The Two-Variable Fragment with Counting Revisited

Ian Pratt-Hartmann

School of Computer Science
University of Manchester
Manchester M13 9PL
United Kingdom
http://www.cs.man.ac.uk/~ipratt

Abstract. The satisfiability and finite satisfiability problems for the two-variable fragment of first-order logic with counting were shown in [5] to be in NExpTime. This paper presents a simplified proof via a result on integer programming due to Eisenbrand and Shmonina [2].

Keywords: Logic, complexity, counting quantifiers

1 Introduction

The two-variable fragment with counting quantifiers, here denoted $\mathcal{C}^2$, is the set of function-free, first-order formulas containing at most two variables, but with the counting quantifiers $\exists_{\leq C}$, $\exists_{\geq C}$ and $\exists_{=C}$ (for every $C > 0$) allowed. Thus, for example, the sentences

No professor supervises more than three students
Every student is supervised by at most one professor

may be formalized using the respective $\mathcal{C}^2$-formulas:

$$\neg\exists x(\text{professor}(x) \wedge \exists_{\geq 4} y(\text{student}(y) \wedge \text{supervises}(x, y)))$$
$$\forall x(\text{student}(x) \rightarrow \exists_{\leq 1} y(\text{professor}(y) \wedge \text{supervises}(y, x))).$$

The *satisfiability problem for* $\mathcal{C}^2$, denoted Sat-$\mathcal{C}^2$, is the problem of determining whether a given $\mathcal{C}^2$-formula has a model. The *finite satisfiability problem for* $\mathcal{C}^2$, denoted Fin-Sat-$\mathcal{C}^2$, is the problem of determining whether a given $\mathcal{C}^2$-formula has a finite model. Since $\mathcal{C}^2$ lacks the finite model property, these problems are distinct. Both problems, however, were shown in [5] to be in NExpTime, thus improving earlier results in [3] and [4]. The proof given in that paper features a long, combinatorial argument to show that, if a $\mathcal{C}^2$-formula has a model at all, then it has a model in which only a small number of distinct 'local configurations' arise. The present paper presents a shorter and more perspicuous proof via a result on integer programming due to Eisenbrand and Shmonina [2].

2 Preliminaries

In the sequel, all signatures will be silently assumed to be purely relational. This results in no loss of generality, as function-symbols are not allowed in $\mathcal{C}^2$, and individual constants can easily be simulated by means of unary predicates. We further assume, also without loss of generality, that all predicates have arity 1 or 2. Finally, we assume all structures to be finite or countably infinite. If φ is a $\mathcal{C}^2$-formula, we write $\|\varphi\|$ to denote the total number of symbols in φ. Here, we assume numerical subscripts in counting quantifiers to be coded as *binary* strings. Thus, for example, the number of symbols contributed by a quantifier $\exists_{\leq C}$ is approximately $\lceil \log C \rceil$, where $\lceil r \rceil$ denotes the smallest integer greater than or equal to r. In this paper, all logarithms are base 2.

We begin with the reduction of $\mathcal{C}^2$-formulas to 'Scott-form'.

Lemma 1. *Let ψ be a $\mathcal{C}^2$-formula. We can generate, in time bounded by a polynomial function of $\|\psi\|$, a quantifier-free $\mathcal{C}^2$-formula α, a list of positive integers $C_1, \ldots, C_m$ and a list of binary predicates $f_1, \ldots, f_m$ $(m \geq 1)$ such that the formulas ψ and*

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

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

Proof. Routine adaptation of the re-naming technique used in [7].

Henceforth, then, we may restrict attention to $\mathcal{C}^2$-formulas of the form (1), since the truth of ψ in a model of size C or less can evidently be checked in time bounded by an exponential function of $\|\psi\|$. In the ensuing analysis of such formulas, the binary predicates $f_1, \ldots, f_m$ will play a special role. We adopt the following (non-standard) terminology.

Definition 1. *Let Σ be a signature, and $f_1, \ldots, f_m$ $(m \geq 1)$ a tuple of distinct binary predicates in Σ. The pair $\langle \Sigma, (f_1, \ldots, f_m) \rangle$ is called a* classified signature*, and the $f_1, \ldots, f_m$ are referred to as its* featured predicates.

Let Σ be a signature (not necessarily classified). We follow standard terminology, and say that a *1-type* (over Σ) is a maximal consistent set of equality-free literals over Σ involving only the variable x. Likewise, a *2-type* (over Σ) is a maximal consistent set of equality-free literals over Σ involving only the variables x and y. Reference to Σ is suppressed where clear from context. If $\mathfrak{A}$ is any structure interpreting Σ, and $a \in A$, then there exists a unique 1-type $\pi(x)$ over Σ such that $\mathfrak{A} \models \pi[a]$; we denote π 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 Σ such that $\mathfrak{A} \models \tau[a,b]$; we denote τ by $\mathrm{tp}^{\mathfrak{A}}[a,b]$. We do not define $\mathrm{tp}^{\mathfrak{A}}[a,b]$ if $a = b$. If π is a 1-type, we say that π is *realized* in $\mathfrak{A}$ if there exists $a \in A$ with $\mathrm{tp}^{\mathfrak{A}}[a] = \pi$. If τ is a 2-type, we say that τ is *realized* in $\mathfrak{A}$ if there exist distinct $a, b \in A$ with $\mathrm{tp}^{\mathfrak{A}}[a,b] = \tau$.

