Two-Variable Logic with Counting is Decidable*

Erich Grädel† Martin Otto† Eric Rosen‡

Preliminary version
June 1996

Abstract

We prove that the satisfiability problem for C^2 is decidable.

C^2 is first-order logic with only two variables in the presence of arbitrary counting quantifiers $\exists^{\geqslant m}$, $m \geqslant 1$. It considerably extends L^2, plain first-order with only two variables, which is known to be decidable by a result of Mortimer. Unlike L^2, C^2 does not have the finite model property. As C^2 extends L^2 by expressive means for counting, significant applications arise from the fact that C^2 embeds corresponding counting extensions of modal logics.

1 Introduction

Let L^2 be the fragment of first-order logic (with equality), that only has the variable symbols x and y — i.e. the closure of atomic formulae involving no variables apart from x and y under Boolean operations and $\exists x$, $\exists y$. Throughout this paper we restrict attention to finite vocabularies consisting of relation symbols and constants.

By [7], L^2 has the *finite model property*, i.e. every satisfiable sentence has a finite model. Consequently, the satisfiability problem for L^2 is decidable. Standard terminology uses the following:

- $sat(X)$ for the the set of $\psi \in X$ that are satisfiable;
- $fin\text{-}sat(X)$ for the set of $\psi \in X$ that have a finite model;

*This research was partially supported by the German-Israeli Foundation for Scientific Research and Development.

†Lehrgebiet Mathematische Grundlagen der Informatik, RWTH Aachen, D-52056 Aachen, {graedel,otto}@informatik.rwth-aachen.de

‡Dept. of Computer Science, Technion–Israel Institute of Technology, Haifa 32000, Israel, erosen@csa.cs.technion.ac.il

- *inf-sat*(X) for the set of $\psi \in X$ that have an infinite model;
- *inf-axioms*(X) for $sat(X) -$ *fin-sat*(X), the *infinity axioms* of X.

By the finite model property, $sat(L^2) =$ *fin-sat*(L^2). Note that *fin-sat*(X) is recursively enumerable for any formula class with effective semantics, and that $sat(X)$ is co-r.e. for any X that can effectively be embedded into first-order logic FO (by completeness). So, as a fragment of FO that has the finite model property, L^2 is decidable. In fact Mortimer also proves a recursive bound on the size of small finite models, so there is a more direct argument for decidability. Mortimer's proof and the quality of this bound have recently been improved upon in [3] so that the complexity of $sat(L^2)$ could be determined precisely.

Some of the important applications of this result arise in the context of modal logics that can be embedded into L^2. From a practical point of view, L^2 remains too weak for many applications, though. Although several extensions of modal logic like propositional dynamic logic PDL, computation tree logic CTL, or propositional μ-calculus L_μ (which feature expressive means that make them useful as process logics) are known to be decidable, it was shown in [4] that several corresponding extensions of L^2 are no longer decidable.

Unlike FO, L^2 is not closed under assertions that there are at least m elements satisfying some property for $m > 2$. It is therefore natural to extend L^2 to allow arbitrary counting quantifiers $\exists^{\geqslant m}$, $m \geqslant 1$.

Definition 1.1 C^2 is that extension of L^2 which admits all counting quantifiers $\exists^{\geqslant m}$, $m \geqslant 1$ rather than just $\exists$.

For instance, the C^2-sentence $\forall x \exists^{\leqslant m} y Exy$ defines the class of all graphs whose degree is bounded by m. Here we have already allowed derived quantifiers $\forall$ for $\neg\exists^{\geqslant 1}$ and $\exists^{\leqslant m}$ for $\neg\exists^{\geqslant m+1}$. Quantifiers $\exists^{=m}$, $\exists^{>m}$ etc. may similarly be admitted in C^2 without increasing its expressive power.

We point out one particularly interesting piece of evidence for the expressive power of C^2, which far exceeds that of L^2. Immerman and Lander [6] show that the C^2-theory of a finite graph (which is actually axiomatized a single sentence of C^2) exactly characterises the stable colouring of that graph. By results of Babai and Kučera [1] it follows that almost all finite graphs are characterized up to isomorphism by their C^2-theory — almost all in the sense of asymptotic probabilities: the proportion of graphs with vertices $0, \ldots, n-1$ having this property tends to 1 as n goes to infinity. For more information on the expressive power and model-theoretic properties of C^2 (and extensions thereof) we refer to [8, 9].

That C^2 does not have the finite model property is witnessed by

$$\forall x \exists^{=1} y Exy \wedge \forall x \exists^{\leqslant 1} y Eyx \wedge \exists x \forall y \neg Eyx,$$

which asserts that E is the graph of an injective function from the universe to itself which fails to be surjective. This is clearly an infinity axiom. Therefore *fin-sat*$(C^2) \subsetneq$ *sat*(C^2), and there are two different decidability issues to be settled.

In the present exposition we give a full proof for the decidability of *sat*(C^2). As $C^2 \subseteq$ FO, it suffices to establish that *sat*(C^2) is r.e., and for this — as *fin-sat*(C^2) clearly is r.e. — it suffices to show *inf-sat*(C^2) to be r.e.

In the full version of this paper we will also address the finite satisfiability problem. Further, we will point out some consequences of our decidability results, e.g. for modal logics with graded modalities.

2 Notation and basic definitions

Let τ always be some finite vocabulary of unary and binary relation symbols and constants. It is not difficult to see that in two-variable logics (even in the presence of constants) predicates of higher arity can be eliminated preserving satisfiability. Details will be given in the full version of this paper.

We shall suppress τ in our notation — everything will be quite uniform also with respect to τ so that it may be regarded as fixed in most arguments. If K is any set of parameters, we may treat these as new constant symbols and write $\tau_K = \tau \dot\cup K$ for the extended vocabulary. Usually K is a subset of some structure, and we do not distinguish between the element $k \in K$ and the constant k which is the syntactic name for k. We denote structures as $\mathfrak{A} = (A, \ldots)$ where A is the universe of $\mathfrak{A}$.

An *atomic n-type* over vocabulary τ is a maximally consistent set of atomic and negated atomic τ-formulae in n variable symbols.

