $(C_{a,1}, \ldots, C_{a,m})$. By

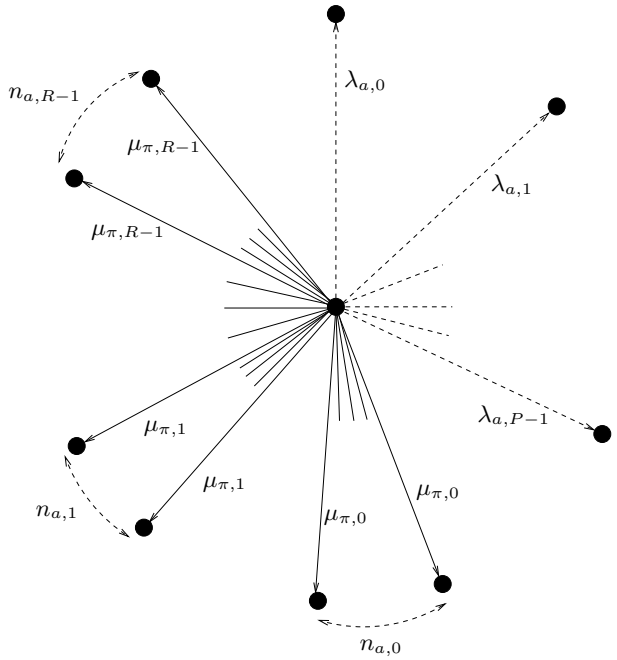


Figure 5.2: The mosaic piece corresponding to $a \in A_\pi$. For each $j$ $(0 \leq j < P)$, $a$ may or may not send a message labelled $\lambda_{a,j}$ (hence the dotted lines); if it does, then $\lambda_{a,j} \in \Lambda_{\pi,j}$. For each $k$ $(0 \leq k < R)$, $a$ sends $n_{a,k}$ messages labelled $\mu_{\pi,k}$; but the numbers $n_{a,k}$ can be zero.

inspection of Fig. 5.2,

$$\mathbf{C}_a = \sum\{\mathbf{f}_{\pi,s'}(a) : |s'| = p\} + \sum\{\mathbf{g}_{\pi,t'}(a) : |t'| = q\},$$

and so, by Lemma 41,

$$\mathbf{C}_a = \mathbf{C}. \tag{5.29}$$

If the mosaic piece corresponding to $a$ contains a message labelled with some message-type $\nu$, we will say that *$a$ sends a message of type $\nu$*. Equation (5.29) is evidently a necessary condition for mosaic pieces that are going to be assembled into a model $\mathfrak{A}$ of $\phi$. We build $\mathfrak{A}$ in four steps as follows.

**Step 1** (Fixing the 1-types): For all 1-types $\pi$ and all $a \in A_\pi$, set $\mathrm{tp}^{\mathfrak{A}}[a] = \pi$. Since the $A_\pi$ are pairwise disjoint, no clashes arise.

**Step 2** (Fixing the invertible message-types): Let $\lambda$ be any invertible message-type. By construction, exactly $|A_\lambda| = x_\lambda$ elements of $A$ send some message labelled with $\lambda$, and each of those elements sends exactly one such message. Hence, the number of messages labelled with $\lambda$ (over all $a \in A$) is $x_\lambda$; likewise, the number of messages labelled with $\lambda^{-1}$ is $x_{\lambda^{-1}}$. By the constraints (5.17),

we may put the $\lambda$-labelled messages and the $\lambda^{-1}$-labelled messages in 1–1 correspondence. If $a \in A$ sends a $\lambda$-labelled message, let $b \in A$ send the corresponding $\lambda^{-1}$-labelled message, and set $\mathrm{tp}^{\mathfrak{A}}[a,b] = \lambda$. For this assignment to make sense, we need to check that $a$ and $b$ are distinct. But, by construction, we must have $x_\lambda > 0$, whence, by the constraints (5.18), $\mathrm{tp}_1(\lambda) \neq \mathrm{tp}_2(\lambda)$, so that $A_{\mathrm{tp}_1(\lambda)}$ and $A_{\mathrm{tp}_2(\lambda)}$ are disjoint sets containing $a$ and $b$, respectively. Thus, the assignment $\mathrm{tp}^{\mathfrak{A}}[a,b] = \lambda$ makes sense, and does not clash with the 1-type assignments in Step 1. We can think of the element $b$ as 'receiving' the message sent by $a$ (and vice versa). Moreover, by construction, for every 1-type $\pi'$, $a$ sends at most one message labelled with an invertible message-type $\lambda'$ such that $\mathrm{tp}_2(\lambda') = \pi'$. Therefore, there is no chance that these assignments clash with each other. Note, incidentally, that this method of avoiding clashes means that $\mathfrak{A}$ will turn out to be chromatic.

