EPiC Series in Computing
Volume 57, 2018, Pages 214–232
LPAR-22. 22nd International Conference on Logic for Programming, Artificial Intelligence and Reasoning

Two-variable First-Order Logic with Counting in Forests

Witold Charatonik[1]*, Yegor Guskov[2],
Ian Pratt-Hartmann[2,3], and Piotr Witkowski[1]*

[1] Institute of Computer Science, University of Wrocław, Poland
Witold.Charatonik@cs.uni.wroc.pl Piotr.Witkowski@cs.uni.wroc.pl
[2] School of Computer Science, University of Manchester, UK
Yegor.Guskov@gmail.com ipratt@cs.man.ac.uk
[3] Instytut Informatyki, Uniwersytet Opolski

Abstract

We consider an extension of two-variable, first-order logic with counting quantifiers and arbitrarily many unary and binary predicates, in which one distinguished predicate is interpreted as the mother-daughter relation in an unranked forest. We show that both the finite satisfiability and the general satisfiability problems for the extended logic are decidable in NExpTime. We also show that the decision procedure for finite satisfiability can be extended to the logic where two distinguished predicates are interpreted as the mother-daughter relations in two independent forests.

1 Introduction

Two-variable logics. The two-variable fragment of first-order logic, here denoted FO^2, is the set of function-free, first-order formulas (with equality) featuring at most two variables. The two-variable fragment with counting, here denoted $\mathcal{C}^2$, is the set of function-free, first-order formulas featuring at most two variables, but with the counting quantifiers $\exists_{[\leq C]}$, $\exists_{[\geq C]}$ and $\exists_{[=C]}$, (for every $C \geq 0$) allowed. The following facts are known. The logic FO^2 has the finite model property [7], and its satisfiability (= finite satisfiability) problem is NExpTime-complete [5]. The logic $\mathcal{C}^2$ is expressive enough for the finite model property to fail; nevertheless, its satisfiability and finite satisfiability problems remain NExpTime-complete [6, 8, 9].

Contributions, techniques. A *forest* is a well-founded directed graph (possibly infinite) in which each vertex has at most one incoming edge. It is impossible, in first-order logic, to express the fact that the graph of a given binary relation is a forest. This suggests enriching FO^2 and $\mathcal{C}^2$ by adding such a facility. Denote by $\mathcal{C}^2[\downarrow]$ and $\mathcal{C}^2[\downarrow_1, \downarrow_2]$ the extensions of $\mathcal{C}^2$ in which respectively one or two distinguished binary predicates are required to be interpreted as unranked forests. Within the logic $\mathcal{C}^2[\downarrow]$ one can state that the graph of a given binary relation is connected, since every connected graph has a spanning tree.

*Supported by Polish National Science Centre grant No 2016/21/B/ST6/01444.

A restriction of the logic $\mathcal{C}^2[\downarrow_1, \downarrow_2]$ was investigated in [4], in which the forests in question are required to have bounded rank: for some fixed number m, no element of either forest has more than m daughters. It was shown that the finite satisfiability problem for this logic is in NExpTime. In the present paper, the restriction to forests of bounded rank is lifted, without affecting the complexity of finite satisfiability. It is also shown that the satisfiability problem for the logic $\mathcal{C}^2[\downarrow]$ is in NExpTime.

Theorem 1. *The finite satisfiability problems for $\mathcal{C}^2[\downarrow]$ and $\mathcal{C}^2[\downarrow_1, \downarrow_2]$ are in* NExpTime.

Theorem 2. *The satisfiability problem for $\mathcal{C}^2[\downarrow]$ is in* NExpTime.

Theorem 1 strengthens the result of [4], while simplifying its proof. The key to this simplification is a closer look at the algorithm given in [10] for deciding finite satisfiability in $\mathcal{C}^2$, which the proof in [4] modifies. By undertaking a more fine-grained analysis of the complexity of this algorithm, we can reduce the problem of determining finite satisfiability of a formula in $\mathcal{C}^2[\downarrow]$ (or $\mathcal{C}^2[\downarrow_1, \downarrow_2]$) to that of determining the finite satisfiability of an *exponentially larger* $\mathcal{C}^2$-formula. This allows us to treat the result of [10] as, essentially, a black box, which was not possible in [4]. Theorem 2 uses new techniques to reduce the general satisfiability problem for $\mathcal{C}^2[\downarrow]$ to the corresponding finite satisfiability problem. In the context of extensions of $\mathcal{C}^2$, proving general satisfiability usually is not an easy task. In fact, most such extensions, including [4, 3, 1], do not touch this problem; an exception is [11]. Thus, the contributions of the present paper are: (a) we lift the earlier restriction to forests of bounded rank; and (b) we consider general satisfiability for $\mathcal{C}^2[\downarrow]$ (thus allowing interpretation over a single, infinite forest). The techniques employed here also work for other extensions of $\mathcal{C}^2$ over trees; in a separate paper, we prove the decidability of the finite satisfiability problem for $\mathcal{C}^2[\downarrow, \rightarrow]$, where the additional predicate $\rightarrow$ is interpreted as the next-sister-relation.

Related work. The languages considered in the present paper contain both counting quantifiers and arbitrarily many binary predicates; however, they employ only a single navigational predicate, namely the mother-daughter relation $\downarrow$. When counting quantifiers are absent, and only unary predicates appear in the signature, we can extend the palette of navigational possibilities while retaining decidability of finite satisfiability. Thus, for example [2] considers the logic FO^2 and its guarded fragment, GF^2, interpreted over a single finite tree, and accessed by various combinations of navigational predicates including $\downarrow$ (mother-daughter), $\downarrow^+$ (mother-descendant), $\rightarrow$ (next-sister) and $\rightarrow^+$ (older-sister). Complexity of satisfiability for these logics varies from PSpace to ExpSpace. Note that all these logics allow vocabularies with arbitrarily many unary-, but no uninterpreted (i.e. non-navigational) binary predicates, and some of them additionally impose the unary alphabet restriction (exactly one predicate holds on each node of a structure). With these extended sets of navigational possibilities, the addition of both counting quantifiers and uninterpreted binary predicates is known to produce significant increases in complexity. Thus, for example, taking the logic $\mathrm{FO}^2[*, 0, \downarrow, \downarrow^+, \rightarrow, \rightarrow^+]$ as a starting point (two-variable first-order logic with arbitrarily many unary predicates, no binary predicates, and the indicated navigational predicates) it was shown in [1] that extending this logic with *either* additional binary predicates *or* counting quantifiers does not increase the complexity of finite satisfiability (ExpSpace); however, extending with both yields a logic whose decidability status is unknown, but is at least as hard as the non-emptiness problem in Vector Addition Tree Automata (VATA).

2 The finite satisfiability problem

Overview of the decision procedure. Although $\mathcal{C}^2$ cannot express that a distinguished predicate $\mathfrak{t}^{\mathfrak{M}}$ is a forest in every finite model $\mathfrak{M}$, it can express that the graph of $\mathfrak{t}^{\mathfrak{M}}$ consists of a number of trees and a number of cycles. Under certain conditions, by a simple rewiring of the model we may implant such a cycle into a tree thus removing the cycle; by repeating this procedure we remove all cycles one by one and obtain a model where $\mathfrak{t}$ is interpreted as a forest.

Given an input $\mathcal{C}^2[\downarrow]$ formula ψ whose finite satisfiability is to be determined we produce an equisatisfiable $\mathcal{C}^2$ formula ψ_S by adding to ψ some conjuncts that encode a forest. The parameter S of ψ_S, called a *shrubbery*, is a description of tree components in a model, and enables the implantation procedure. The shrubbery S is of size exponential in the signature of ψ and can be simply guessed; this however gives ψ_S of exponential size, which might lead to an exponentially slower algorithm. One of contributions of this paper is the notion of *effective size* of a formula. We show that ψ_S has polynomial effective size and that the satisfiability of such formulas can be tested without moving to a higher complexity class.