Let $\boldsymbol{\alpha}$ denote the finite set of atomic 1-types over τ in the single variable x. Let $\boldsymbol{\beta}$ stand for the finite set of those atomic 2-types over τ in two variables x and y which include $\neg x = y$. $\boldsymbol{\alpha}^K$ will denote the set of atomic 1-types over τ_K in the variable x. $\boldsymbol{\alpha}^K$ is finite if K is.

We use letters α and β to denote typical elements of $\boldsymbol{\alpha}$ (or $\boldsymbol{\alpha}^K$) and $\boldsymbol{\beta}$, respectively. For τ-structures $\mathfrak{A}$, some fixed subset $K \subseteq A$, and distinct elements a and b in A let

$$\begin{aligned} \mathrm{atp}_{\mathfrak{A}}(a) &\in \boldsymbol{\alpha}, \\ \mathrm{atp}^K_{\mathfrak{A}}(a) &\in \boldsymbol{\alpha}^K, \\ \mathrm{atp}_{\mathfrak{A}}(a,b) &\in \boldsymbol{\beta} \end{aligned}$$

denote the respective atomic types. It is clear that for $\alpha \in \boldsymbol{\alpha}$ and quantifier-free η in the single variable x, it can effectivley be determined whether $\alpha \models \eta$, i.e. whether realizations of α necessarily satisfy η. (To see this put η in disjunctive normal form.)

Similarly $\beta \models \eta$, for $\beta \in \boldsymbol{\beta}$ and η quantifier-free, and $\alpha \models \eta$, for $\alpha \in \boldsymbol{\alpha}^K$ and quantifier-free η in the single variable x but possibly with parameters from K, are decidable.

We shall see that the analysis of structures with respect to C^2 can make good use of one further fundamental notion of type, which explicitly incorporates some limited counting information.

Definition 2.1 For an element a of $\mathfrak{A}$ let the *counting star* of a, denoted $\mathrm{ctp}_{\mathfrak{A}}(a)$, be the function

$$\begin{aligned} \gamma{:}\boldsymbol{\beta} &\longrightarrow \{0, 1, 2^+\} \\ \beta &\longmapsto \dot{\#}_{b \in A}(\mathrm{atp}_{\mathfrak{A}}(a, b) = \beta). \end{aligned}$$

$\dot{\#}_{x \in S} \ldots$ counts the number of $x \in S$ satisfying $\ldots$ according to 0, 1, 2^+ (= many). We use $\boldsymbol{\gamma}$ to denote the finite set of all non-degenerate satisfiable counting stars, i.e. of all $\mathrm{ctp}_{\mathfrak{A}}(a)$ that are realized in some structure $\mathfrak{A}$ that has at least two elements. For any particular $\mathfrak{A}$ (with at least two elements) let $\boldsymbol{\gamma}_{\mathfrak{A}} \subseteq \boldsymbol{\gamma}$ denote the set of those $\gamma \in \boldsymbol{\gamma}$ that are realized in $\mathfrak{A}$.

Any atomic 2-type $\beta \in \boldsymbol{\beta}$ uniquely determines two atomic 1-types $\mathrm{atp}_1(\beta)$ and $\mathrm{atp}_2(\beta)$ as the 1-types of x and y respectively that are prescribed in β. Note that through the $\mathrm{atp}_1(\beta)$, for those β with $\gamma(\beta) \neq 0$, each $\gamma \in \boldsymbol{\gamma}$ determines *its atomic* 1*-type* $\mathrm{atp}(\gamma)$. Semantically $\mathrm{atp}(\gamma)$ is the unique $\alpha \in \boldsymbol{\alpha}$ realized by all a that realize γ.

We shall also make use of the mapping on atomic 2-types β that exchanges the rôles of x and y:

$$\beta \longmapsto \overline{\beta} := \beta\tfrac{yx}{xy}.$$

The truncated counting information in the counting stars plays an essential rôle with respect to the following useful normal form, which only involves counting up to 2.

A normal form

Definition 2.2 Say that a sentence $\varphi \in C^2$ is in normal form if it is a conjunction of sentences of the following kinds: $\forall x \forall y \eta$ and $\forall x \exists^{=1} y \eta$, where the η are quantifier-free.

Lemma 2.3 *There is a recursive reduction* NF *from* C^2*-sentences to* C^2*-sentences in normal form (over an extended vocabulary), which is sound for satisfiability:* $\varphi \in sat(C^2)$ *if and only if* $\mathrm{NF}(\varphi) \in sat(C^2)$.

Proof. The proof is given in two parts. We first show that for any sentence $\varphi \in C^2$ of vocabulary τ (which without loss of generality contains a constant c) there are sentences φ_0 and θ in an expanded vocabulary satisfying the following.

(i) θ is a conjunction of sentences of the form $\forall x \exists^{\geqslant m} y \eta$ and $\forall x \exists^{<m} y \eta$ for quantifier-free η and $m \geqslant 1$. φ_0 is quantifier-free.

(ii) Each τ-structure has a unique expansion to a model of θ.

(iii) θ implies equivalence of φ and φ_0.

Note that (ii) and (iii) imply that φ is satisfiable if and only $\theta \wedge \varphi_0$ is satisfiable. In the second step we shall transform $\theta \wedge \varphi_0$ into normal form to finish the proof of the lemma.

θ and φ_0 are constructed inductively with respect to the number of quantifiers in φ. If φ is quantifier-free, we are done. Otherwise consider a subformula ψ of type $\exists^{\geqslant m} y \chi$, where y is free in χ, but x may or may not be free in χ. These two cases are treated separately. Consider firstly $\psi(x) = \exists^{\geqslant m} y \chi(x, y)$, with displayed variables occurring free. Introduce a new unary predicate P and let θ_1 be the conjunction of

$$\forall x \exists^{\geqslant m} y (Px \rightarrow \chi(x, y))$$
$$\forall x \exists^{<m} y (\neg Px \wedge \chi(x, y))$$

which is equivalent to $\forall x (Px \leftrightarrow \exists^{\geqslant m} y \chi(x, y))$. Let φ' be the result of replacing the subformula $\psi(x)$ in φ by the atom Px. Then φ' has fewer quantifiers than φ, θ_1 is of the desired form, and, since any model of θ_1 must interpret P as $\{x | \psi(x)\}$ it follows that $\theta_1 \models \varphi \leftrightarrow \varphi'$.