Notation 1 *Let τ be a 2-type over a signature Σ. The result of transposing the variables x and y in τ is also a 2-type, denoted τ^{-1}; and the set of literals in τ 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 τ not featuring the variable x, and then replacing y throughout by x.

Remark 1. If τ is any 2-type over a signature Σ, $\mathfrak{A}$ is a structure interpreting Σ, 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 $\mathrm{tp}^{\mathfrak{A}}[b] = \mathrm{tp}_2(\tau)$.

The following terminology, relating to classified signatures, is non-standard:

Definition 2. *Let $\mathfrak{A}$ be a structure interpreting a classified signature $\langle \Sigma, \bar{f} \rangle$ and C a positive integer. We say that $\mathfrak{A}$ is C-*bounded *if, for all $a \in A$ and all featured predicates f in $\bar{f}$,*

$$1 \leq |\{b \in A \setminus \{a\} \mid \mathfrak{A} \models f[a,b]\}| \leq C.$$

We say that $\mathfrak{A}$ is bounded *if it is C-bounded for some C.*

Thus, $\mathfrak{A}$ is C-bounded just in case, for every featured predicate f, no element of A is non-reflexively related to more than C elements of A by f, and every element of A is non-reflexively related to some element of A by f.

Remark 2. If φ is of the form (1), $C \geq \max_h C_h$ and $\mathfrak{A} \models \varphi$, then $\mathfrak{A}$ is C-bounded.

Definition 3. *Let $\langle \Sigma, \bar{f} \rangle$ be a classified signature, and let τ be a 2-type over Σ. We say that τ is a* message-type *(over Σ) if $f(x,y) \in \tau$ for some featured predicate f. If τ is a message-type such that τ^{-1} is also a message-type, we say that τ is* invertible. *On the other hand, if τ is a 2-type such that neither τ nor τ^{-1} is a message-type, we say that τ is a* silent *2-type.*

Thus, a 2-type τ 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 μ, then we may imagine that a sends a message (of type μ) to b. If μ is invertible, then b replies by sending a message (of type μ^{-1}) back to a. If $\mathrm{tp}^{\mathfrak{A}}[a,b]$ is silent, then neither element sends a message to the other.

3 A Result on Solutions to Integer Programming Problems

Our strategy in analysing the problems Sat-$\mathcal{C}^2$ and Fin-Sat-$\mathcal{C}^2$ is to reduce them to integer programming problems. Having done so, we shall employ a variant of a result of Eisenbrand and Shmonina [2] (also used in [6] in connection with the one-variable fragment with counting).

Lemma 2. *Let $\mathcal{E}$ be a set of m linear inequalities of the form*

$$a_0 + a_1 x_1 + \cdots + a_n x_n \leq b_0 + b_1 x_1 + \cdots + b_n x_n,$$

in variables $x_1, \ldots, x_n$, where $a_0, b_0 \in \mathbb{N}$ and $a_i, b_i \in \{0, 1\}$ for all i $(1 \leq i \leq n)$. If $\mathcal{E}$ has a solution over $\mathbb{N}$, then it has a solution over $\mathbb{N}$ in which at most $5m(\log m + 1)$ variables take non-zero values.

Proof. Routine adaptation of [6, Theorem 1].

Notice that the bound in Lemma 2 depends only on the number of equations, and not on the number of variables, nor indeed on the sizes of the constant terms.

We need to generalize this result slightly to deal with infinite solutions.

Notation 2 *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 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 3. *Let $\mathcal{E}$ be a system of m linear inequalities as in Lemma 2. If $\mathcal{E}$ has a solution over $\mathbb{N}^*$, then $\mathcal{E}$ has a solution over $\mathbb{N}^*$ in which at most $5m(\log m + 1)$ variables take non-zero values.*

Proof. Pick some solution of $\mathcal{E}$ over $\mathbb{N}^*$, and list those inequalities whose right-hand sides are infinite for this solution. For each such inequality, pick one variable x_i with infinite value whose coefficient b_i is 1. By re-ordering the variables if necessary, let $x_1, \ldots, x_k$ be the selected variables, $x_{k+1}, \ldots, x_\ell$ the other variables taking infinite values, and $x_{\ell+1}, \ldots, x_n$ the variables taking finite values. Let $\mathcal{E}'$ be the set of inequalities in $\mathcal{E}$ whose right- (and therefore left-) hand sides are finite for the given solution. Clearly, the coefficients $a_1, \ldots, a_\ell$ and $b_1, \ldots, b_\ell$ are all zero for these inequalities. Assuming $\ell < m$, $\mathcal{E}'$ therefore has a solution $(0, \ldots, 0, x'_{\ell+1}, \ldots, x'_n)$ over $\mathbb{N}$ with at most $5(m - \ell)(\log(m - \ell) + 1)$ non-zero values. But then $(\aleph_0, \ldots, \aleph_0, 0, \ldots, 0, x'_{\ell+1}, \ldots, x'_n)$, with k $\aleph_0$s, is a solution for $\mathcal{E}$.