The implantation procedure is very similar to the one used in [4]. However, the techniques developed here not only lift the restriction in [4] to ranked trees, but also enable us to use the result of [10] directly, without having to reconstruct the proof given there (as was done in [4]).

2.1 Preliminaries

In the sequel, *formula* means a formula of $\mathcal{C}^2$. A formula φ is in *normal form* if it conforms to the pattern:

$$\forall x \forall y(x = y \vee \alpha) \wedge \bigwedge_{h=1}^{m} \forall x \exists_{[\prec_h C_h]} y(\beta_h \wedge x \neq y), \tag{1}$$

where α and the β_h are quantifier-free, equality-free formulas, m is a positive integer, the $\prec_h$ are either $=$ or $\leq$, and the C_h are (bit-strings representing) non-negative integers. The integer m is the *multiplicity* of φ, and the integer $C = \max(C_1, \ldots, C_m)$ the *ceiling* of φ. The *modulus* of φ is the formula $\theta = \beta_1 \vee \ldots \vee \beta_m$. We assume without loss of generality that $C \geq 1$. Observe that if a $\prec_h$ is $=$ then ψ is not satisfiable over any domain of cardinality less than or equal C_h. The following lemma uses a familiar technique originally employed in [12] in the context of FO^2.

Lemma 1 ([10, Lemma 1])**.** *Given a $\mathcal{C}^2$-formula φ, we can compute, in polynomial time, a normal-form $\mathcal{C}^2$-formula ψ, with ceiling C, such that, for any set A of cardinality greater than C, ψ is satisfiable (in either $\mathcal{C}^2$, $\mathcal{C}^2[\downarrow]$ or $\mathcal{C}^2[\downarrow_1, \downarrow_2]$) over the domain A if and only if φ is.*

In the sequel, we shall require fine control over the various parameters in a normal form formula, and in order to obtain this, we generalize the notion slightly. Let θ be a quantifier-free formula. A formula ψ with free variable x is θ-*eclipsed* if it is a Boolean combination of formulas which are either (i) quantifier-free or (ii) of the form $\exists_{[\bowtie B]} y\, \chi$, where B is a non-negative integer, the symbol $\bowtie$ is chosen from $\{\leq, =, \geq\}$, and χ is a quantifier-free formula such that $\models \chi \rightarrow \theta$. A formula ψ is in *weak normal form* if it conforms to the pattern:

$$\varphi \wedge \bigwedge_{g=1}^{\ell} \left(\exists_{[\bowtie_g B_g]} x.\, \xi_g \right) \wedge \forall x.\eta, \tag{2}$$

where φ is in normal form with modulus θ, $B_1, \ldots, B_\ell$ are non-negative integers, $\xi_1, \ldots, \xi_\ell$ are quantifier-free formulas with free variable x, the symbols $\bowtie_1, \ldots, \bowtie_\ell$ are chosen from $\{\leq, =, \geq\}$, and η is θ-eclipsed. We define the *multiplicity* and *ceiling* of ψ to be the multiplicity and ceiling of φ respectively. We take formulas that are trivially equivalent to (weak) normal form formulas to be in (weak) normal form by courtesy. Thus, any normal-form formula is automatically also in weak normal form.

The satisfiability and finite satisfiability problems for $\mathcal{C}^2$ are decidable and in fact NExpTime-complete [6, 8, 9]. In this paper, we require fine control over both the parameters of the formula and the features of the model in question. Let ψ be a (weak) normal-form $\mathcal{C}^2$-formula over a signature Σ having ceiling C and multiplicity m. The *size* of ψ, denoted $|\psi|$, is the number of symbols it contains (with quantifier subscripts coded in binary). We say that the *effective size* of ψ is the quantity $|\Sigma| + \log(|\psi|) + \log(mC)$.

Theorem 3. *There exists a non-deterministic procedure which, given a $\mathcal{C}^2$-formula ψ in weak normal form, runs in time bounded by a fixed exponential function of the effective size of ψ, and which has a successful run if and only if ψ is finitely satisfiable.*

Proof. An immediate corollary of the proof of [10, Theorem 1], in which the finite satisfiability of a normal-form formula ψ with multiplicity m and ceiling C is reduced to the solvability over $\mathbb{N}$ of a guessed system $\mathcal{E}$ of Boolean combinations of linear inequalities, together with a check that $\mathcal{E}$ verifies the satisfiability of ψ. The number of linear inequalities in $\mathcal{E}$ is, by inspection [10, p. 51, formulas (**C1**)–(**C6**)], exponential in $|\Sigma| + \log(mC)$. Furthermore, checking that $\mathcal{E}$ verifies the satisfiability of ψ requires time polynomial in $|\psi|$. Thus, the procedure runs in (nondeterministic) exponential time. The relaxation to *weak* normal form requires just two changes. First, to deal with the conjuncts $\exists_{[\bowtie_g B_g]} x.\, \xi_g$, we must add to the system $\mathcal{E}$ an additional collection of ℓ (in)equalities. Second, when checking that $\mathcal{E}$ verifies the satisfiability of ψ, we take into account the eclipsed conjunct $\forall x.\eta$. Both changes are completely routine. □

A *forest* is a directed graph $G = (V, E)$ in which E is inverse-functional and well-founded. That is: for all $v \in V$ there exists at most one $u \in V$ such that uEv, and V contains no infinite reverse chains $v_0, v_1, \ldots$ with $v_{i+1} E v_i$ for all $i \geq 0$. It follows from well-foundedness that E is anti-symmetric (for all $u, v \in V$, uEv implies $\neg vEu$), and hence irreflexive. A connected forest is called a *tree*.

In this section, we employ the distinguished binary predicate $\mathfrak{t}$ whose interpretation, in the logic $\mathcal{C}^2[\downarrow]$, is constrained to be a forest. In order to be able to discuss both $\mathcal{C}^2$ and $\mathcal{C}^2[\downarrow]$ together, we call any structure $\mathfrak{A}$ in which the graph $(A, \mathfrak{t}^{\mathfrak{A}})$ is a forest *dendral*. Thus, $\mathcal{C}^2[\downarrow]$ is the logic $\mathcal{C}^2$ restricted to dendral structures. The universe of a structure $\mathfrak{A}$ is denoted A.

For technical reasons we employ a further distinguished binary predicate, $\mathfrak{s}$, and two distinguished unary predicates, s^+, s^-. We define the $\mathcal{C}^2$-formula Δ, featuring these predicates, to be the conjunction $\Delta_0 \wedge \Delta_1 \wedge \Delta_2 \wedge \Delta_3$, where

$$\begin{aligned}
\Delta_0 &:= \forall x \forall y(\mathfrak{t}(x,y) \to \neg \mathfrak{t}(y,x)) \wedge \forall x \exists_{[\leq 1]} y.\mathfrak{t}(y,x) \\
\Delta_1 &:= \forall x \forall y(\mathfrak{s}(x,y) \to \mathfrak{t}(x,y)) \wedge \forall x \exists_{[\leq 1]} y.\mathfrak{s}(x,y) \\
\Delta_2 &:= \forall x(s^+(x) \to (\exists_{[=1]} y.\mathfrak{s}(x,y) \wedge \forall y \neg \mathfrak{s}(y,x))) \\
\Delta_3 &:= \forall x((\exists_{[=1]} y.\mathfrak{s}(y,x) \wedge \forall y \neg \mathfrak{s}(x,y)) \to s^-(x)).
\end{aligned}$$

For any finite model $\mathfrak{A}$ of Δ_0, the relation $\mathfrak{t}^{\mathfrak{A}}$ is anti-symmetric and inverse functional. Hence, it is easy to see that the graph $G = (A, \mathfrak{t}^{\mathfrak{A}})$ consists of components which are either trees or which contain $\mathfrak{t}$-cycles, namely sequences of distinct elements $a_0, \ldots, a_{n-1}$ ($n \geq 3$) such that, writing $a_n = a_0$, we have $\mathfrak{A} \models \mathfrak{t}[a_i, a_{i+1}]$ for all i ($0 \leq i < n$). Elements belonging to tree-components will be said to be *dendral*; elements belonging to $\mathfrak{t}$-cycles will be said to be *cyclic*. Some elements may be neither dendral nor cyclic; however, if there are no cyclic elements, then all elements are dendral. Any element $a \in A$ such that there is no b with $\mathfrak{A} \models \mathfrak{t}[b, a]$ will be called a *root*. By definition, elements not incident on any $\mathfrak{t}$-edge (which form isolated vertices of G) are roots. It is obvious that all roots are dendral. Of course, if $\mathfrak{A}$ is dendral, then $\mathfrak{A} \models \Delta_0$, and, moreover, every element of $\mathfrak{A}$ is dendral.