**Step 3** (Fixing the non-invertible message-types): As a preliminary, for every 1-type $\pi$, we decompose $A_\pi$ into three pairwise disjoint (possibly empty) sets $A_{\pi,0}$, $A_{\pi,1}$ and $A_{\pi,2}$ satisfying the condition that, if $|A_\pi| \geq 3mC$, then $|A_{\pi,j}| \geq mC$ for all $j$ $(0 \leq j \leq 2)$. Now let $\mu$ be any non-invertible message-type, let $\pi = \mathrm{tp}_1(\mu)$, and let $\rho = \mathrm{tp}_2(\mu)$. (Note that $\pi$ and $\rho$ may be identical.) Let $t$ be the bit-string of length $q$ such that $\mu = \mu_{\pi,t}$, and suppose some element $a$ sends $n_{a,t} > 0$ messages labelled $\mu$. It follows that $a \in A_\pi$, and also that there is a vector $\mathbf{u} > \mathbf{0}$ such that $\mathbf{g}_{\pi,t}(a) = \mathbf{u}$, and hence such that $\mathbf{g}^{-1}_{\pi,t}(\mathbf{u})$ is non-empty. By Equation (5.24), $z_{\pi,t,\mathbf{u}}$ is positive, whence, by the constraints (5.22), $\sum\{y_{\rho,\epsilon,\mathbf{u}'} \mid \mathbf{u}' \leq \mathbf{C}\}$ is also positive, and therefore, by our choice of solution, greater than or equal to $3mC$. But recall that, since $\rho$ is a 1-type, $|A_\rho| = \sum\{y_{\rho,\epsilon,\mathbf{u}'} \mid \mathbf{u}' \leq \mathbf{C}\}$, so that each of the sets $A_{\rho,0}$, $A_{\rho,1}$ and $A_{\rho,2}$ contains at least $mC$ elements. Since $a \in A_\pi$, let $j$ $(0 \leq j \leq 2)$ be such that $a \in A_{\pi,j}$, let $k = j + 1 \pmod 3$, and select $n_{a,t}$ elements $b$ from $A_{\rho,k}$ which have not yet been chosen to receive any other messages (invertible or non-invertible) sent by $a$. Since the total number of messages sent by $a$ is certainly at most $mC$, we never run out of choices. For each of these elements $b$, set $\mathrm{tp}^{\mathfrak{A}}[a,b] = \mu$. Since $\pi = \mathrm{tp}_1(\mu)$ and $\rho = \mathrm{tp}_2(\mu)$, these assignments cannot clash with those made in Step 1, and by construction, they cannot clash with assignments corresponding to other messages sent by $a$. We need only check that they cannot clash with assignments corresponding to messages sent by $b$. Specifically, we must ensure that, if $\mathrm{tp}^{\mathfrak{A}}[a,b] = \mu$ is assigned as just described, it is not possible for $a$ to be chosen to receive a $\mu'$-labelled message sent by $b$, where $\mu'$ is some non-invertible message-type. But any $\mu'$-labelled message sent by $b \in A_{\rho,k}$, with $\mathrm{tp}_2(\mu') = \pi$, could only be sent to an element in $A_{\pi,j'}$, where $j' = k + 1 \pmod 3$; and by assumption, $A_{\pi,j}$ and $A_{\pi,j'}$ are disjoint, (Fig. 5.3). Observe that this conclusion follows even if $\pi = \rho$.