4 The Main Result

The principal challenge in establishing upper complexity bounds for Sat-$\mathcal{C}^2$ and Fin-Sat-$\mathcal{C}^2$ consists in the very general nature of the structures we must work with. The following two notions help to reduce this generality slightly.

Definition 4. *Let $\mathfrak{A}$ be a structure interpreting a classified signature $\langle \Sigma, \bar{f} \rangle$. We say that $\mathfrak{A}$ is* chromatic *if, 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 all distinct and both $\mathrm{tp}^{\mathfrak{A}}[a, a']$ and $\mathrm{tp}^{\mathfrak{A}}[a', a'']$ are invertible message-types, then $\mathrm{tp}^{\mathfrak{A}}[a] \neq \mathrm{tp}^{\mathfrak{A}}[a'']$.*

Thus, a structure is chromatic just in case distinct elements connected by a chain of 1 or 2 invertible message-types always have distinct 1-types.

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

Definition 5. *Let $\mathfrak{A}$ be a structure interpreting a signature Σ, and Z a positive integer. We say that $\mathfrak{A}$ is Z-*differentiated *if, for every 1-type π over Σ, the number u of elements in A having 1-type π satisfies either $u \leq 1$ or $u > Z$.*

Thus, in a Z-differentiated structure, every 1-type is realized either at most once or more than Z times (possibly infinitely often).

The following lemmas have straightforward proofs [5, Lemmas 2 and 3].

Lemma 4. *Let $\mathfrak{A}$ be a C-bounded structure interpreting a classified signature $\langle \Sigma, \bar{f} \rangle$, and $m = |\bar{f}|$. Then $\mathfrak{A}$ can be expanded to a chromatic structure $\mathfrak{A}'$ by interpreting $\lceil \log((mC)^2 + 1) \rceil$ new unary predicates.*

Lemma 5. *Let $\mathfrak{A}$ be a chromatic structure interpreting a classified signature $\langle \Sigma, \bar{f} \rangle$, and Z a positive integer. Let Σ' be the signature obtained by adding $\lceil \log Z \rceil$ new unary predicates to Σ. Then $\mathfrak{A}$ can be expanded to a chromatic, Z-differentiated structure interpreting the classified signature $\langle \Sigma', \bar{f} \rangle$.*

Our next task is to acquire the means to talk about ‘local configurations’ in bounded structures interpreting a classified signature.

Notation 3 *Fix a classified signature $\langle \Sigma, \bar{f} \rangle$ with $\bar{f} = (f_1, \ldots, f_m)$ and $|\Sigma| = s$. We assume a* standard enumeration

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

of the 1-types over Σ, with arbitrary ordering, where $L = 2^s$. We likewise assume a standard enumeration

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

of the message-types *over $\langle \Sigma, \bar{f} \rangle$, where $\mu_1, \ldots, \mu_{M^*}$ are the invertible message-types, and $\mu_{M^*+1}, \ldots, \mu_M$ the non-invertible message-types.* (*Otherwise, the ordering in this enumeration is again arbitrary.*)

s	the number of symbols in Σ
$\pi_1, \ldots, \pi_L$	the 1-types over Σ
$\mu_1, \ldots, \mu_{M^*}$	the invertible message-types over $\langle \Sigma, \bar{f} \rangle$
$\mu_{M^*+1}, \ldots, \mu_M$	the non-invertible message-types over $\langle \Sigma, \bar{f} \rangle$
$\sigma_1, \ldots, \sigma_N$	the C-bounded star-types over $\langle \Sigma, \bar{f} \rangle$

Table 1. Quick reference guide to symbols used in connection with a classified signature $\langle \Sigma, \bar{f} \rangle$.

The above notation, which will be used throughout this section, is summarized in the first four rows of Table 1. We remark that $M \leq m 2^{4s-1}$.

Definition 6. *Let $\mathfrak{A}$ be a bounded structure interpreting a classified signature $\langle \Sigma, \bar{f} \rangle$, and let a be an element of A. The* star-type *of a in $\mathfrak{A}$, denoted* $\mathrm{st}^{\mathfrak{A}}[a]$*, is the M-tuple $\sigma = (v_1, \ldots, v_M)$ of natural numbers where, for all j $(1 \leq j \leq M)$,*

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

Evidently, σ satisfies the condition

$$v_j > 0 \textit{ implies } \mathrm{tp}_1(\mu_j) = \mathrm{tp}^{\mathfrak{A}}[a],$$

for all j $(1 \leq j \leq M)$. Accordingly, we take a star-type *over $\langle \Sigma, \bar{f} \rangle$ to be any M-tuple σ of natural numbers satisfying the condition*

$$v_j > 0 \textit{ and } v_{j'} > 0 \textit{ implies } \mathrm{tp}_1(\mu_j) = \mathrm{tp}_1(\mu_{j'}),$$

for all j, j' $(1 \leq j < j' \leq M)$. We denote the number v_j by $\sigma[j]$, for all j $(1 \leq j \leq M)$. A bounded structure $\mathfrak{A}$ is said to realize *a star-type σ if, for some $a \in A$, $\mathrm{st}^{\mathfrak{A}}[a] = \sigma$.*

Thus, $\mathrm{st}^{\mathfrak{A}}[a]$ is a description of a's 'local environment' in $\mathfrak{A}$. We remark that, if $\mathfrak{A}$ is not bounded, and $a \in A$, then the cardinalities $|\{b \in A \setminus \{a\} : \mathrm{tp}^{\mathfrak{A}}[a, b] = \mu_j\}|$ may be infinite. For this reason, we restrict attention to bounded structures when talking about star-types of elements.