Assuming that Δ_0 holds, Δ_1 then states that the graph of $\mathfrak{s}$ consists of zero or more disjoint, linear sequences of $\mathfrak{t}$-edges. A maximal sequence $a_0, \ldots, a_n$ ($n \geq 1$) such that $\mathfrak{A} \models \mathfrak{s}[a_i, a_{i+1}]$ for all i ($0 \leq i < n$), will be called an $\mathfrak{s}$-*chain*. The formula Δ_2 ensures that $\mathfrak{s}$-chains begin with all elements satisfying s^+, while Δ_3 ensures that $\mathfrak{s}$-chains end only with elements satisfying s^-.

Any formula with two free variables is assumed to have those variables taken in the order x, y. Thus, we write $\mathfrak{A} \models \theta[a, b]$, where a, b are elements of A, to indicate that θ is satisfied in $\mathfrak{A}$ under the assignment $a \mapsto x$ and $b \mapsto y$. We denote by $\theta(y, x)$ the result of transposing the variables in θ. Finally, if ξ has x as its only free variable, we denote by $\xi(y)$ the result of replacing x by y in ξ.

Let Σ be a signature of unary and binary predicates. A *1-type* is a maximal consistent set of literals over Σ involving only the variable x. Likewise, a *2-type* is a maximal consistent set of literals over Σ involving only the variables x and y and containing the literal $x \neq y$. If τ is a 2-type, we denote by τ^{-1} the 2-type obtained by exchanging the variables x and y in τ, and call τ^{-1} the *inverse* of τ. We denote by $\mathrm{tp}_1(\tau)$ the 1-type obtained by removing from τ any literals containing y; and we denote by $\mathrm{tp}_2(\tau)$ the 1-type obtained by first removing from τ any literals containing x, and then replacing all occurrences of y by x. Evidently, $\mathrm{tp}_2(\tau) = \mathrm{tp}_1(\tau^{-1})$. We equivocate freely between finite sets of formulas and their conjunctions; thus, we treat 1-types and 2-types as formulas, where convenient. Let $\mathfrak{A}$ be any structure interpreting Σ. If $a \in A$, then there exists a unique 1-type π such that $\mathfrak{A} \models \pi[a]$; we denote π by $\mathrm{tp}^{\mathfrak{A}}[a]$ and say that a *realizes* π. If, in addition, $b \in A \setminus \{a\}$, then there exists a unique 2-type τ such that $\mathfrak{A} \models \tau[a, b]$; we denote τ by $\mathrm{tp}^{\mathfrak{A}}[a, b]$ and say that the pair a, b *realizes* τ. Evidently, in that case, $\tau^{-1} = \mathrm{tp}^{\mathfrak{A}}[b, a]$; $\mathrm{tp}_1(\tau) = \mathrm{tp}^{\mathfrak{A}}[a]$; and $\mathrm{tp}_2(\tau) = \mathrm{tp}^{\mathfrak{A}}[b]$.

The following terminology helps us to characterize configurations of pairs of elements in structures. Let θ be a quantifier-free, equality-free formula over Σ, and $\mathfrak{A}$ any structure interpreting Σ. A θ-*ray* in $\mathfrak{A}$ is an ordered pair of distinct elements $\langle a, b \rangle \in A^2$ such that $\mathfrak{A} \models \theta[a, b]$; we say that the ray in question is *emitted by* a and *absorbed by* b, or simply that a sends a θ-ray to b. A θ-*ray-type* is a 2-type ρ over Σ such that $\models \rho \rightarrow \theta$. (Thus, a θ-ray-type is the 2-type of some possible θ-ray.) We refer to $\mathrm{tp}_1(\rho)$ and $\mathrm{tp}_2(\rho)$ as the *emission*-1-type and *absorption*-1-type of ρ, respectively. For Y a positive integer, we say that $\mathfrak{A}$ has θ-*degree* Y if no element of $\mathfrak{A}$ emits more than Y θ-rays. As an illustration of these concepts, let ψ be a formula over Σ of the form (1), and suppose $\mathfrak{A} \models \psi$. Setting $\theta := \beta_1 \vee \cdots \vee \beta_m$ and $Y = C_1 + \cdots + C_m$, we see that $\mathfrak{A}$ has θ-degree Y. A θ-ray $\langle a, b \rangle$ is said to be *invertible* if also $\mathfrak{A} \models \theta[b, a]$. Similarly, with ray-types: a θ-ray-type ρ is said to be *invertible* if $\models \rho^{-1} \rightarrow \theta$. A 2-type τ is said to be θ-*dark* if neither τ nor τ^{-1} is a θ-ray-type.

The following two lemmas can be established by simple counting arguments.

Lemma 2. *Let θ be a quantifier-free, equality-free formula over a signature Σ. Let $\mathfrak{A}$ be a structure interpreting Σ, of θ-degree Y, and suppose B, B' are subsets of A of cardinality $2(Y+1)$. Then there exist $a \in B$, $b \in B'$ such that $a \neq b$ and $tp^{\mathfrak{A}}[a, b]$ is θ-dark.*

Lemma 3. *Let θ be a quantifier-free, equality-free formula over a signature Σ. Let $\mathfrak{A}$ be a structure interpreting Σ, of θ-degree Y, and suppose B, B' are subsets of A of cardinalities $(3Y+2)$ and $2(Y+1)$ respectively. Then, for any $b' \in A$, there exist $a \in B$, $b \in B'$, such that $a \neq b$, $tp^{\mathfrak{A}}[a, b]$ is θ-dark, and b' sends no θ-ray to a.*

The following terminology helps us to characterize structures in terms of the configurations of elements that arise in them. Let $\mathfrak{A}$ be a structure interpreting signature Σ, X a positive integer and θ a quantifier-free, equality-free formula over Σ. We call $\mathfrak{A}$ X-*differentiated* if, for every 1-type π over Σ, the set $A_\pi = \{a \in A \mid \text{tp}^{\mathfrak{A}}[a] = \pi\}$ satisfies either $|A_\pi| \leq 1$ or $|A_\pi| > X$. We call $\mathfrak{A}$: (i) θ-*semichromatic* if no invertible θ-ray has the same emission- and absorption 1-type; (ii) θ-*chromatic* if it is θ-semichromatic and no element emits two or more invertible θ-rays having the same absorption-type as each other; and (iii) θ-*superchromatic* if it is θ-semichromatic and no element emits two or more θ-rays at least one of which is invertible, having the same absorption-type as each other. Note that a θ-semichromatic structure $\mathfrak{A}$ is θ-chromatic if there is no triple of distinct elements b_1, a, b_2, with $\text{tp}^{\mathfrak{A}}[b_1] = \text{tp}^{\mathfrak{A}}[b_2]$, such that $\text{tp}^{\mathfrak{A}}[b_1, a]$ and $\text{tp}^{\mathfrak{A}}[a, b_2]$ are invertible θ-ray-types; likewise, $\mathfrak{A}$ is θ-superchromatic if there is no triple of distinct elements b_1, a, b_2, with $\text{tp}^{\mathfrak{A}}[b_1] = \text{tp}^{\mathfrak{A}}[b_2]$, such that $\text{tp}^{\mathfrak{A}}[b_1, a]$ is an invertible θ-ray-type and $\text{tp}^{\mathfrak{A}}[a, b_2]$ a θ-ray-type. We can write the definitions of these concepts as various types of $\mathcal{C}^2$-formulas.