If $\psi = \exists^{\geqslant m} y \chi(y)$ does not have x as a free variable, then one may similarly use a unary P and the constant c to simulate a Boolean value in quantifier-free fashion. We may take the conjunction of the following for θ_1:

$$\forall x \exists^{\geqslant m} y (Px \rightarrow \chi(y))$$
$$\forall x \exists^{<m} y (\neg Px \wedge \chi(y)).$$

θ_1 forces P to be the full, respectively empty, predicate according to the truth value of ψ. For φ' we now take the result of substituting the atom Pc for ψ in φ. Again $\theta_1 \models \varphi \leftrightarrow \varphi'$.

An inductive application of this procedure eventually yields θ (as the conjunction of the θ_i of each step) and a quantifier-free φ_0, as desired.

It remains to transform a sentence $\theta \wedge \varphi_0$ as obtained into proper normal form without affecting satisfiability. As φ_0 cannot have any free variables, it may as well be universally quantified to form an $\forall\forall$-conjunct. Using $\forall\forall$-conjuncts to eliminate other quantifier-free constituents, we may actually assume that the quantifier-free parts of the $\forall\exists^{\geqslant m}$- and $\forall\exists^{<m}$-conjuncts in θ are atomic fromulae Pxy.

In order to translate $\forall x \exists^{\geq m} y Pxy$ into normal form we use m new binary predicates $P_1, \ldots, P_m$ and the conjunction of

$$\forall x \forall y (\bigvee_i P_i xy \ \rightarrow Pxy)$$
$$\bigwedge_{i \neq j} \forall x \forall y (P_i xy \rightarrow \neg P_j xy)$$
$$\bigwedge_i \forall x \exists^{=1} y \, P_i xy$$

For $\forall x \exists^{<m} y Pxy$, where $m > 1$ we similarly use m new binary $P_1, \ldots, P_{m-1}$ and the conjunction of

$$\forall x \forall y (Pxy \rightarrow \bigvee_i P_i xy)$$
$$\bigwedge_i \forall x \exists^{=1} y \, P_i xy$$

$\forall x \exists^{<1} y Pxy$, finally is equivalent with $\forall x \forall y \neg Pxy$ which is in normal form.

It is clear that these replacements are sound for satisfiability and yield a sentence in normal form. □

This reduction is particularly important, since counting stars alone determine which sentences in normal form are satisfied in a structure.

Lemma 2.4 *Given φ in normal form and $\boldsymbol{\gamma}_{\mathfrak{A}}$ (for any $\mathfrak{A}$ with at least two elements), it is effectively decidable whether $\mathfrak{A} \models \varphi$.*

Sketch of proof. Recall that for $\beta \in \boldsymbol{\beta}$ and quantifier-free $\eta(x, y)$ it can directly be checked whether $\beta \models \eta(x, y)$ and whether $\mathrm{atp}_1(\beta) \models \eta(x, x)$. And similarly, for $\gamma \in \boldsymbol{\gamma}$ it can be checked whether $\mathrm{atp}(\gamma) \models \eta(x, x)$.

Consider now separately the constituent sentences of φ in normal form. For an $\forall\forall$-sentence $\forall x \forall y \eta(x, y)$ it suffices to check that for all $\gamma \in \boldsymbol{\gamma}_{\mathfrak{A}}$ and any β for which $\gamma(\beta) > 0$ it is true that $\beta \models \eta(x, y)$ and $\mathrm{atp}_1(\beta) \models \eta(x, x)$.

For an $\forall \exists^{=1}$-sentence $\forall x \exists^{=1} y \eta(x, y)$ similarly it merely has to be checked that all $\gamma \in \boldsymbol{\gamma}_{\mathfrak{A}}$ satisfy the follwing: either $\sum_{\beta \models \eta(x,y)} \gamma(\beta) = 1$ and $\mathrm{atp}(\gamma) \models \neg \eta(x, x)$, or $\mathrm{atp}(\gamma) \models \eta(x, x)$ and $\sum_{\beta \models \eta(x,y)} \gamma(\beta) = 0$. For the correctness of these sums note that different β are mutually exclusive. □

3 Analysis of infinite structures

We turn to the analysis of infinite structures $\mathfrak{A}$. In a definable way we shall separate $\mathfrak{A}$ into a finite part and a rather homogeneous infinite part. We shall find *finite descriptions* for the infinite part, that suffice to check for satisfaction of sentences in normal form.

Let $\mathfrak{A}$ be a fixed infinite τ-structure. Recall that $\boldsymbol{\gamma}$ is the set of all counting stars. Split $\boldsymbol{\gamma}$ into three disjoint subsets

$$\boldsymbol{\gamma} = \boldsymbol{\gamma}_0 \,\dot\cup\, \boldsymbol{\gamma}_{\text{fin}} \,\dot\cup\, \boldsymbol{\gamma}_{\text{inf}}$$

according to whether γ is realized in $\mathfrak{A}$ not at all, or finitely often, or infinitely often. Thus, the set $\boldsymbol{\gamma}_{\mathfrak{A}}$ of counting stars realized in $\mathfrak{A}$ is $\boldsymbol{\gamma}_{\text{fin}} \,\dot\cup\, \boldsymbol{\gamma}_{\text{inf}}$.

The kings. Let $\mathfrak{K} \subseteq \mathfrak{A}$, *the kings* of $\mathfrak{A}$, be the finite substructure with universe

$$K := \{a \in A \mid \text{ctp}_{\mathfrak{A}}(a) \in \boldsymbol{\gamma}_{\text{fin}}\}.$$

Formally we here admit that $K = \emptyset$ and allow ourselves to talk about possibly empty structures. In any case, the kings are the elements of rare kinds. And as usual, they come with a court, which will consist of those elements that have special relationships with the kings.

Relationships with the kings are formalized by taking atomic 1-types with constant names for the kings into account. Recall that $\boldsymbol{\alpha}^K$ denotes the set of these types — $\boldsymbol{\alpha}^K$ is finite since K is finite.