Certain important characteristics of bounded structures depend only on the star-types they realize.

Definition 7. *Let $\langle \Sigma, \bar{f} \rangle$ be a classified signature, with $\bar{f} = (f_1, \ldots, f_m)$, and let σ be a star-type over $\langle \Sigma, \bar{f} \rangle$. We say that σ is C-*bounded*, for $C > 0$, if for all h $(1 \leq h \leq m)$,*

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

Furthermore, we say that σ is chromatic *if, for every 1-type π' over Σ, the sum*

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

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

Lemma 6. *Let $\mathfrak{A}$ be a bounded structure interpreting a classified signature $\langle \Sigma, \bar{f}\rangle$. Then $\mathfrak{A}$ is C-bounded 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 $\mathcal{F}$ is chromatic.*

Proof. Immediate once the definitions are unravelled.

The important point about C-bounded star-types over a finite classified signature $\langle \Sigma, \bar{f}\rangle$ is that there are only finitely many of them. Indeed, for a given $\langle \Sigma, \bar{f}\rangle$, and given C, we may enumerate them in a standard way as

$$\sigma_1, \ldots, \sigma_N, \tag{2}$$

just as we did with the 1-types and message-types (Table 1). Simple calculation shows that $N \leq (C+1)^M$, where M the number of message-types. It is easy to see that N is in generally doubly-exponential in $s = |\Sigma|$; however, the results of Section 3 will ensure that this is no problem. Beware that the listing (2) depends on the bound C of the star-types in question: this parameter is left implicit to reduce notational clutter.

Having obtained characterizations of 'local environments' in structures interpreting a classified signatures, we turn our attention to larger-scale aspects of those structures. We begin by considering the special role played by silent 2-types.

Definition 8. *Let $\langle \Sigma, \bar{f}\rangle$ be a classified signature. Define $\Pi^{(2)}$ to be the set of unordered pairs of (not-necessarily distinct) 1-types over Σ:*

$$\Pi^{(2)} = \{\{\pi, \pi'\} \mid \pi, \pi' \text{ 1-types over } \Sigma\}.$$

We call an element of $\Pi^{(2)}$ a quiet pair *(in $\mathfrak{A}$) if there exist distinct $a, a' \in A$ with $\mathrm{tp}^{\mathfrak{A}}[a] = \pi$ and $\mathrm{tp}^{\mathfrak{A}}[a'] = \pi'$, such that the 2-type $\mathrm{tp}^{\mathfrak{A}}[a, a']$ is silent.*

Quiet pairs can always be found in structures with populous 1-types [5, Lemma 4]:

Lemma 7. *Let $\mathfrak{A}$ be a C-bounded structure interpreting a classified signature $\langle \Sigma, \bar{f}\rangle$, and $m = |\bar{f}|$. Suppose that π and π' are 1-types over Σ (not necessarily distinct), both realized in $\mathfrak{A}$ more than $(mC+1)^2$ times. Then $\{\pi, \pi'\}$ is a quiet pair.*

For the purpose of determining satisfiability of $\mathcal{C}^2$-formulas, we can afford to be somewhat relaxed about the silent 2-types any putative model realizes.

Definition 9. *Let $\langle \Sigma, \bar{f}\rangle$ be a classified signature, $\Pi^{(2)}$ the set of unordered pairs of 1-types over Σ, and Ξ the set of silent 2-types over $\langle \Sigma, \bar{f}\rangle$. A* regulator over *$\langle \Sigma, \bar{f}\rangle$ is a partial function $\theta :\subseteq \Pi^{(2)} \to \Xi$ such that*

$$\{\mathrm{tp}_1(\theta(\{\pi, \pi'\})), \mathrm{tp}_2(\theta(\{\pi, \pi'\}))\} = \{\pi, \pi'\},$$

for every $\{\pi, \pi'\} \in \mathrm{dom}(\theta)$. Further, let $\mathfrak{A}$ be a structure interpreting $\langle \Sigma, \bar{f}\rangle$. We say that θ is a regulator for $\mathfrak{A}$, *if $\mathrm{dom}(\theta)$ is the set of quiet pairs in $\mathfrak{A}$, and for every $\{\pi, \pi'\}$ in this set, and any pair of distinct $a, a' \in A$ with $\mathrm{tp}^{\mathfrak{A}}[a] = \pi$, $\mathrm{tp}^{\mathfrak{A}}[a'] = \pi'$ and $\mathrm{tp}^{\mathfrak{A}}[a, a']$ silent, either $\mathrm{tp}^{\mathfrak{A}}[a, a'] = \theta(\{\pi, \pi'\})$ or $\mathrm{tp}^{\mathfrak{A}}[a', a] = \theta(\{\pi, \pi'\})$. Finally, we call $\mathfrak{A}$* regular *if it has a regulator.*