Lemma 4. *Let $\mathfrak{A}$ be a structure interpreting a signature Σ, and θ a quantifier-free, equality-free formula over Σ. There exist θ-eclipsed formulas χ_θ^-, χ_θ and χ_θ^+ such that: (i) $\mathfrak{A}$ is θ-semichromatic if and only if $\mathfrak{A} \models \forall x.\chi_\theta^-$, (ii) $\mathfrak{A}$ is θ-chromatic if and only if $\mathfrak{A} \models \forall x.\chi_\theta$; (iii) $\mathfrak{A}$ is θ-superchromatic if and only if $\mathfrak{A} \models \forall x.\chi_\theta^+$. All formulas have size at most $O((|\theta|+|\Sigma|)\cdot 2^{|\Sigma|})$.*

For structures with bounded θ-degree Y, θ-(super)chromaticity and X-differentiation can be ensured by expanding the signature with a suitable collection of unary predicates:

Lemma 5. *Let θ be a quantifier-free, equality-free formula interpreting a signature Σ, and suppose $\mathfrak{A}$ is a structure interpreting Σ with θ-degree Y. Then $\mathfrak{A}$ can be expanded to a θ-chromatic structure over a signature extending Σ with at most $\lceil \log(Y^2+1) \rceil$ new unary predicates; moreover, $\mathfrak{A}$ can be expanded to a θ-superchromatic structure over a signature extending Σ with at most $\lceil \log(2Y^2+1) \rceil$ new unary predicates.*

Lemma 6 ([10, Lemma 5])**.** *Let $\mathfrak{A}$ be a structure and X a positive integer. Then $\mathfrak{A}$ can be expanded to an X-differentiated structure $\mathfrak{A}'$ by interpreting at most $\lceil \log X \rceil$ additional unary predicates. If $\mathfrak{A}$ is θ-(super)chromatic, for some θ, then so is $\mathfrak{A}'$.*

We now construct apparatus for describing the 'local environment' of elements in superchromatic structures interpreting Σ. Let θ be a quantifier-free, equality-free formula over Σ, and let the θ-ray-types be listed in some fixed order (depending on Σ and θ) as $\rho_1, \ldots, \rho_M$. A θ-*star-type* is an $(M+1)$-tuple $\sigma = \langle \pi, v_1, \ldots, v_M \rangle$, where π is a 1-type over Σ and the v_j are cardinal numbers (not-necessarily finite) such that $v_j \neq 0$ implies $\text{tp}_1(\rho_j) = \pi$ for all j $(1 \leq j \leq M)$. We denote the 1-type π by $\text{tp}(\sigma)$. To motivate this terminology, suppose $\mathfrak{A}$ is a structure interpreting Σ. For any $a \in A$, we define $\text{st}_\theta^{\mathfrak{A}}[a] = \langle \text{tp}^{\mathfrak{A}}[a], v_1, \ldots, v_M \rangle$, where $v_j = |\{b \in A : b \neq a \text{ and } \text{tp}^{\mathfrak{A}}[a, b] = \rho_j\}|$. Evidently, $\text{st}_\theta^{\mathfrak{A}}[a]$ is a θ-star-type; we call it the θ-*star-type of* a in $\mathfrak{A}$, and say that a *realizes* $\text{st}_\theta^{\mathfrak{A}}[a]$. Intuitively, the θ-star-type of an element records the number of θ-rays of each type emitted by some element. It helps to think, informally, of a θ-star-type σ as *emitting* a collection of θ-rays of various types, as shown in Fig. 1b). To understand the significance of θ-star-types, consider again the formula ψ given in (1), and again let $\theta := \beta_1 \vee \cdots, \vee \beta_m$. If $\mathfrak{A}$ is a structure interpreting the signature of ψ, whether $\mathfrak{A} \models \psi$ is determined entirely by the 2-types and the θ-star-types realized in $\mathfrak{A}$. More formally, we say that a 2-type τ is *compatible with* ψ if $\tau \wedge \alpha \wedge \alpha(y, x)$ is consistent; similarly a star-type $\sigma = \langle \pi, v_1, \ldots, v_M \rangle$ is compatible with ψ if (i) each of the ray-types emitted by σ is compatible with ψ and, (ii) for all h $(1 \leq h \leq m)$, the total number of rays whose type entails β_h is,

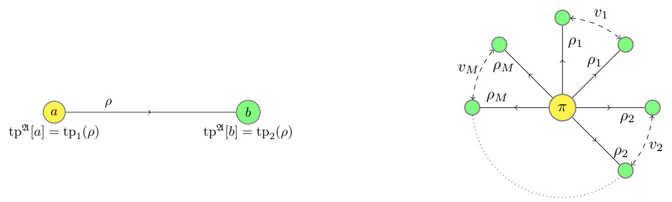

(a) (b)

Figure 1: Depiction of: (a) an element a sending a ray of type ρ to an element b in a structure $\mathfrak{A}$; and (b) a star-type $\langle \pi, v_1, v_2, \ldots, v_M \rangle$, emitting v_j rays of type ρ_j for all j $(1 \leq j \leq M)$.

respectively, equal to C_h (if $\prec_h$ is $=$) or bounded by C_h (if $\prec_h$ is $\leq$):

$$\sum_{j=1}^{m} \{v_j \mid 1 \leq j \leq M \text{ and } \models \rho_j \to \beta_h\} \prec_h C_h$$

Thus, $\mathfrak{A} \models \psi$ just in case all the θ-star-types and θ-dark 2-types realized in $\mathfrak{A}$ are compatible with ψ. These notions are extended to a formula ψ in weak normal form in the obvious way.

A θ-star-type σ is: (i) *semichromatic* if it does not emit any invertible θ-rays with absorption type $\mathrm{tp}(\sigma)$; (ii) *chromatic* if it is semichromatic and does not emit any two invertible θ-rays that have the same absorption-type as each other; *superchromatic* if it is semichromatic and does not emit two θ-rays, at least one of which is invertible, that have the same absorption-type as each other. Thus, a structure interpreting Σ is θ-(semi/super-) chromatic if and only if every θ-star-type it realizes is (semi/super-) chromatic.

We finish these preliminaries with some notation for labelling elements in structures. Let $\bar{d} = d_1, \ldots, d_n$ be a sequence of unary predicates. For all k $(0 \leq k < 2^n)$, we abbreviate by $\bar{d}\langle k \rangle$ the quantifier-free, equality-free formula $\delta_1 \wedge \cdots \wedge \delta_n$, where, for all j $(1 \leq j \leq n)$, δ_j is $d_j(x)$ if the j-th bit in the n-digit binary representation of k is 1, and $\neg d_j(x)$ otherwise. We call $\bar{d}\langle k \rangle(x)$ the k-th *labelling formula* (*over* $d_1, \ldots, d_n$). Evidently, if $A = \{a_0, \ldots, a_{M-1}\}$ is a set of cardinality $M \leq 2^n$, then we can interpret the predicates in d_j $(1 \leq j \leq n)$ over A so as to ensure that, for all k $(0 \leq k < M)$, a_k satisfies $\bar{d}\langle k \rangle$.