Extended counting stars. The extended counting star is defined according to

$$\text{ctp}^K_{\mathfrak{A}}(a) := (\text{ctp}_{\mathfrak{A}}(a), \text{atp}^K_{\mathfrak{A}}(a)).$$

Let $\boldsymbol{\gamma}^K$ stand for the finite set of all those elements of $\boldsymbol{\gamma} \times \boldsymbol{\alpha}^K$ that obey the following restriction:

$$(\gamma, \alpha) \in \boldsymbol{\gamma}^K \quad \text{if for all } \beta\text{: } \gamma(\beta) \geqslant \dot\#_{k \in K}(\alpha \models \beta[x, k]).$$

It is clear that any $\text{ctp}^K_{\mathfrak{A}}(a) \in \boldsymbol{\gamma}^K$. Let Π be the natural projection

$$\begin{array}{rccc} \Pi\text{:} & \boldsymbol{\gamma}^K & \longrightarrow & \boldsymbol{\gamma} \\ & (\gamma, \alpha) & \longmapsto & \gamma. \end{array}$$

We also split $\boldsymbol{\gamma}^K$ according to the number of elements of $\mathfrak{A}$ that realize (γ, α) into

$$\boldsymbol{\gamma}^K = \boldsymbol{\gamma}^K_0 \,\dot\cup\, \boldsymbol{\gamma}^K_{\text{fin}} \,\dot\cup\, \boldsymbol{\gamma}^K_{\text{inf}}$$

The court. Let *the court* be the substructure $\mathfrak{C} \subseteq \mathfrak{A}$ with universe

$$C := \{a \in A \mid \mathrm{ctp}^K_{\mathfrak{A}}(a) \in \boldsymbol{\gamma}^K_{\mathrm{fin}}\}.$$

It is clear that $\mathfrak{K} \subseteq \mathfrak{C} \subseteq \mathfrak{A}$, where $\mathfrak{K}$ and $\mathfrak{C}$ are finite and possibly empty. If we let $\boldsymbol{\gamma}^K_{\mathfrak{A}}$ be the set of $(\gamma, \alpha) \in \boldsymbol{\gamma}^K$ that are realized in $\mathfrak{A}$, then

$$\boldsymbol{\gamma}_{\mathfrak{A}} = \Pi(\boldsymbol{\gamma}^K_{\mathfrak{A}}) \qquad \text{and also} \qquad \boldsymbol{\gamma}_{\mathrm{inf}} = \Pi(\boldsymbol{\gamma}^K_{\mathrm{inf}}).$$

The latter assertion follows from the first one, if we notice that the fibres of the projection Π are finite: $\Pi^{-1}(\gamma) \subseteq \boldsymbol{\alpha}^K$ is finite.

The counting stars of kings only depend on their court and a very rough knowledge of the rest of the society:

Remark 3.1 *Given $\mathfrak{K}, \mathfrak{C}$ such that $\mathfrak{K} \subseteq \mathfrak{C}$ and $\boldsymbol{\gamma}^K_{\mathrm{inf}} \subseteq \boldsymbol{\gamma}^K$, it is possible to determine (recursively) $\mathrm{ctp}^K_{\mathfrak{A}}(k)$ for all $k \in K$.*

Proof. Let $\mathrm{ctp}^K_{\mathfrak{A}}(k) = (\gamma, \alpha)$. Then $\alpha = \mathrm{atp}^K_{\mathfrak{A}}(a) = \mathrm{atp}^K_{\mathfrak{K}}(a)$, and the interesting information to be reconstructed is $\gamma(\beta)$ for each $\beta \in \boldsymbol{\beta}$. We distinguish two cases: if there is some $(\gamma', \alpha') \in \boldsymbol{\gamma}^K_{\mathrm{inf}}$ such that $\alpha' \models \beta[k, a]$ then $\#_{a \in A}(\mathfrak{A} \models \beta[k, a]) = \omega$, so that $\gamma(\beta) = 2^+$. If there is no such $(\gamma', \alpha') \in \boldsymbol{\gamma}^K_{\mathrm{inf}}$, then it must be the case that $\#_{a \in A}(\mathfrak{A} \models \beta[k, a]) = \#_{a \in C}(\mathfrak{A} \models \beta[k, a])$, which may be determined in $\mathfrak{C}$. □

Reduced stars. In some of the considerations to follow it will be important to classify elements of $A \setminus C$ according to how many β-edges they can have to elements *other than kings*. It is clear that this information can directly be extracted from $\mathrm{ctp}^K_{\mathfrak{A}}(a)$. For notational convenience let us introduce a function

$$\begin{aligned} \mathit{red}\colon\ \boldsymbol{\gamma}^K &\longrightarrow \boldsymbol{\gamma} \\ (\gamma, \alpha) &\longmapsto \gamma^- \end{aligned}$$

$$\text{where} \quad \gamma^-(\beta) = \begin{cases} 2^+ & \text{if } \gamma(\beta) = 2^+ \\ \gamma(\beta) - \#_{k \in K}(\alpha \models \beta[x, k]) & \text{otherwise.} \end{cases}$$

Note that by definition of $\boldsymbol{\gamma}^K$, $\mathit{red}(\gamma, \alpha)$ is a well-defined counting star — in particular no negative values can occur for $\gamma^-(\beta)$. In fact, for $a \notin K$, a realisation of $\mathit{red}(\mathrm{ctp}^K_{\mathfrak{A}}(a))$ is obtained by introducing 2 extra β-edges from a to new vertices for those β that have $\gamma(\beta) = 2^+$, and removing all former kings.

Finite characteristics. For infinite structures $\mathfrak{A}$ as considered above we abstract the following finite data as characteristic information, which we shall call the characteristic of $\mathfrak{A}$, or $char(\mathfrak{A})$ for short.

$$\boxed{\begin{array}{l} \mathfrak{K} \subseteq \mathfrak{C} \quad \text{kings and court} \\ \\ F\colon C \;\to\; \boldsymbol{\gamma}^K \\ \qquad c \;\mapsto\; \mathrm{ctp}^K_{\mathfrak{A}}(c) \\ \\ \boldsymbol{\gamma}^K_{\mathrm{inf}} \subseteq \boldsymbol{\gamma}^K \end{array}} \qquad (*)$$

Let M be the recursive set of all tuples $(\mathfrak{K}, \mathfrak{C}, F, X)$, where $\mathfrak{K} \subseteq \mathfrak{C}$ are finite τ-structures (possibly empty), F is a mapping $F\colon C \to \boldsymbol{\gamma}^K$ and $\emptyset \neq X \subseteq \boldsymbol{\gamma}^K$.