Roughly, a regular structure $\mathfrak{A}$ is one in which, for any quiet pair $\{\pi, \pi'\}$, we can identify a silent 2-type, $\theta(\{\pi, \pi'\})$, that relates—in one direction or the other—*all* the pairs of distinct elements a and a' having respective 1-types π and π' such that $\mathrm{tp}^{\mathfrak{A}}[a, a']$ is silent.

Lemma 8. *Let φ be any formula of the form* (1) *over a signature Σ, let $\bar{f} = (f_1, \ldots, f_m)$, and suppose $\mathfrak{A}$ is a structure over the classified signature $\langle \Sigma, \bar{f} \rangle$ such that $\mathfrak{A} \models \varphi$. Then there exists a regular structure $\mathfrak{B}$ over $\langle \Sigma, \bar{f} \rangle$ with the same domain, such that $\mathfrak{B} \models \varphi$. Moreover if $\mathfrak{A}$ is chromatic* (Z*-differentiated, for some* $Z > 0$)*, then so is $\mathfrak{B}$.*

Proof. Consider any quiet pair $\{\pi, \pi'\}$ in $\mathfrak{A}$, and pick distinct b, b' such that $\mathrm{tp}^{\mathfrak{A}}[b] = \pi$ and $\mathrm{tp}^{\mathfrak{A}}[b'] = \pi'$, with $\xi = \mathrm{tp}^{\mathfrak{A}}[b, b']$ silent. Suppose now that there exist distinct $a, a' \in A$ such that $\mathrm{tp}^{\mathfrak{A}}[a] = \pi$ and $\mathrm{tp}^{\mathfrak{A}}[a'] = \pi'$, but $\mathrm{tp}^{\mathfrak{A}}[a, a'] \neq \xi$ and $\mathrm{tp}^{\mathfrak{A}}[a', a] \neq \xi$. Let us alter $\mathfrak{A}$ to obtain a model $\mathfrak{A}'$ (say) by setting $\mathrm{tp}^{\mathfrak{A}'}[a, a'] = \mathrm{tp}^{\mathfrak{A}}[b, b']$; evidently, $\mathfrak{A}' \models \varphi$. Carrying out this transformation uniformly yields the required model $\mathfrak{B}$.

With the above apparatus at our disposal, we are in a position to characterize entire structures in terms of the patterns of local configurations they exhibit.

Definition 10. *Let $\langle \Sigma, \bar{f} \rangle$ be a classified signature, C a positive integer, and $\sigma_1, \ldots, \sigma_N$ the standard enumeration of C-bounded star-types over $\langle \Sigma, \bar{f} \rangle$. A* frame *is a quintuple $\mathcal{F} = \langle \Sigma, \bar{f}, C, K, \theta \rangle$, where K is a non-empty subset of $\{1, \ldots, N\}$, and θ is a regulator over $\langle \Sigma, \bar{f} \rangle$. We call $\mathcal{F}$* chromatic *if every σ_k ($k \in K$) is chromatic. Further, let $\mathfrak{A}$ be a bounded structure interpreting $\langle \Sigma, \bar{f} \rangle$. We say that $\mathcal{F}$* describes $\mathfrak{A}$ *just in case $\{\sigma_k \mid k \in K\}$ is exactly the set of star-types realized in $\mathfrak{A}$, and θ is a regulator for $\mathfrak{A}$.*

Lemma 9. *Let $\mathfrak{A}$ be a C-bounded regular structure over a classified signature $\langle \Sigma, \bar{f} \rangle$. Then $\mathfrak{A}$ is described by a frame of the form $\mathcal{F} = \langle \Sigma, \bar{f}, C, K, \theta \rangle$. Further, if $\mathfrak{A}$ is chromatic, then so is $\mathcal{F}$.*

Proof. Lemma 6.

Let φ be a formula of the form (1). If $\mathcal{F}$ describes $\mathfrak{A}$, then $\mathcal{F}$ contains all the information needed to determine whether $\mathfrak{A} \models \varphi$:

Definition 11. *Let φ be any formula of the form* (1) *over a signature Σ, let $\bar{f} = (f_1, \ldots, f_m)$, and let $\mathcal{F} = \langle \Sigma, \bar{f}, C, K, \theta \rangle$ be a frame, where $C \geq C_h$ for all h ($1 \leq h \leq m$). We write $\mathcal{F} \models \varphi$ if the following conditions are satisfied:*

1. *for all $k \in K$ 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 $k \in K$ 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.*
3. *for all $\{\pi, \pi'\} \in \mathrm{dom}(\theta)$, $\models \bigwedge \theta(\pi, \pi') \rightarrow \alpha(x, y) \wedge \alpha(y, x)$.*

Lemma 10. *Let φ, $\mathcal{F}$ be as in Definition 11, and suppose $\mathfrak{A}$ is a bounded structure over $\langle \Sigma, \bar{f} \rangle$ such that $\mathcal{F}$ describes $\mathfrak{A}$. Then $\mathfrak{A} \models \varphi$ if and only if $\mathcal{F} \models \varphi$.*

Proof. Immediate once the definitions are unravelled.