2.2 Shrubberies and their logical encodings

Turning to the logic $\mathcal{C}^2[\downarrow]$, for the rest of this section, we assume that all signatures feature the distinguished binary predicate $\mathfrak{t}$. Let Σ be a signature. A 2-type over Σ containing either of the atoms $\mathfrak{t}(x, y)$ or $\mathfrak{t}(y, x)$ will be said to be *arboreal*. A *shrubbery* over Σ is a triple $S = (V, E, L)$, where (V, E) is a non-empty, finite forest and L a labelling function defined on $V \cup E$ such that:

(i) for all $v \in V$, $L(v)$ is a 1-type (over Σ);

(ii) for all $(u, v) \in E$, $L(u, v)$ is either a 2-type (over Σ) containing the atom $\mathfrak{t}(x, y)$ or is the special symbol \$.

We refer to any edge labelled \$ as a *special* edge; all other edges are *ordinary*. We define $|S|$, the *size* of S, to be the cardinality of the set V, and we assume that V is enumerated in some fixed way as $\{v_0, \ldots, v_{|S|-1}\}$. For any positive integer X, we say that S is X-*differentiated* if, for every 1-type π over Σ, either $|L^{-1}(\pi)| \leq 1$ or $|L^{-1}(\pi)| > X$—that is, if either at most one or more than X vertices in V are labelled with any particular 1-type.

By way of motivation, we note that, if $\mathfrak{A}$ is dendral, then we can always construct a shrubbery S by taking any subgraph G of the forest $(A, \mathfrak{t}^{\mathfrak{A}})$, with vertices and edges labelled by their 1- and 2-types, and then optionally collapsing some number of linear paths in G to single edges (i.e. taking a topological minor), labelling the newly-formed edges with the special symbol, \$ (Fig. 2). Indeed, if $\mathfrak{A}$ is X-differentiated, then, by selecting G appropriately, and being careful which chains we collapse, we can ensure that S also is X-differentiated. This observation is formalized in Lemma 7.

Suppose that $S = (V, E, L)$ is a shrubbery over Σ. We define a formula $\Delta^S = \Delta^S_1 \wedge \cdots \wedge \Delta^S_8$ encoding S. Our formula features collections $p_1, \ldots, p_n$, $s_1, \ldots, s_n$ of new unary predicates, where $n = \lceil \log(|S| + 1) \rceil$, in addition to the predicates $\mathfrak{t}$, $\mathfrak{s}$, s^+ and s^- featured in Δ. Recall that $\bar{p}\langle k \rangle$ denotes the k-th labelling formula ($0 \leq k < 2^n$) over $p_1, \ldots, p_n$, and similarly for $\bar{s}\langle k \rangle$. For ease of reading, we present the conjuncts of Δ^S using their English glosses under the (helpful) assumption that the formula Δ is true; of course, however, these are really $\mathcal{C}^2$-formulas in weak normal form (modulo trivial logical manipulation). Thus, for instance, we have $\Delta^S_1 \equiv \bigwedge_{k=0}^{|S|-1} \exists_{[=1]} x.\bar{p}\langle k \rangle$; Δ^S_2–Δ^S_8 are purely universal. The complete list is as follows.

Δ^S_1: For $0 \leq k < |S|$, there exists a unique element satisfying $\bar{p}\langle k \rangle$, intuitively, the k-th element in the enumeration of V.

Δ^S_2: No element corresponding to a root of the forest (V, E) has any incoming $\mathfrak{t}$-edges.

Δ^S_3: The element corresponding to any vertex of V has the 1-type given to that vertex by L, and the elements corresponding to any ordinary edge of E have the 2-type given to that edge by L.

Δ^S_4: Any vertex u such that $\langle u, v \rangle$ is a special edge corresponds to an element satisfying s^+ and hence starts an $\mathfrak{s}$-chain.

Δ^S_5: For $0 \leq k < |S|$, if the first element of an $\mathfrak{s}$-chain satisfies $\bar{p}\langle k \rangle$, then all subsequent elements satisfy $\bar{s}\langle k \rangle$.

Δ^S_6: If (v_i, v_j) is a special edge of S, then any $\mathfrak{s}$-chain with non-initial elements satisfying $\bar{s}\langle i \rangle$, has final element satisfying $\bar{p}\langle j \rangle$.

Δ^S_7: The only 1-types realized in the structure are those labelling the vertices of S, while the only arboreal 2-types realized in the structure are those labelling the ordinary edges of S.

Δ^S_8: The only 1-types realized more than once in the structure are those labelling more than one vertex of S.

Observe that Δ^S_4, Δ^S_5 and Δ^S_6 together ensure that vertices of S linked by special edges correspond to elements of $\mathfrak{A}$ joined by $\mathfrak{s}$-chains.

2.3 The reduction

We present a non-deterministic procedure, `FinSatF(`ψ`)`, for determining whether a given weak normal-form $\mathcal{C}^2[\downarrow]$-formula, ψ, has a finite, dendral model. Since ψ is in weak normal form, we may write it as

$$\varphi \wedge \bigwedge_{g=1}^{\ell} \left(\exists_{[\bowtie_g B_g]} x.\, \xi_g \right) \wedge \forall x.\eta, \tag{3}$$

where φ is a normal-form formula with multiplicity m, ceiling C, and modulus θ'. Define $Y = mC + 1$ and $X = 3Y + 2$, and let Σ^- be the signature of ψ (which we may assume contains the distinguished predicate $\mathfrak{t}$), and let Σ be Σ^- together with $\lceil \log(2Y^2 + 1) \rceil + \lceil \log X \rceil$ fresh unary predicates. Let $\theta := \theta' \vee \mathfrak{t}(y, x)$, and let χ^+_θ be the formula of Lemma 4 encoding the property of θ-superchromaticity over Σ. Recall the sentence Δ governing the predicates $\mathfrak{s}$, s^+ and s^- (which, we may assume, do not belong to Σ). A simple check shows that the formula $\psi_S := \psi \wedge \forall x\, \chi^+_\theta \wedge \Delta \wedge \Delta^S$ is in weak normal form with modulus $\theta \vee \mathfrak{s}(x, y) \vee \mathfrak{s}(y, x)$. (The additional disjuncts are required by Δ.) Let `FinSat`$(\cdot)$ be the procedure for testing finite satisfiability of a $\mathcal{C}^2$-formula in weak normal form, as guaranteed by Theorem 3. The procedure `FinSatF`(ψ) consists of two steps:

1. Non-deterministically guess an X-differentiated shrubbery S over Σ of size at most $5(X \cdot 2^{|\Sigma|} + 2^{4|\Sigma|})$, and compute the formula $\psi_S := \psi \wedge \forall x\, \chi^+_\theta \wedge \Delta \wedge \Delta^S$, in weak normal form.
2. Run `FinSat`(ψ_S) and report the result.

Recall that `FinSat` runs in time bounded by an exponential function of the effective size of its argument.

2.4 Correctness: direction 1

We show that, if ψ has a finite, dendral model, then the procedure `FinSatF`(ψ) has a successful run. We begin with a property secured by the formula $\Delta \wedge \Delta^S$.