Consider the situation in which rather than $\mathfrak{A}$ itself, only $char(\mathfrak{A})$ is presented and we want to know whether $\mathfrak{A} \models \varphi$ for φ in normal form. By Lemma 2.4 it suffices to determine $\boldsymbol{\gamma}_{\mathfrak{A}} = \boldsymbol{\gamma}_{\mathrm{fin}} \cup \boldsymbol{\gamma}_{\mathrm{inf}}$. But $\boldsymbol{\gamma}_{\mathrm{fin}} = \{\mathrm{ctp}_{\mathfrak{A}}(k) \mid k \in K\}$ may, by Remark 3.1, effectively be determined from the knowledge of $\mathfrak{K}$, $\mathfrak{C}$, and $\boldsymbol{\gamma}_{\mathrm{inf}}$, which is just the projection of the given $\boldsymbol{\gamma}^K_{\mathrm{inf}}$. This yields the following.

Remark 3.2 *$\mathfrak{A} \models \varphi$ can be decided in terms of $char(\mathfrak{A})$ for φ in normal form.*

It follows that the set of C^2-sentences in normal form that are in *inf-sat*(C^2) is r.e., provided we can show that $\{char(\mathfrak{A}) \mid \mathfrak{A}$ an infinite τ-structure $\}$ is a recursive subset of M. This is shown in the following section.

4 Decidability of the characteristics

Theorem 4.1 *Given $(\mathfrak{K}, \mathfrak{C}, F, X) \in M$, it is decidable whether there is an infinite $\mathfrak{A}$ such that $(\mathfrak{K}, \mathfrak{C}, F, X) = char(\mathfrak{A})$.*

The proof is separated into two parts: in the first step we isolate three necessary conditions; these are shown to be sufficient in the second step.

Three necessary conditions

Let $(\mathfrak{K}, \mathfrak{C}, F, X) = char(\mathfrak{A})$ for some infinite $\mathfrak{A}$. We may write $\boldsymbol{\gamma}^K_{\mathrm{inf}}$ and $\boldsymbol{\gamma}_{\mathrm{inf}}$ for X and its projection $\Pi(X)$. Then the following conditions C1, C2, and C3 are satisfied.

Condition C1: Compatibility of F with $\mathfrak{K}$ and $\mathfrak{C}$.

C1 actually is a group of rather simple conditions. They assure for vertices in $\mathfrak{C}$ that F specifies atomic types (with parameters in K) in accordance with the actual atomic types in $\mathfrak{C}$; and that the counting stars specified by F are such that (1) no vertex already has more outgoing β-edges within $\mathfrak{C}$ than are allowed by its counting star, and (2) if a vertex has fewer outgoing β-edges within $\mathfrak{C}$ than required by its counting star, then there are extended counting stars in $\boldsymbol{\gamma}^K_{\text{inf}}$ which accept incoming β-edges.

C1(a): $\mathfrak{K} \subseteq \mathfrak{C}$ and for all $k \in \mathfrak{K}$: if $F(k) = (\gamma_0, \alpha_0)$, then

(i) $\gamma_0 \notin \boldsymbol{\gamma}_{\text{inf}}$.

(ii) $\alpha_0 = \text{atp}_{\mathfrak{K}}(k)$.

(iii) for all $\beta \in \boldsymbol{\beta}$:

- if $\gamma_0(\beta) \in \{0, 1\}$,
  then $\gamma_0(\beta) = \#_{c \in C}(\text{atp}_{\mathfrak{C}}(k, c) = \beta)$ and for all $(\gamma, \alpha) \in \boldsymbol{\gamma}^K_{\text{inf}}$: $\alpha \models \neg\overline{\beta}[x, k]$.
- if $\gamma_0(\beta) = 2^+$ and $\#_{c \in C}(\text{atp}_{\mathfrak{C}}(k, c) = \beta) < 2$,
  then there is some $(\gamma, \alpha) \in \boldsymbol{\gamma}^K_{\text{inf}}$ such that $\alpha \models \overline{\beta}[x, k]$.

C1(b): for all $c \in \mathfrak{C} \setminus \mathfrak{K}$: if $F(c) = (\gamma_0, \alpha_0)$, then

(i) $(\gamma_0, \alpha_0) \notin \boldsymbol{\gamma}^K_{\text{inf}}$, but $\gamma_0 \in \boldsymbol{\gamma}_{\text{inf}}$.

(ii) $\alpha_0 = \text{atp}^K_{\mathfrak{C}}(c)$.

(iii) for all $\beta \in \boldsymbol{\beta}$:

- $\gamma_0(\beta) \geqslant \#_{c' \in C}(\text{atp}_{\mathfrak{C}}(c, c') = \beta)$.
- if $\gamma_0(\beta) > \#_{c' \in C}(\text{atp}_{\mathfrak{C}}(c, c') = \beta)$,
  then there is some $(\gamma, \alpha) \in \boldsymbol{\gamma}^K_{\text{inf}}$ whose reduced counting star $\gamma^- := \mathit{red}(\gamma, \alpha)$ satisfies $\gamma^-(\overline{\beta}) > 0$.

Proof of necessity. In both cases conditions (i) and (ii) are obvious. For (iii) in (a) notice that

$$\#_{a \in A}(\text{atp}_{\mathfrak{A}}(k, a) = \beta) \text{ is infinite} \quad \text{iff} \quad \exists(\gamma, \alpha) \in \boldsymbol{\gamma}^K_{\text{inf}} \ : \ \alpha \models \overline{\beta}[x, k].$$

The implication from left to right uses finiteness of $\boldsymbol{\gamma}^K_{\text{inf}}$ and the fact that any a with $\text{ctp}^K_{\mathfrak{A}}(a) \notin \boldsymbol{\gamma}^K_{\text{inf}}$ belongs to the finite $\mathfrak{C}$.

If $\#_{a\in A}(\mathrm{atp}_{\mathfrak{A}}(k,a)=\beta)$ is finite (which in particular must be the case if $\gamma_0(\beta)=0,1$), then $\#_{a\in A}(\mathrm{atp}_{\mathfrak{A}}(k,a)=\beta)=\#_{a\in C}(\mathrm{atp}_{\mathfrak{C}}(k,a)=\beta)$.