Lemma 9 tells us that every bounded regular structure is described by some frame. However not every frame describes a structure; and it is important for us to define a class of frames which do. The following notation will prove useful to this end.

Notation 4 *Let $\langle \Sigma, \bar{f} \rangle$ be a classified signature and $C > 0$. With reference to the standard enumerations of Table 1, and, for integers i, k in the ranges $1 \leq i \leq L$, $1 \leq k \leq N$, we write:*

$$o_{ik} = \begin{cases} 1 \textit{ if for some } j \ (1 \leq j \leq M), \ \sigma_k[j] > 0 \textit{ and } \mathrm{tp}_1(\mu_j) = \pi_i \\ 0 \textit{ otherwise;} \end{cases}$$

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

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

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

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

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

To understand the meanings of these constants, suppose $\mathfrak{A}$ is a C-bounded structure interpreting $\langle \Sigma, \bar{f} \rangle$. Then, for all i, j and k in the appropriate ranges:

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

The following notion now gives us a way of providing a 'statistical summary' of structures. Recall the extended natural numbers introduced in Notation 2.

Definition 12. *Let $\langle \Sigma, \bar{f} \rangle$ be a classified signature, C a positive integer, and $\mathfrak{A}$ a C-bounded structure interpreting $\langle \Sigma, \bar{f} \rangle$. Let $\sigma_1, \ldots, \sigma_N$ be the standard enumeration of the C-bounded star-types. The* C-histogram *of $\mathfrak{A}$ is the N-tuple* $\mathrm{Hist}_C(\mathfrak{A}) = (w_1, \ldots, w_N)$ *of elements of $\mathbb{N}^*$, where, for all k $(1 \leq k \leq N)$,*

$$w_k = |\{a \in A : \mathrm{st}^{\mathfrak{A}}[a] = \sigma_k\}|.$$

The following notation will be useful when talking about (putative) histograms of structures.

Notation 5 *Fix some frame $\mathcal{F}$ (and hence the associated constants of Notation 4), and let $w_1, \ldots, w_N$ be variables. We employ the letters u_i $(1 \leq i \leq L)$, v_j $(1 \leq j \leq M^*)$ and $x_{ii'}$ $(1 \leq i \leq L,\ 1 \leq i' \leq L)$ as shorthand for the following expressions:*

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

To understand the meanings of these expressions, suppose first that $\mathfrak{A}$ is a bounded, regular structure, described by $\mathcal{F} = \langle \Sigma, \bar{f}, C, K, \theta \rangle$, and that $\mathrm{Hist}_C(\mathfrak{A}) = (w_1, \ldots, w_N)$. Then

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 $\pi_{i'}$.

We can now give our long-awaited criterion for a frame to describe a structure.

Definition 13. *Let $\mathcal{F} = \langle \Sigma, \bar{f}, C, K, \theta \rangle$ be a frame, Z a positive integer, $m = |\bar{f}|$, and L, M^*, M, N the constants defined in Notation 4. A* Z-solution *of $\mathcal{F}$ is an N-tuple $\bar{w} = (w_1, \ldots, w_N)$ of elements of $\mathbb{N}^*$ such that, for all k $(1 \leq k \leq N)$, $w_k > 0$ if and only if $k \in K$, and such that the following conditions are satisfied for all i $(1 \leq i \leq L)$, all i' $(1 \leq i' \leq L)$, and all j $(1 \leq j \leq M^*)$:*

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

(C2) *if $u_i = 0$, then $\sum\{w_k \mid s_{ik} > 0\} = 0$; if $u_i = 1$, then $\sum\{w_k \mid s_{ik} > 1\} = 0$;*

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

(C4) *if $u_i \leq 1$, then for all positive integers $D \leq mC$, we have either $x_{i'i} \geq D$ or $\sum\{w_k \mid o_{ik} = 1$ and $r_{i'k} \geq D\} = 0$;*

(C5) *if $\{\pi_i, \pi_{i'}\} \not\in \mathrm{dom}(\theta)$, then either $u_i \leq 1$ or $u_{i'} \leq 1$;*

(C6) *if $\{\pi_i, \pi_{i'}\} \not\in \mathrm{dom}(\theta)$, then for all positive integers $D \leq mC$, we have either $x_{i',i} \leq D$ or $\sum\{w_k \mid o_{ik} = 1$ and $r_{i'k} \leq D\} = 0$.*

We say that $\bar{w}$ is finite *if each of its elements is in* $\mathbb{N}$. *If $\mathcal{F}$ has a (finite) Z-solution, we say that $\mathcal{F}$ is* (finitely) Z-solvable.