Lemma 7. *If $\mathfrak{A}$ is a finite, X-differentiated dendral structure interpreting Σ, then there exists an X-differentiated shrubbery S over Σ, of size at most $5(X \cdot 2^{|\Sigma|} + 2^{4|\Sigma|})$, such that $\mathfrak{A}$ can be expanded to a model $\mathfrak{A}^+ \models \Delta \wedge \Delta^S$.*

Proof. Let $L = 2^{|\Sigma|}$ be the number of 1-types realized in $\mathfrak{A}$, and $M \leq 2^{4|\Sigma|}$ the number of arboreal 2-types realized in $\mathfrak{A}$. For every 1-type π realized in $\mathfrak{A}$ exactly once, mark the unique element satisfying π. For every 1-type π realized in $\mathfrak{A}$ more than once, select X elements satisfying π, and mark them. For every arboreal 2-type τ realized in $\mathfrak{A}$, pick distinct elements a, b realizing it (in either direction), and mark those elements. Let V_0 be the set of marked elements; bearing in mind that arboreal 2-types are never equal to their inverses, it is enough to pick $|V_0| \leq XL + M$ elements. Let W be the set of elements of V_0 together with all their ancestors in the forest $(A, \mathfrak{t}^{\mathfrak{A}})$, and let F be the restriction of $\mathfrak{t}^{\mathfrak{A}}$ to W. Thus, $H = (W, F)$ is a forest in which every leaf vertex is in V_0, so that H has at most $XL + M$ branches. Let V_1 be the set of vertices in W having at least two daughters in H; thus, $|V_1| \leq XL + M$. Let V_2 be the set of daughters of any vertex of $V_0 \cup V_1$ in H; thus $|V_2| \leq |V_0 \cup V_1| + (XL + M)$. Let $V = V_0 \cup V_1 \cup V_2$; thus $|V| \leq 5(XL + M) \leq 5(X \cdot 2^{|\Sigma|} + 2^{4|\Sigma|})$. Let V be enumerated as $v_0, \ldots, v_{|V|-1}$. Observe that, for all $w \in W \setminus V$, w has exactly one F-predecessor and exactly one F-successor. That is: the elements of $W \setminus V$ correspond to linear strands in the forest H.

Remembering that $V \subseteq A$, for every $v \in V$, define $L(v) = \mathrm{tp}^{\mathfrak{A}}[v]$. Let E_0 be the restriction of F to V, and, for any edge $\langle u, v \rangle \in E_0$, define $L(u, v) = \mathrm{tp}^{\mathfrak{A}}[u, v]$. By construction of the graph H, $L(u, v)$ contains the atom $\mathfrak{t}(x, y)$. Let E_1 be the set of ordered pairs $\langle u, v \rangle$ from V such that there exists a path in the forest H of the form $u = a_0, \ldots, a_m = v$ ($m \geq 2$) such that for all i ($1 \leq i < m$), $a_i \in W \setminus V$. In that case, call the sequence of elements $\{a_0, \ldots, a_m\}$ a *special chain*, denoted $S(u)$, and we define $L(u, v) = \$$. In other words, special chains are linear strands in the forest F leading from one element of V to another, having no elements of V between the two termini. Let $E = E_0 \cup E_1$. Thus (V, E) is a finite forest and $S = (V, E, L)$ a

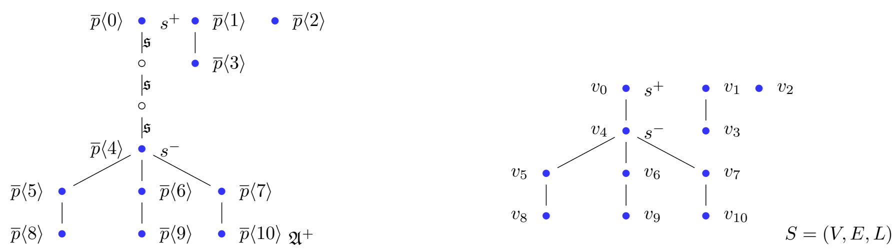

Figure 2: Extraction of a shrubbery S from a dendral structure $\mathfrak{A}$ and its encoding in $\mathfrak{A}^+$ (proof of Lemma 7). White nodes represent elements of $W \setminus V$, in this case forming a single $\mathfrak{s}$-chain.

shrubbery. Fig 2 shows a small example. Observe that the edges in E_0 are ordinary edges of S (labelled with 2-types), and those in E_1, *special edges* (labelled with \$).

Recalling the enumeration $v_0, \ldots, v_{|V|-1}$ of V, we proceed to expand $\mathfrak{A}$ to a model $\mathfrak{A}^+$ satisfying Δ^S. First we interpret the predicates $p_1, \ldots, p_n$ so that v_k satisfies the labelling formula $\bar{p}\langle k\rangle$ for all k ($0 \leq k < |V|$), and so that all elements of $A \setminus V$ satisfy $\bar{p}\langle |V|\rangle$. Recalling that each special edge (u, v) corresponds to a special chain $u = a_0, \ldots, a_m = v$, we take $u = a_0$ to satisfy the unary predicate s^+, $v = a_m$ to satisfy the unary predicate s^-, and each pair $\langle a_i, a_{i+1}\rangle$ ($0 \leq i < m$) to satisfy the binary predicate $\mathfrak{s}$. Finally, we interpret the predicates $s_1, \ldots, s_n$ so that each element a_i ($1 \leq i \leq m$) satisfies the labelling formula $\bar{s}\langle k\rangle$, where $u = v_k$. Thus, the $\bar{s}$ index of each element in a special chain (except the first) equals the $\bar{p}$-index of the first element. All elements of A which are not members of special chains can be taken to satisfy $\bar{s}\langle |V|\rangle$. It is straightforward to check that $\mathfrak{A}^+ \models \Delta \wedge \Delta^S$. □

We are now in a position to show that, if ψ, as given in (3), has a finite, dendral model, then the procedure `FinSatF`(ψ) has a successful run. Suppose $\mathfrak{A}^- \models \psi$, with $\mathfrak{A}^-$ finite and dendral, interpreting the signature Σ^-. Let $Y = mC + 1$ and $X = 3Y + 2$. Since $\mathfrak{A}^- \models \varphi$, $\mathfrak{A}^-$ has θ-degree at most Y. By Lemmas 5 and 6, we can expand $\mathfrak{A}^-$ to a θ-superchromatic, X-differentiated structure $\mathfrak{A}$ over Σ. Observe that Σ has the requisite number of spare unary predicates. By Lemma 4, $\mathfrak{A} \models \forall x.\chi^+_\theta$. By Lemma 7, let S be an X-differentiated shrubbery S of size at most $5(X \cdot 2^{|\Sigma|} + 2^{4|\Sigma|})$, such that $\mathfrak{A}$ can be further expanded to a model $\mathfrak{A}^+ \models \Delta \wedge \Delta^S$. Since ψ_S has a finite model, `FinSat`(ψ_S) has a successful run, and so therefore does `FinSatF`(ψ).

2.5 Correctness: direction 2

We show that, if the procedure `FinSatF`(ψ) has a successful run, then ψ has a finite, dendral model. Recall that, in a finite model $\mathfrak{A} \models \Delta_0$, an element is said to be dendral if it belongs to a tree-component of the graph $(A, \mathfrak{t}^{\mathfrak{A}})$. Let X be a positive integer. We say that a finite model $\mathfrak{A} \models \Delta_0$ is X-*viable* if:

(i) every 1-type realized in $\mathfrak{A}$ is realized by a dendral element;

(ii) every 1-type realized by more than one element in $\mathfrak{A}$ is realized by at least X dendral elements; and

(iii) every arboreal 2-type realized in $\mathfrak{A}$ is realized by a pair of dendral elements.