For necessity of (b)(iii), observe that, if $\gamma_0(\beta)>\#_{c'\in C}(\mathrm{atp}_{\mathfrak{C}}(c,c')=\beta)$, then there must be an $a\in A\setminus C$ such that $\mathrm{atp}_{\mathfrak{A}}(c,a)=\beta$. But $(\gamma,\alpha):=\mathrm{ctp}^K_{\mathfrak{A}}(a)\in\boldsymbol{\gamma}^K_{\mathrm{inf}}$ and, as $c\notin K$, the reduced counting star $\gamma^-=red(\gamma,\alpha)$ admits the incoming β-edge (c,a), whence $\gamma^-(\overline{\beta})>0$. □

Condition C2: A closure property of $\boldsymbol{\gamma}^K_{\mathrm{inf}}$.

This condition assures that $\boldsymbol{\gamma}^K_{\mathrm{inf}}$ is *closed* in the sense that for any extended star type in the infinite part that requires outgoing β-edges to elements other than kings, there is another extended star type in the infinite part that can receive incoming β-edges from elements other than kings.

C2: For all $(\gamma_0,\alpha_0)\in\boldsymbol{\gamma}^K_{\mathrm{inf}}$ and $\beta\in\boldsymbol{\beta}$: if $\gamma_0(\beta)>\#_{k\in K}(\alpha_0\models\beta[x,k])$, then there is some $(\gamma,\alpha)\in\boldsymbol{\gamma}^K_{\mathrm{inf}}$ whose reduced counting star $\gamma^-:=red(\gamma,\alpha)$ satisfies $\gamma^-(\overline{\beta})>0$.

Proof of necessity. Consider all $b\in A\setminus C$ with $\mathrm{ctp}^K_{\mathfrak{A}}(b)=(\gamma_0,\alpha_0)$ — there are infinitely many of them. $\gamma_0(\beta)>\#_{k\in K}(\alpha_0\models\beta[x,k])$ implies that for each such b there is some $a\in A\setminus K$ for which $\mathrm{atp}_{\mathfrak{A}}(b,a)=\beta$. If for some b there even is such an $a\in A\setminus C$, then $(\gamma,\alpha):=\mathrm{ctp}^K_{\mathfrak{A}}(a)\in\boldsymbol{\gamma}^K_{\mathrm{inf}}$ is as desired.

If on the other hand for all b there are only $a\in C\setminus K$ with $\mathrm{atp}_{\mathfrak{A}}(b,a)=\beta$, then for some of these finitely many a it must be the case that $\gamma:=\mathrm{ctp}_{\mathfrak{A}}(a)\in\boldsymbol{\gamma}_{\mathrm{inf}}$ has $\gamma(\overline{\beta})=2^+$. Let α be such that $(\gamma,\alpha)\in\boldsymbol{\gamma}^K_{\mathrm{inf}}$. As $\gamma(\overline{\beta})=2^+$, also the reduced $\gamma^-=red(\gamma,\alpha)$ has $\gamma^-(\overline{\beta})=2^+$, and (γ,α) is as desired. □

Condition C3: A homogeneity property of $\boldsymbol{\gamma}_{\mathrm{inf}}$.

This last condition asserts that for all pairs of distinct vertices in the infinite part there is at least one β to connect them which is not limited (i.e. has value 2^+) according to the specified star types at either end.

C3: For all $\gamma,\gamma'\in\boldsymbol{\gamma}_{\mathrm{inf}}$ there is some $\beta\in\boldsymbol{\beta}$ such that $\gamma(\beta)=\gamma'(\overline{\beta})=2^+$.

This uses a result from Ramsey theory, see Theorem 1 in Chapter 5 of [5].

Fact 4.2 *If the edges of the complete bipartite graph $K_{\omega,\omega}$ are coloured with finitely many colours, then it contains monochromatically coloured copies of the complete bipartite graph $K_{n,n}$, for all n.*

Proof of necessity of C3. Note that C3 applies even to the case $\gamma=\gamma'$. Embed $K_{\omega,\omega}$ injectively into $A\setminus K$ in such a way that the parts (of the bipartition) are mapped into

$\{a \in A \mid \mathrm{ctp}_{\mathfrak{A}}(a) = \gamma\}$ and $\{a \in A \mid \mathrm{ctp}_{\mathfrak{A}}(a) = \gamma'\}$, respectively. The $\beta \in \boldsymbol{\beta}$ induce a finite colouring of the edges, and any β that admits a β-coloured copy of $K_{2,2}$ is as desired. □

Remark 4.3 *Conditions C1–C3 are recursive for given* $(\mathfrak{K}, \mathfrak{C}, F, X) \in M$.

Sufficiency of C1–C3

It remains to prove the following.

Proposition 4.4 *For any* $(\mathfrak{K}, \mathfrak{C}, F, X) \in M$ *that satisfies C1–C3, there is an infinite* $\mathfrak{A}$ *such that* $(\mathfrak{K}, \mathfrak{C}, F, X) = \mathit{char}(\mathfrak{A})$.

Proof. Let $(\mathfrak{K}, \mathfrak{C}, F, X)$ with C1–C3 be given. Beyond the elements of $\mathfrak{C}$, the desired structure $\mathfrak{A}$ has to have, for each of the finitely many $(\gamma, \alpha) \in X$, an infinite sequence of vertices that realize exactly that (γ, α). We therefore put $A := C \,\dot\cup\, (\omega \times X)$ for the universe of $\mathfrak{A}$. Think of the new vertices in $V := \omega \times X$ as divided into finitely many infinite boxes, according to their second components, which specify the extended star types these vertices shall eventually realize.

The task now is to declare atomic types for all pairs of vertices of A in such a way that

(a) a consistent interpretation of a τ-structure $\mathfrak{A}$ over A with $\mathfrak{K} \subseteq \mathfrak{C} \subseteq \mathfrak{A}$ is obtained,

(b) all elements $c \in C$ satisfy the counting star and atomic τ_K-type that is prescribed by F: $\mathrm{ctp}^K_{\mathfrak{A}}(c) = F(c)$,

(c) all elements of V satisfy the counting star and atomic τ_K-type suggested by their X-component: $\mathrm{ctp}^K_{\mathfrak{A}}(v) = (\gamma, \alpha)$ for $v = (m, (\gamma, \alpha))$, $m \in \omega$.