Remark 4. Noting that the constants $r_{i'k}$ in Definition 13 are bounded by mC, we see that conditions **(C4)** and **(C6)** may be more simply formulated as the collections of conditions

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

(C6*) if $\{\pi_i, \pi_{i'}\} \not\in \text{dom}(\theta)$ and $o_{ik} = 1$, then $r_{i'k} \geq x_{i'i}$,

respectively, for all i $(1 \leq i \leq L)$, i' $(1 \leq i' \leq L)$ and k $(1 \leq k \leq N)$. The reason for the rather awkward formulation adopted above will emerge presently.

The two main lemmas of this section may now be stated. They tell us that, for sufficiently large Z, we may treat (finitely) Z-solvable, chromatic frames as substitutes for (finite) bounded, Z-differentiated chromatic structures.

Lemma 11. *Let $\mathcal{F} = \langle \Sigma, \bar{f}, C, K, \theta \rangle$ be a frame, $m = |\bar{f}|$, and $Z \geq (mC + 1)^2$ be an integer. If $\mathfrak{A}$ is a (finite) bounded, Z-differentiated, structure described by $\mathcal{F}$, then* $\text{Hist}_C(\mathfrak{A})$ *is a (finite) Z-solution for $\mathcal{F}$.*

Proof (Sketch). See [5, Lemma 13] for full details. It is a routine matter to check the conditions **(C1)**–**(C6)**. Observe that condition **(C3)** is immediate from the assumption that $\mathfrak{A}$ is Z-differentiated. We note in addition that the same assumption may be used in conjunction with Lemma 7 (of this paper) to show that condition **(C5)** obtains. For suppose $\{\pi_i, \pi_{i'}\} \not\in \text{dom}(\theta)$. Since $\mathcal{F}$ describes $\mathfrak{A}$, $\{\pi_i, \pi_{i'}\}$ cannot be a quiet pair; hence either $u_i \leq Z$ or $u_{i'} \leq Z$; whence $u_i \leq 1$ or $u_{i'} \leq 1$.

Lemma 12. *Let $\mathcal{F} = \langle \Sigma, \bar{f}, C, K, \theta \rangle$ be a chromatic frame, $m = |\bar{f}|$, and $Z \geq 3mC$ be an integer. If $\mathcal{F}$ has a (finite) Z-solution, then there exists a (finite) bounded structure $\mathfrak{A}$ such that $\mathcal{F}$ describes $\mathfrak{A}$.*

Proof (Sketch). See [5, Lemma 14] for full details. For every $k \in K$, let A_k be a set of cardinality w_k, and let A be the disjoint union of the A_k. We imagine A_k as a set of elements having star-type σ_k, and show that, under the conditions **(C1)**–**(C6)**, these star-type instances can be assembled into a well-defined model $\mathfrak{A}$ with domain A. The construction depends crucially on the assumptions that the frame $\mathcal{F}$ is chromatic, and that condition **(C3)** obtains.

The next lemma tells us that, if Z-solvability is what interests us, we may restrict attention to small frames:

Lemma 13. *Let Z be a positive integer, $\mathcal{F}' = \langle \Sigma, \bar{f}, C, K', \theta \rangle$ a (finitely) Z-solvable frame, $m = |\bar{f}|$ and $s = |\Sigma|$. Then there exists a non-empty $K \subseteq K'$ such that the frame $\mathcal{F} = \langle \Sigma, \bar{f}, C, K, \theta \rangle$ is also (finitely) Z-solvable, and $|K| \leq p(mC)2^{p(s)}$, where p is a fixed polynomial.*

Proof. There are fixed polynomials p', q' such that $p'(mC)2^{q'(s)}$ bounds the number of equations **(C1)**–**(C6)** in Definition 13. (Note that this claim would in general be false if we had replaced **(C4)** and **(C6)** by their simpler variants, **(C4*)** and **(C6*)**.) By Lemmas 2 and 3, there is a polynomial p such that $\mathcal{F}$ has a (finite) solution $w_1, \ldots, w_N$ with at most $p(mC)2^{p(s)}$ non-zero values (but not none). Now let $K = \{k \in K' | w_k \neq 0\}$.

It is well known that the problem of determining whether a system $\mathcal{E}$ of linear inequalities has a solution over $\mathbb{N}$ is NPTime-complete [1], and similarly for solutions over $\mathbb{N}^*$. Indeed, if $\mathcal{E}$ has a solution over $\mathbb{N}$, then it has a solution whose size (measured in terms of the number of bits required) is bounded by a polynomial function of the total number of bits used to encode $\mathcal{E}$.

Theorem 1. *The problems Sat-$\mathcal{C}^2$ and Fin-Sat-$\mathcal{C}^2$ are in* NExpTime.

Proof. Let a $\mathcal{C}^2$-formula ψ be given. By Lemma 1, we may compute a formula φ of the form (1) in polynomial time, such that φ and ψ are satisfiable over the same domains of size greater than $C = \max(\{C_h | 1 \leq h \leq m\})$. Let $Z = (mC + 1)^2$: note that $Z \geq (mC)^2 + 1$, and also $Z \geq 3mC$. Let Σ be the signature of φ together with $2\lceil \log(Z) \rceil$ new unary predicates, and let $\bar{f} = (f_1, \ldots, f_m)$. Thus $\langle \Sigma, \bar{f} \rangle$ is a classified signature. Write $s = |\Sigma|$.