Lemma 8. *Let $\mathfrak{A}^+$ be a finite model of $\Delta \wedge \Delta^S$, where S is an X-differentiated shrubbery over Σ, and let $\mathfrak{A}$ be the reduct of $\mathfrak{A}^+$ to Σ. Then $\mathfrak{A}$ is X-viable.*

(a) Noninvertible case

(b) Invertible case

Figure 3: Model rewiring

Sketch proof. Write $S = (V, E, L)$, enumerate V as $\{v_0, \ldots, v_{|S|-1}\}$, and let A be the domain of $\mathfrak{A}^+$. Since $\mathfrak{A}^+ \models \Delta_1^S$, there exists, for each k ($0 \leq k < |S|$), a unique $b_k \in A$ satisfying $\bar{p}\langle k \rangle$. Thus, the map $\iota : v_k \mapsto b_k$ is an embedding of V in A. The main idea of the proof is to show that every element in the image of ι is dendral, making essential use of the formulas Δ, Δ_2, Δ_3, Δ_4, Δ_5 and Δ_6. The X-viability of $\mathfrak{A}^+$ then follows almost immediately from Δ_1^S, Δ_3^S, Δ_7^S, Δ_8^S and the assumed X-differentiation of $\mathfrak{A}$. □

Lemma 9. *Let Σ be a signature, θ a quantifier-free, equality-free formula over Σ, and $\mathfrak{A}$ a finite, θ-superchromatic structure of θ-degree Y, interpreting Σ. Let $X = 3Y + 2$, and suppose $\mathfrak{A}$ is X-viable and contains at least one $\mathfrak{t}$-cycle. Then there exists an X-viable structure $\mathfrak{A}'$ interpreting Σ over the same domain, realizing the same 2-types and the same θ-star-types, and containing fewer $\mathfrak{t}$-cycles.*

Proof. Let $\langle a', b' \rangle$ be any edge of some $\mathfrak{t}$-cycle in $\mathfrak{A}$, and let $\pi = \mathrm{tp}^{\mathfrak{A}}[a']$, $\pi' = \mathrm{tp}^{\mathfrak{A}}[b']$, and $\tau = \mathrm{tp}^{\mathfrak{A}}[a', b']$. We proceed to break the $\mathfrak{t}$-cycle containing a' and b', thus rendering all of its elements dendral. Since $\mathfrak{A} \models \mathfrak{t}[a', b']$ and $\mathfrak{t}(y, x)$ is a conjunct in θ, τ^{-1} is necessarily a θ-ray; this θ-ray may be invertible or non-invertible.

We consider first the case where τ^{-1} is non-invertible. Since $\mathfrak{A}$ is X-viable, π and π' are realized more than once in $\mathfrak{A}$, and so must be realized at least X times by dendral elements. Setting B to be the set of dendral elements of 1-type π and B' the set of dendral elements of 1-type π', by Lemma 3, there exist dendral elements a, b, such that $\mathrm{tp}^{\mathfrak{A}}[a] = \pi$, $\mathrm{tp}^{\mathfrak{A}}[b] = \pi'$, $\mathrm{tp}^{\mathfrak{A}}[a, b]$ is θ-dark, and b' sends no θ-ray to a. Thus, $\mathrm{tp}^{\mathfrak{A}}[a, b']$ is either a non-invertible θ-ray-type or is θ-dark. Indeed, since a is dendral, but b' is not, $\mathrm{tp}^{\mathfrak{A}}[a, b']$ is not arboreal. Thus, we may define the structure $\mathfrak{A}'$ to be exactly like $\mathfrak{A}$ except that

$$\mathrm{tp}^{\mathfrak{A}'}[a, b] = \mathrm{tp}^{\mathfrak{A}}[a, b'] \qquad \mathrm{tp}^{\mathfrak{A}'}[a, b'] = \mathrm{tp}^{\mathfrak{A}}[a', b']$$
$$\mathrm{tp}^{\mathfrak{A}'}[a', b'] = \mathrm{tp}^{\mathfrak{A}}[a, b],$$

as illustrated in Fig. 3a. It is immediate that $\mathfrak{A}'$ and $\mathfrak{A}$ realize the same 2-types, and clear by inspection of Fig. 3a that the star-type of every element is the same in $\mathfrak{A}'$ as in $\mathfrak{A}$. On the other hand, the $\mathfrak{t}$-cycle containing a' and b' has been broken in $\mathfrak{A}'$, and all its elements have become dendral.

We consider next the case where τ^{-1}—and hence $\tau = \mathrm{tp}^{\mathfrak{A}}[a', b']$—is an invertible θ-ray-type. Since $\mathfrak{A}$ is X-viable, let a, b be dendral elements such that $\mathrm{tp}^{\mathfrak{A}}[a, b] = \tau$. Since $\mathfrak{A}$ is θ-superchromatic, we know that $\mathrm{tp}^{\mathfrak{A}}[a, b']$ and $\mathrm{tp}^{\mathfrak{A}}[a', b]$ are θ-dark. In that case, we simply set

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

as illustrated in Fig. 3b. It is again immediate that $\mathfrak{A}'$ and $\mathfrak{A}$ realize the same 2-types, and clear by inspection of Fig. 3b that the star-type of every element is the same in $\mathfrak{A}'$ as in $\mathfrak{A}$. On the other hand, the $\mathfrak{t}$-cycle containing a' and b' has again been broken in $\mathfrak{A}'$, and all its elements have become dendral. □

We are now in a position to show that, if the procedure `FinSatF`(ψ) has a successful run, then ψ has a finite, dendral model. For suppose S is the shrubbery guessed in the first step: let $\mathfrak{A}^+$ be a finite model of ψ_S, and let $\mathfrak{A}$ be the reduct of $\mathfrak{A}^+$ to Σ. Since $\mathfrak{A}^+ \models \Delta \wedge \Delta^S$ and S is, by assumption, $(3Y+2)$-differentiated, it follows by Lemma 8 that $\mathfrak{A}$ is $(3Y+2)$-viable. Since $\mathfrak{A}^+ \models \forall x\, \chi_\theta^+$, $\mathfrak{A}$ is also θ-superchromatic. Now apply Lemma 9 to obtain the structure $\mathfrak{A}'$. Since this process leaves the θ-star-types of elements unchanged, and never causes dendral elements to become non-dendral, $\mathfrak{A}'$ is a θ-superchromatic, $(3Y+2)$-viable structure. Thus, we may continue this process until we obtain a structure $\mathfrak{A}^*$ containing no cycles at all. Then $\mathfrak{A}^*$ is the desired dendral model of ψ.

This shows that `FinSatF`(ψ) yields the correct result. To analyse the running time, let the multiplicity and ceiling of ψ be m and C, respectively, and let the signature of ψ be Σ. Examination of the size of the formula ψ_S then establishes the following lemma.

Lemma 10. *Let ψ be a weak normal-form $\mathcal{C}^2[\downarrow]$-formula with multiplicity m and ceiling C over signature Σ. We can non-deterministically compute a $\mathcal{C}^2$-formula ψ_S in weak normal-form with multiplicity m_S and ceiling C_S over signature Σ_S, such that: (i) ψ has a finite dendral model if and only if, for some run of the computation, ψ_S has a finite model; (ii) $|\psi_S|$ is at most $|\psi| \cdot 2^{O(|\Sigma|)}(mC)^{O(1)}$; (iii) $m_S = m+2$ and $C_S = C$; (iv) $|\Sigma_S|$ is $O(|\Sigma| + \log(mC))$. The computation of ψ_S requires time polynomial in $|\psi_S|$.*

Lemmas 1 and 10 and Theorem 3 imply the first part of Theorem 1.

2.6 Generalization to two forests

The above argument can be unproblematically extended to the logic $\mathcal{C}^2[\downarrow_1, \downarrow_2]$, where two distinguished predicates, $\mathfrak{t}_1$ and $\mathfrak{t}_2$ are, required to be interpreted as forests.

Lemma 11. *The finite satisfiability problem for the logic $\mathcal{C}^2[\downarrow_1, \downarrow_2]$ is in* NExpTime.

In the proof of this lemma we use the following well-known fact about graph-colouring.