With respect to (c), let $\widehat{F}$ be the extension of the given F to all of A according to $\widehat{F}(m, (\gamma, \alpha)) := (\gamma, \alpha)$. Then (b) and (c) require that $\mathrm{ctp}^K_{\mathfrak{A}}(a) = \widehat{F}(a)$ for all $a \in A$. Indeed $\mathit{char}(\mathfrak{A}) = (\mathfrak{K}, \mathfrak{C}, F, X)$ then follows. In particular it should be noted that (b) and (c) imply that $\mathfrak{K} \subseteq \mathfrak{C} \subseteq \mathfrak{A}$ really become the kings and court of $\mathfrak{A}$.

Giving a full interpretation as a τ-structure to $\mathfrak{A}$ should be thought of as allocating atomic 2-types $\beta \in \boldsymbol{\beta}$ to any pair of distinct vertices in A. Thus one may think of successively *putting β-edges for suitable β* between any two distinct vertices. Note that atomic 1-types get settled automatically through the allocation of atomic 2-types.

In the following we first give a rough and intuitive sketch of how (a)–(c) can be achieved, and afterwards a more detailed and definite description. For these arguments it may be useful to consult the following diagram, which indicates the scenario we start with and also highlights some aspects of the tasks discussed below.

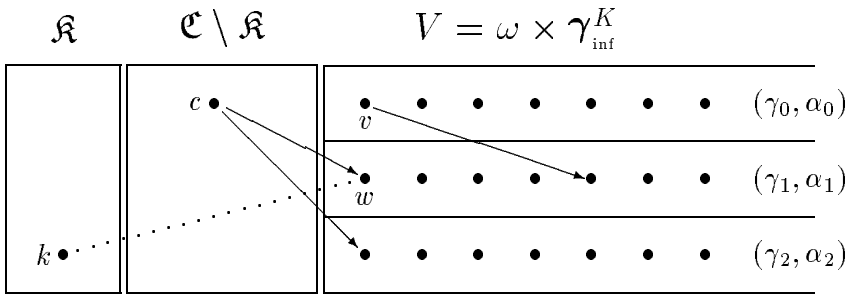

Edges involving two vertices from C are completely specified in $\mathfrak{C}$, of course. Observe also that all edges between K and V are forced by requirement (c): $\mathrm{atp}^K_{\mathfrak{A}}(w) \models \beta[k, x]$ for exactly one $\beta \in \boldsymbol{\beta}$, and $\mathrm{atp}^K_{\mathfrak{A}}(w)$ is specified by the box that w belongs to.

In fact it suffices to satisfy (b) for the elements of $C \setminus K$, because (c) then implies that (b) is also fulfilled for all $k \in K$. This follows by Remark 3.1 together with condition C1(a): if F specifies the counting star of $k \in K$ to be γ and if $\gamma(\beta) = 0$ or 1, then C1(a) guarantees that (1) the corresponding number of β-edges is already put right in $\mathfrak{C}$ and that (2) no more β-edges from k to elements of V can be introduced if the $\mathrm{atp}^K_{\mathfrak{A}}(v)$ are put right according to specification. If on the other hand $\gamma(\beta) = 2^+$, then either there are already at least two β-edges from k to elements of $\mathfrak{C}$, or condition C1(a) (iii) implies that infinitely many more will be introduced if the $\mathrm{atp}^K_{\mathfrak{A}}(v)$ for $v \in V$ are settled according to specification.

It is important to observe that as yet, i.e. with β-edges attributed to pairs over C according to the given $\mathfrak{C}$, no vertex has more outgoing β-edges than it is meant to have according to $\widehat{F}$. This follows from condition C1.

> *In the allocation of new edges, this property has to be preserved — at both ends of any new edge!*

With this provision in mind, the remaining tasks will be settled in the following order:

T1: introduction of sufficiently many outgoing β-edges of respective kinds at each $c \in C \setminus K$ to get (b) right.

T2: introduction of sufficiently many outgoing β-edges of respective kinds at each $v \in V$ to get (c) right.

T3: declaration of all remaining edges between pairs of distinct vertices (without affecting their prospective counting stars any more).

A rough sketch. Compare the diagram above for the argument. The outgoing edges for T1 and T2 (we only talk of those that do not go to kings by now) can all be chosen to go to vertices in V (rather than possibly to $C \setminus K$). Consider T1 for $c \in C \setminus K$. If this c requires β-edges to vertices outside C, then condition C1(b) (iii) says that there is some box in the infinite part, whose vertices w can accept an incoming β-edge; this precisely means that if (γ_1, α_1) is the specification for w, then the reduced star type γ^- of w has $\gamma^-(\beta) > 0$ — note that γ^- rather than γ has to be considered as α_1 may already specify incoming β-edges from kings (the dotted connection in the diagram could be a β-edge). Similarly for some $v \in V$ that requires β-edges to vertices other than kings: condition C2 now guarantess that such edges can be directed to vertices in V, whose reduced star type lets them accept incoming β-edges.

All edges introduced in phases T1 and T2 can be chosen independently in the sense that no two such edges ever go to the same vertex in V — thereby preventing the danger than any individual $v \in V$ gets overloaded. (The right order for going through the $v \in V$ for T2 is to treat them accoding to increasing first component, or column-wise in the diagram.)

At the end of this phase, all vertices have a correct number of outgoing β-edges for all β: their counting stars would be all right, only there remain edges between elements of V to be declared in phase T3. This is where condition C3 becomes essential, as it guarantees edges β that may be used between any pair of distinct vertices from $A \setminus K$ without affecting the count that is taken at either end: simply because multiplicity 2^+ cannot be spoilt by the introduction of extra edges of that kind.

The explicit construction. A detailed strategy to achieve T1–T3 can actually be given by means of choice functions telling into which one of the infinite compartments of V, $\omega \times \{(\gamma, \alpha)\}$, edges are to be directed during stages T1 and T2, and which β-edges to choose in T3. Such choice functions, f, g and h may be fixed as follows. We already write $\boldsymbol{\gamma}^K_{\text{inf}}$ and $\boldsymbol{\gamma}_{\text{inf}}$ for the given $X \subseteq \boldsymbol{\gamma}^K$ and its projection $\Pi(X)$, since this is what we want these sets to become.