We claim that φ is (finitely) satisfiable if and only if there exists a chromatic (finitely) Z-solvable frame $\mathcal{F} = \langle \Sigma, \bar{f}, C, K, \theta \rangle$ such that $|K| \leq p(mC)2^{p(s)}$ and $\mathcal{F} \models \varphi$, where p is some fixed polynomial, independent of φ. Suppose first that φ has a (finite) model $\mathfrak{A}'$. Evidently, $\mathfrak{A}'$ is C-bounded. By Lemmas 4, 5 and 8, φ has a (finite) C-bounded, chromatic, Z-differentiated, regular model $\mathfrak{A}$ over $\langle \Sigma, \bar{f} \rangle$. By Lemma 9, there exists a chromatic frame $\mathcal{F} = \langle \Sigma, \bar{f}, C, K, \theta \rangle$ describing $\mathfrak{A}$; by Lemma 10, $\mathcal{F} \models \varphi$; and by Lemma 11, $\mathcal{F}$ has a (finite) Z-solution. Taking p to be the fixed polynomial of Lemma 13, we may assume without loss of generality that $|K| \leq p(mC)2^{p(s)}$. Conversely, suppose that $\mathcal{F} = \langle \Sigma, \bar{f}, K, C, \theta \rangle$ is a chromatic frame such that $\mathcal{F} \models \varphi$, and $\mathcal{F}$ has a (finite) Z-solution. By Lemma 12, there exists a (finite) structure $\mathfrak{A}$ such that $\mathcal{F}$ describes $\mathfrak{A}$, and by Lemma 10, $\mathfrak{A} \models \varphi$.

Consider the following non-deterministic procedure, where q_1, q_2 and q_3 are fixed polynomials, and $n = \|\varphi\|$.

1. Guess a chromatic frame $\mathcal{F} = \langle \Sigma, \bar{f}, C, K, \theta \rangle$ with $|K| \leq 2^{q_1(n)}$ and check that $\mathcal{F} \models \varphi$;
2. Guess a system of at most $2^{q_2(n)}$ linear inequalities $\mathcal{E}$ (propositionally) entailing the conditions **(C1)**–**(C6)** for $\mathcal{F}$ to have a Z-solution.
3. Guess a tuple $\bar{w}$ of elements of $\mathbb{N}^*$ whose size (number of bits) is bounded by $2^{q_3(n)}$.
4. If $\bar{w}$ is a solution for $\mathcal{E}$, succeed; else fail.

For all polynomials q_1, q_2 and q_3, this procedure runs in time bounded by an exponential function of $\|\varphi\|$. But the claim of the previous paragraph shows that, for suitable q_1, q_2 and q_3, it has a successfully terminating run if and only φ is satisfiable. This proves that Sat-$\mathcal{C}^2$ is in NExpTime. To do the same for Fin-Sat-$\mathcal{C}^2$, we simply modify line 3 to insist that $\bar{w}$ be a tuple of natural numbers.

It is well known that the satisfiability (= finite satisfiability) problem for the two-variable fragment of first-order logic *without* counting quantifiers is already NExpTime-hard. Thus, the NExpTime bound of Theorem 1 is tight.

Corollary 1. *Let φ be a formula of $\mathcal{C}^2$. If φ is finitely satisfiable, then it is satisfiable in a structure of size bounded by a doubly exponential function of $\|\varphi\|$.*

Proof. In the proof of Theorem 1, if the system $\mathcal{E}$ of equations in line 3 of the procedure has a solution over $\mathbb{N}$, then it has a solution every element of which has size (number of bits) bounded by a polynomial function of $\|\mathcal{E}\|$, and hence by a singly exponential function of $\|\varphi\|$.

It was shown in [3] that there exists a sequence $\{\varphi_n\}$ of finitely satisfiable $\mathcal{C}^2$-formulas where $\|\varphi_n\|$ is bounded above by a polynomial function of n, but the size of the smallest model of φ_n is bounded below by 2^{2^n}. Thus, the doubly-exponential bound of Corollary 1 is tight.

Acknowledgment

The author is indebted to Dr. Yevgeny Kazakov for helpful discussions on the topic of this paper.

References

1. I. Borosh and L. Treybig. Bounds on the positive integral solutions of linear Diophantine equations. *Proceedings of the American Mathematical Society*, 55(2):299–304, 1976.
2. F. Eisenbrand and G. Shmonina. Carathéodory bounds for integer cones. *Operations Research Letters*, 34(5):564–568, 2006.
3. 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.
4. 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.
5. I. Pratt-Hartmann. Complexity of the two-variable fragment with counting quantifiers. *Journal of Logic, Language and Information*, 14:369–395, 2005.
6. Ian Pratt-Hartmann. On the computational complexity of the numerically definite syllogistic and related logics. *Bulletin of Symbolic Logic*, 14(1):1–28, 2008.
7. Dana Scott. A decision method for validity of sentences in two variables. *Journal of Symbolic Logic*, 27:477, 1962.