**Step 4** (Fixing the remaining 2-types): Recall that a guard-atom is any atom $p(x,y)$ or $p(y,x)$, where $p$ is a binary predicate. If $\mathrm{tp}^{\mathfrak{A}}[a,b]$ has not been defined, set it to be the 2-type

$$\pi \cup \rho[y/x] \cup \{\neg\gamma \mid \ \gamma \text{ is a guard-atom not involving } \approx\},$$

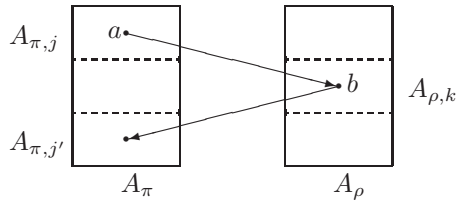


Figure 5.3: Fixing the non-invertible message-types.

where $\pi = \mathrm{tp}^{\mathfrak{A}}[a]$, $\rho = \mathrm{tp}^{\mathfrak{A}}[b]$, and $\rho[y/x]$ is the result of replacing $x$ by $y$ in $\rho$. Note that neither this 2-type nor its inverse is a message-type. Note also that, since the integers $C_1, \ldots, C_m$ and $m$ are by assumption all positive, $a$ and $b$ certainly send some messages, so that the constraints (5.19) and (5.20) ensure that both $\alpha \wedge \bigwedge \pi$ and $\alpha \wedge \bigwedge \rho$ are satisfiable.

This completes the definition of $\mathfrak{A}$; it remains to show that $\mathfrak{A} \models \phi$. Referring to (5.1), we consider first the conjuncts:

$$\forall x \alpha \wedge \bigwedge_{1 \leq h \leq l} \forall x \forall y (e_h(x,y) \rightarrow (\beta_h \vee x \approx y)).$$

We see from the constraints (5.19) and (5.20) that no 2-type assignment in Steps 2 and 3 violates these conjuncts. And it is obvious that no assignment in Step 4 does so. (This is where we use the guardedness of $\phi$, of course.) Finally, we consider the conjuncts

$$\bigwedge_{1 \leq i \leq m} \forall x \exists_{=C_i} y (f_i(x,y) \wedge x \not\approx y).$$

To see that these conjuncts are all satisfied, it suffices to note Equation (5.29) and the fact that none of the 2-types assigned in Step 4 is a message-type. □

The constraints $\mathcal{E}$ all have the forms

$$\begin{array}{rcl} x_1 + \cdots + x_n & = & x \\ x_1 + \cdots + x_n & \geq & 1 \\ x & = & 0 \\ x > 0 \;\Rightarrow\; x_1 + \cdots + x_n & > & 0, \end{array} \tag{5.30}$$

where $n > 0$, $x, x_1, \ldots, x_n$ are variables. We now investigate the problem of determining whether $\mathcal{E}$ has a solution. The following lemma essentially repeats Lutz, Sattler and Tendera [22], Proposition 11. (Those authors in turn credit Calvanese [5].) We give a proof for convenience. An *integer programming problem* is a system of linear equalities and inequalities interpreted over $\mathbb{Z}$. A *linear programming problem* is a system of linear equalities and inequalities interpreted over $\mathbb{Q}$.

**Lemma 44.** *Let $\phi$ and $\mathcal{E}$ be as above. An algorithm exists to determine whether $\mathcal{E}$ has a solution over $\mathbb{N}$ in time bounded by a polynomial function of $\|\mathcal{E}\|$, and hence by an exponential function of $\|\phi\|$.*

*Proof.* Evidently, $\mathcal{E}$ can be regarded as a very large disjunction of integer programming problems, each one of which has size bounded by $\|\mathcal{E}\|$. By a well-known theorem (Borosh and Treybig [4]), there is a monotonic function $h$, computable in polynomial time, such that, if an integer programming problem of size $n$ has a solution, then it has a solution in which every value is bounded by $h(n) > 0$. Hence, $\mathcal{E}$ has a solution over $\mathbb{N}$ if and only if it has a solution over $\mathbb{N}$ in which every value is bounded by $H = h(\|\mathcal{E}\|)$.

Now consider the integer programming problem $\mathcal{E}_H$ defined by replacing every constraint of the form $x > 0 \Rightarrow x_1 + \cdots + x_n > 0$ in $\mathcal{E}$ by the corresponding inequalities

$$\begin{aligned} Hy &\geq x \\ x_1 + \cdots + x_n &\geq y, \end{aligned}$$

where $y$ is a new variable. Every solution of $\mathcal{E}_H$ over $\mathbb{N}$ is clearly a solution of $\mathcal{E}$. Conversely, suppose $\theta : V \to \mathbb{N}$ is a solution of $\mathcal{E}$ in which all values are bounded by $H$. Let $y$ be any of the new variables of $\mathcal{E}_H$, introduced to eliminate the constraint $x > 0 \Rightarrow x_1 + \cdots + x_n > 0$; and extend $\theta$ to give a value to $y$ as follows:

$$\theta(y) = \begin{cases} 0 & \text{if } \theta(x) = 0 \\ 1 & \text{otherwise.} \end{cases}$$