- $f : (C \setminus K) \times \boldsymbol{\beta} \longrightarrow \boldsymbol{\gamma}^K_{\text{inf}}$,

such that if $F(c) = (\gamma_0, \alpha_0)$ and $\gamma_0(\beta) > \#_{c' \in C}(\text{atp}_{\mathfrak{C}}(c, c') \models \beta)$, then $(\gamma, \alpha) := f(c, \beta)$ is as guaranteed in C1(b), i.e. for the reduced counting star $\gamma^- := red(\gamma, \alpha)$: $\gamma^-(\overline{\beta}) > 0$.

- $g : \boldsymbol{\gamma}^K_{\text{inf}} \times \boldsymbol{\beta} \longrightarrow \boldsymbol{\gamma}^K_{\text{inf}}$,

such that if $\gamma_0(\beta) > \#_{k \in K}(\alpha_0 \models \beta[x, k])$, then $(\gamma, \alpha) := g((\gamma_0, \alpha_0), \beta)$ is as guaranteed in C2, i.e. for the reduced counting star $\gamma^- := red(\gamma, \alpha)$: $\gamma^-(\overline{\beta}) > 0$.

- $h : \boldsymbol{\gamma}_{\text{inf}} \times \boldsymbol{\gamma}_{\text{inf}} \longrightarrow \boldsymbol{\beta}$,

such that $\beta := h(\gamma, \gamma')$ is as guaranteed by C3, i.e. $\gamma(\beta) = \gamma'(\overline{\beta}) = 2^+$.

With these choice functions, we can give a definite description of the desired $\mathfrak{A}$, or of the allocation of edges β to all remaining pairs of distinct elements, according to T1–T3 above. To be quite definite about the sequence in which tasks are settled, let $C \setminus K = \{c_1, \ldots, c_r\}$, $\boldsymbol{\beta} = \{\beta_1, \ldots, \beta_s\}$, and $\boldsymbol{\gamma}^K_{\inf} = \{(\gamma_1, \alpha_1), \ldots, (\gamma_t, \alpha_t)\}$ be enumerations of the respective sets without repetitions.

T1: For $i = 1, \ldots, r$ / for $j = 1, \ldots, s$:
if $\Pi(F(c_i)) = \gamma$, and if $\gamma(\beta_j)$ exceeds the number of outgoing β_j-edges that c_i already has, then choose m minimal in ω such that $v = (m, f(c_i, \beta_j)) \in V$ does not yet have any incoming edges from vertices in $C \setminus K$, and put a β_j-edge, i.e. make $\mathrm{atp}_{\mathfrak{A}}(c_i, v) = \beta_j$ for that $v \in V$. This procedure is carried out one or two times, depending on whether $\gamma(\beta_j)$ exceeds the number of outgoing β_j-edges that c originally has within $\mathfrak{C}$ by 1 or 2.

This is compatible with the requirements on $\mathrm{ctp}^K(v)$, since f is such that $\mathrm{ctp}^K(v)$ admits at least one $\overline{\beta}_j$-edge to vertices outside K.

T2: For $n = 0, 1, 2, \ldots$ / for $i = 1, \ldots, t$ / for $j = 1, \ldots, s$:
if $\gamma_i(\beta_j)$ exceeds the number of outgoing β_j-edges that $v = (n, (\gamma_i, \alpha_i))$ already has, choose m minimal in ω such that $w = (m, g((\gamma_i, \alpha_i), \beta_j)) \in V$ does not yet have any incoming edges from outside K, and put a β_j-edge. That is, put $\mathrm{atp}_{\mathfrak{A}}(v, w) = \beta_j$ for that $w \in V$. Again, this procedure may have to be applied once or twice to the same v, depending on the number of β_j-edges required.

Compatibility with the requirements on $\mathrm{ctp}^K(w)$ is guaranteed by the choice of g.

T3: It remains to settle all remaining atomic 2-types, namely those between two distinct vertices v and v' in V, or between $v \in V$ and $c \in C \setminus K$, that have not yet been connected by any β-edge. For definiteness let $v = (m, (\gamma_i, \alpha_i))$ and $v' = (m', (\gamma_{i'}, \alpha_{i'}))$ with $i < i'$ or $m < m'$ in the first case, and $v = (m, (\gamma_i, \alpha_i))$ and $\gamma_{i'} := \Pi(F(c))$ in the second case. We then put $\mathrm{atp}_{\mathfrak{A}}(v, v') = \beta$ respectively $\mathrm{atp}_{\mathfrak{A}}(v, c) = \beta$, for $\beta := h(\gamma_i, \gamma_{i'})$.

Compatibility with $\mathrm{ctp}(v) := \gamma$ and $\mathrm{ctp}(v') := \gamma'$ or $\mathrm{ctp}(c) := \gamma'$ is clear, since the β selected by h is such that $\gamma(\beta) = \gamma'(\overline{\beta}) = 2^+$.

This finishes the proof of sufficiency: $\mathfrak{A}$ as constructed has $\mathit{char}(\mathfrak{A}) = (\mathfrak{K}, \mathfrak{C}, F, X)$. □

References

[1] L. Babai, L. Kučera, *Canonical labellings of graphs in linear average time*, IEEE Symp. on Foundations of Computer Science (1980), pp. 39-46.

[2] E. Börger, E. Grädel, and Y. Gurevich, *The Classical Decision Problem*, Springer, 1996, to appear.

[3] E. Grädel, P. Kolaitis, and M. Vardi, *On the complexity of the decision problem for two-variable first-order logic.* in preparation.

[4] E. Grädel, M. Otto, and E. Rosen, *Undecidability results on two-variable logics.* Preprint, 1996.

[5] R. Graham, B. Rothschild, and J. Spencer, *Ramsey Theory.* John Wiley and Sons, 1980.

[6] N. Immerman and E. Lander, *Describing graphs: a first-order approach to graph canonization*, in A. Selman, ed., Complexity Theory Retrospective, Springer, 1990, pp. 59–81.

[7] M. Mortimer, *On languages with two variables*, Zeitschr. f. math. Logik u. Grundlagen d. Math., 21 (1975), pp. 135–140.

[8] M. Otto, *Bounded variable logics and counting — a study in finite models.* Habilitationsschrift RWTH Aachen, 1995.

[9] ——, *Ptime canonization for two variables with counting*, Proceedings 10th IEEE Symposium on Logic in Computer Science, LICS95, San Diego (1995), 342–352.