It is routine to check that extending $\theta$ in this way for all the new variables $y$ in $\mathcal{E}_H$ yields a solution of $\mathcal{E}_H$. Hence $\mathcal{E}$ can be transformed, in time bounded by a polynomial function of $\|\mathcal{E}\|$, into a constraint set $\mathcal{E}_H$, in which all constraints are of the forms

$$\begin{aligned} x_1 + \cdots + x_n &= x & \qquad x &= 0 \\ x_1 + \cdots + x_n &\geq 1 & \qquad Hx_1 &\geq x_2 \\ x_1 + \cdots + x_n &\geq x, \end{aligned}$$

such that $\mathcal{E}$ has a solution (over $\mathbb{N}$) if and only if $\mathcal{E}_H$ has. It is obvious that, if $\mathcal{E}_H$ has a solution over the non-negative rationals, then it has a solution over $\mathbb{N}$ as well. (Simply multiply by the product of all the denominators.) Hence, we can equivalently regard $\mathcal{E}_H$ as a linear programming problem. But linear programming is in PTIME, by Khachiyan's theorem [17]. □

**Theorem 16.** *The finite satisfiability problem for $\mathcal{GC}^2$ is in* ExpTime.

*Proof.* Lemmas 32, 39, 43 and 44. □

## 5.8 The satisfiability problem

In the previous section, we considered only the *finite* satisfiability problem for $\mathcal{GC}^2$. However, the technique employed easily yields a corresponding result on the satisfiability problem for $\mathcal{GC}^2$.

**Notation 6.** Let $\mathbb{N}^*$ denote the set $\mathbb{N} \cup \{\aleph_0\}$. We extend the ordering $>$ and the arithmetic operations $+$ and $\cdot$ from $\mathbb{N}$ to $\mathbb{N}^*$ in the obvious way. Specifically, we define $\aleph_0 > n$ for all $n \in \mathbb{N}$; we define $\aleph_0 + \aleph_0 = \aleph_0 \cdot \aleph_0 = \aleph_0$ and $0 \cdot \aleph_0 = \aleph_0 \cdot 0 = 0$; we define $n + \aleph_0 = \aleph_0 + n = \aleph_0$ for all $n \in \mathbb{N}$; and we define $n \cdot \aleph_0 = \aleph_0 \cdot n = \aleph_0$ for all $n \in \mathbb{N}$ such that $n > 0$. Under this extension, $>$ remains a total order, and $+$, $\cdot$ remain associative and commutative.

Using the arithmetic in Notation 6, consider again the constraints $\mathcal{E}$ given in (5.7)–(5.14), (5.16)–(5.21) and (5.22), but now with the variables ranging over the whole of $\mathbb{N}^*$.

**Lemma 45.** *Let $\phi$ and $\mathcal{E}$ be as above. Then $\phi$ is satisfiable if and only if $\mathcal{E}$ has a solution over $\mathbb{N}^*$.*

*Proof.* If $\phi$ is satisfiable, then it has a model $\mathfrak{A}$ which is finite or countably infinite. Now assign to the variables in $V$ values in $\mathbb{N}^*$ as directed by Table 5.3. The reasoning of Lemmas 36–39 then goes through, with the obvious changes of formulation, exactly as in the finite case. For the converse, proceed exactly as for Lemmas 40–43. □

**Lemma 46.** *The set of constraints $\mathcal{E}$ has a solution over $\mathbb{N}^*$ if and only if it has a solution over $\{0, \aleph_0\}$.*

*Proof.* Suppose $\mathcal{E}$ has a solution $\theta : V \to \mathbb{N}^*$. Now define $\theta' : V \to \{0, \aleph_0\}$ by $\theta'(v) = \aleph_0 \theta(v)$. By inspection, $\theta'$ is a solution of $\mathcal{E}$. □

Model-theoretically, Lemma 46 is simply a reflection of Lemma 31.

Since the domain $\{0, \aleph_0\}$ has only 2-elements, variables interpreted over it are essentially Boolean. If $x \in V$, let us write $X$ for the corresponding statement $x = 0$, so that the constraints $\mathcal{E}$ are viewed as formulas of propositional logic. For example, a constraint of the form

$$x_1 + \cdots + x_n = x$$

becomes the set of propositional logic formulas

$$\{X_1 \wedge \cdots \wedge X_n \to X\} \cup \{X \to X_i \mid 1 \leq i \leq n\};$$

a constraint of the form

$$x_1 + \cdots + x_n \geq 1$$

becomes the propositional logic formula

$$X_1 \wedge \cdots \wedge X_n \to \bot;$$

and a constraint of the form

$$x > 0 \quad \Rightarrow \quad x_1 + \cdots + x_n > 0$$

becomes the propositional logic formula

$$X_1 \wedge \cdots \wedge X_n \to X.$$

A quick check reveals that all of the resulting formulas are Horn-clauses. This immediately yields:

**Theorem 17.** *The satisfiability problem for* $\mathcal{GC}^2$ *is in* ExpTime.

## 5.9 Bibliographic notes

The language $\mathcal{M}$ of modal logic can be found in Lewis and Langford [21], and originally formed the basis of a collection of axiomatically characterized logics; the alternative, semantic, characterization of these logics in terms of frame classes—which forms the basis of the approach adopted in these notes—is due to Kripke [18]. The complexity-theoretic analysis of these logics was carried out relatively early. Theorem 8 may be regarded as folklore; Theorem 9 is due to Ladner [20].

The special case of Theorem 11 where $\mathcal{F} = \emptyset$ is due to Tobies [35]; the remaining cases can be easily proved using the same technique. Theorems 10 and 12 use rather different methods, and are due to Kazakov and Pratt-Hartmann [16].

The guarded fragment $\mathcal{G}$ was originally introduced by Andréka *et al.* [1], as a means of generalizing modal logic (under the relational semantics), and, at the same time, exhibiting its connection to earlier work in algebraic logic. Theorems 13, 14 and 15 are due to Grädel [11]. Theorem 16 is due to Pratt-Hartmann [31]. Theorem 17 was first proved—using a very different technique—in Kazakov [15]. Kazakov's strategy is to show that satisfiability in $\mathcal{GC}^2$ can be reduced in polynomial time to satisfiability in the 3-variable guarded fragment; the result then follows from Theorem 14. We note in this regard that Kazakov's reduction is not conservative, and therefore yields no proof of Theorem 16.

# Bibliography


[1] Hajnal Andréka, Johan van Benthem, and István Németi. Modal languages and bounded fragments of predicate logic. *Journal of Philosophical Logic*, 27(3):217–274, 1998.

[2] Aristotle. *Prior Analytics*. Hackett, Indianapolis, IN, 1989. (Robin Smith, Tr.).

[3] Patrick Blackburn, Maarten de Rijke, and Yde Venema. *Modal Logic*. Cambridge Tracts in Theoretical Computer Science. Cambridge University Press, Cambridge, 2001.

[4] I. Borosh and L.B. Treybig. Bounds on the positive integral solutions of linear Diophantine equations. *Proceedings of the American Mathematical Society*, 55(2):299–304, 1976.

[5] D. Calvanese. *Unrestricted and finite model reasoning in class-based representation formalisms*. PhD thesis, Dipartimento di Informatica e Sistemistica, Università di Roma, 1996.

[6] A. de Morgan. *Formal Logic: or, the calculus of inference, necessary and probable*. Taylor and Walton, London, 1847.

[7] F. Eisenbrand and G. Shmonina. Carathéodory bounds for integer cones. *Operations Research Letters*, 34(5):564–568, 2006.

[8] G. Georgakopoulos, D. Kavvadias, and C. Papadimitriou. Probabilistic satisfiability. *Journal of complexity*, 1988.

[9] Kurt Gödel. Zum Entscheidungsproblem des logischen Funktionenkalküls. *Monatshefte für Mathematik und Physik*, 40:433–443, 1933.

[10] W. Goldfarb. The unsolvability of the gödel class with identity. *Journal of Symbolic Logic*, 49:1237–1252, 1984.

[11] E. Grädel. On the restraining power of guards. *Journal of Symbolic Logic*, 64:1719–1742, 1999.

[12] E. Grädel, P. Kolaitis, and M. Vardi. On the decision problem for two-variable first-order logic. *Bulletin of Symbolic Logic*, 1997.

[13] Erich Grädel, Martin Otto, and Eric Rosen. Two-variable logic with counting is decidable. In *Proceedings of the 12th IEEE Symposium on Logic in Computer Science*, pages 306–317. IEEE Online Publications, 1997.

[14] E. Hacker and W. Parry. Pure numerical Boolean syllogisms. *Notre Dame Journal of Formal Logic*, 8(4):321–324, 1967.

[15] Y. Kazakov. A polynomial translation from the two-variable guarded fragment with number restrictions to the guarded fragment. In J. J. Alferes and J. Leite, editors, *Logics in Artificial Intelligence: 9th European Conference, JELIA 2004*, volume 3229 of *Lecture Notes in Artificial Intelligence*, pages 372–384, Berlin, 2004. Springer.

[16] Y. Kazakov and I. Pratt-Hartmann. A note on the complexity of the satisfiability problem for graded modal logics. In *Proceedings, Logic in Computer Science*. IEEE, 2009. (forthcoming).

[17] L.G. Khachiyan. A polynomial algorithm in linear programming. *Soviet Mathematics Doklady*, 20:191–194, 1979.

[18] S.A. Kripke. Semantic analysis of modal logic I: normal propositional calculi. *Zeitschrift für mathematische Logik und Grundlagen der Mathematik*, 9:67–96, 1963.

[19] V. Kuncak and M. Rinard. Towards efficient satisfiability checking for Boolean algebra with Presburger arithmetic. In F. Pfenning, editor, *Proceedings, 21st International Conference on Automated Deduction (CADE-21)*, volume 4603 of *Lecture Notes in Computer Science*, pages 215–230, Berlin, 2007. Springer.

[20] R. Ladner. The computational complexity of provability in systems of modal propositional logic. *SIAM J. on Comp.*, 6:467–480, 1977.

[21] C.I. Lewis and C.H. Langford. *Symbolic Logic*. Dover, 1959. (Originally published, 1932).

[22] C. Lutz, U. Sattler, and L. Tendera. The complexity of finite model reasoning in description logics. *Information and Computation*, 199:132–171, 2005.

[23] M. Mortimer. On languages with two variables. *Zeitschrift für mathematische Logik und Grundlagen der Mathematik*, 21:135–140, 1975.

[24] W. Murphree. The numerical syllogism and existential presupposition. *Notre Dame Journal of Formal Logic*, 38(1):49–64, 1997.

[25] Leszek Pacholski, Wieslaw Szwast, and Lidia Tendera. Complexity of two-variable logic with counting. In *Proceedings of the 12th IEEE Symposium on Logic in Computer Science*, pages 318–327. IEEE Online Publications, 1997.

[26] Leszek Pacholski, Wieslaw Szwast, and Lidia Tendera. Complexity results for first-order two-variable logic with counting. *SIAM Journal on Computing*, 29(4):1083–1117, 1999.

[27] Christos H. Papadimitriou. *Computational Complexity*. Addison-Wesley, Reading, MA, 1994.

[28] J. Paris. *The Uncertain Reasoner's Companion*. Cambridge University Press, Cambridge, 1994.

[29] I. Pratt-Hartmann and L.S. Moss. Logics for the relational syllogistic. *Review of Symbolic Logic*, 2009. forthcoming.

[30] Ian Pratt-Hartmann. Complexity of the two-variable fragment with counting quantifiers. *Journal of Logic, Language and Information*, 14:369–395, 2005.

[31] Ian Pratt-Hartmann. Complexity of the guarded two-variable fragment with counting quantifiers. *Journal of Logic and Computation*, 17:133–155, 2007.

[32] Ian Pratt-Hartmann. On the complexity of the numerically definite syllogistic and related fragments. *Bulletin of Symbolic Logic*, 14(1):1–28, 2008.

[33] Ian Pratt-Hartmann. No syllogisms for the numerical syllogistic. In Orna Grumberg, Michael Kaminski, Shmuel Katz, and Shuly Wintner, editors, *Languages: From Formal to Natural*, volume 5533 of *Lecture Notes in Computer Science*, pages 192–203. Springer, Berlin, 2009.

[34] Dana Scott. A decision method for validity of sentences in two variables. *Journal of Symbolic Logic*, 27:477, 1962.

[35] Stephan Tobies. PSPACE reasoning for graded modal logics. *Journal of Logic and Computation*, 2001.