Lemma 12. *If $G = (A, E)$ is a directed graph with out-degree m, then the underlying undirected graph of G has a proper $(2m+1)$-colouring.*

Sketch proof of Lemma 11. Here, we sketch the principal differences to $\mathcal{C}^2[\downarrow]$. Call any structure in which the predicates $\mathfrak{t}_1$ and $\mathfrak{t}_2$ are interpreted as forests *dendral*. Given a formula ψ of $\mathcal{C}^2[\downarrow_1, \downarrow_2]$ in weak normal-form, we again construct a shrubbery together with a $\mathcal{C}^2$-formula $\psi_S = \psi \wedge \forall x\, \chi_\theta^+ \wedge \Delta \wedge \Delta^S$, such that ψ has a finite dendral model if and only if ψ_S has a finite model. The notion of a shrubbery is modified so that it is the union of two (not necessarily disjoint) forests. In the case where $\mathfrak{A}$ has a finite dendral model, a shrubbery is obtained as a subgraph of the coloured graph $(A, \mathfrak{t}_1^{\mathfrak{A}}, \mathfrak{t}_2^{\mathfrak{A}})$, including sufficiently many ordinary edges, and with any very long connecting strands contracted to special edges. (We now need two types of special edges: one for each forest.) The formulas Δ and Δ^S are modified in the obvious way. In particular, the conjunct Δ_0 of Δ states that both $\mathfrak{t}_1$ and $\mathfrak{t}_2$ are irreflexive and inverse functional.

The key idea behind the construction, given in Lemma 9, then proceeds almost identically to the case $\mathcal{C}^2[\downarrow]$. The principal difference is that we must remove both $\mathfrak{t}_1$-cycles and $\mathfrak{t}_2$-cycles. Suppose $a_0, \ldots, a_{n-1}$ is a $\mathfrak{t}_1$-cycle (again writing $a_n = a_0$). If, for every i ($0 \leq i < n$), either $\mathfrak{A} \models \mathfrak{t}_2[a_i, a_{i+1}]$ or $\mathfrak{A} \models \mathfrak{t}_2[a_{i+1}, a_i]$, we observe that, since $\mathfrak{A} \models \Delta_0$, the same possibility holds for each i. That is, either $a_0, \ldots, a_{n-1}$ is a $\mathfrak{t}_2$-cycle or $a_{n-1}, \ldots, a_0$ is. This means that we can break both the $\mathfrak{t}_1$-cycle and the $\mathfrak{t}_2$-cycle simultaneously at the elements a_i and a_{i+1}, using the argument of Lemma 9, which works unproblematically. If, on the other hand, for any i $(0 \leq i < n)$, $\mathfrak{A} \not\models \mathfrak{t}_2[a_i, a_{i+1}]$ and $\mathfrak{A} \not\models \mathfrak{t}_2[a_{i+1}, a_i]$, then we can break the $\mathfrak{t}_1$-cycle at that point, taking $a' = a_i$, $b' = a_{i+1}$ and $\tau = \mathrm{tp}^{\mathfrak{A}}[a', b']$, again selecting dendral elements to absorb the relevant $\mathfrak{t}_1$-ray. We must show that, in doing so, we create no $\mathfrak{t}_2$-cycles. If $\mathrm{tp}[b', a']$ is an invertible θ-ray-type, then we select dendral elements a, b such that $\mathrm{tp}[a, b] = \mathrm{tp}[a', b']$. By θ-superchromaticity, $\mathrm{tp}^{\mathfrak{A}'}[a, b']$ and $\mathrm{tp}^{\mathfrak{A}'}[a', b]$ are θ-dark, and so we may set

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

exactly as in the argument of Lemma 9 (Fig. 3b).

If, on the other hand, $\mathrm{tp}[a_{i+1}, a_i] = \mathrm{tp}[b', a']$ is a non-invertible θ-ray-type, a complication arises. Following Lemma 9 (Fig. 3a), we wish to select dendral elements a, b such that $\mathrm{tp}^{\mathfrak{A}}[a] = \mathrm{tp}^{\mathfrak{A}}[a']$, $\mathrm{tp}^{\mathfrak{A}}[b] = \mathrm{tp}^{\mathfrak{A}}[b']$, and $\mathrm{tp}^{\mathfrak{A}}[a, b]$ is θ-dark, and set

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

This does indeed have the desired effect of eliminating a $\mathfrak{t}_1$-cycle without introducing any $\mathfrak{t}_2$-cycle, provided b' is not the $\mathfrak{t}_2$-mother of a, that is, provided $\mathfrak{A} \not\models \mathfrak{t}_2[b', a]$. This is important, because b might be a $\mathfrak{t}_2$-ancestor of a, in which case the assignment $\mathrm{tp}^{\mathfrak{A}'}[a, b] = \mathrm{tp}^{\mathfrak{A}}[a, b']$ would create a $\mathfrak{t}_2$-loop.

Fortunately, we can easily prevent this situation from arising. Suppose that the $\mathcal{C}^2[\downarrow_1, \downarrow_2]$-formula ψ has a dendral model $\mathfrak{A}_A$, and consider the graph (A, E) where E is the set of ordered pairs $\langle a, c \rangle$ for which there exists $b \in A$ such that either $\mathfrak{A} \models \mathfrak{t}_1[c, b]$ and $\mathfrak{A} \models \mathfrak{t}_2[b, a]$ or $\mathfrak{A} \models \mathfrak{t}_2[c, b]$ and $\mathfrak{A} \models \mathfrak{t}_1[b, a]$. Since each node has at most one mother in each forest, this graph has degree at most 2, and so can be properly 5-coloured by Lemma 12. Now let $p_1, \ldots, p_5$ be fresh predicates encoding these colours, and let ζ be a two-variable formula saying that the graph (A, E) is 5-coloured. Then, if ψ is replaced by $\psi \wedge \zeta$, it can never happen that there is a triple of elements a, b', a' with $\mathfrak{A} \models \mathfrak{t}_1[a', b']$, $\mathfrak{A} \models \mathfrak{t}_2[b', a]$ and $\mathrm{tp}^{\mathfrak{A}}[a'] = \mathrm{tp}^{\mathfrak{A}}[a]$, and the construction of Lemma 9 goes through. On the other hand, $\psi \wedge \zeta$ has a dendral model if and only if ψ has, which proves the theorem. □

This is the second part of Theorem 1. This argument does not work for three distinguished predicates $\mathfrak{t}_1$, $\mathfrak{t}_2$, and $\mathfrak{t}_3$. The reason is that a $\mathfrak{t}_1$-cycle may be composed entirely of $\mathfrak{t}_2$- and $\mathfrak{t}_3$-edges that do not form parts of $\mathfrak{t}_2$- and $\mathfrak{t}_3$-cycles. In this case, the swapping construction of Lemma 9, when used to eliminate a $\mathfrak{t}_1$ cycle, may create new $\mathfrak{t}_2$- and $\mathfrak{t}_3$-cycles. The argument then fails. It is not known whether the finite satisfiability problem for the logic $\mathcal{C}^2[\downarrow_1, \cdots, \downarrow_k]$ is decidable for any $k \geq 3$; similarly for the satisfiability problem.

3 The general satisfiability problem

The purpose of this section is to prove Theorem 2: the satisfiability problem for $\mathcal{C}^2[\downarrow]$ is in NExpTime. We proceed by reduction to the finite satisfiability problem for $\mathcal{C}^2[\downarrow]$. For the remainder of this section, we fix a $\mathcal{C}^2[\downarrow]$-formula φ in normal form (1) over a signature Σ', and let θ be the formula $\bigvee_{h=1}^{m} \beta_h \vee \mathfrak{t}(y, x)$. We take Σ to be the signature Σ' together with

$\lceil\log(2(mC+1)^2+1)\rceil+\lceil\log(3(mC+1))\rceil$ additional unary predicates. Thus, by Lemmas 5 and 6, any model of φ interpreting Σ' can be expanded to a θ-super-chromatic, $3(mC+1)$-differentiated model of φ interpreting Σ. Henceforth, all 1-types and 2-types are to be understood as 1-types and 2-types over Σ. Moreover, since the formula θ will not vary, we speak of θ-ray-types, θ-dark 2-types, θ-star-types, θ-*chromatic* structures etc. simply as *ray-types*, *dark* 2-*types*, *star-types*, *chromatic* structures etc.

Overview of the decision procedure. We reduce the general satisfiability of $\mathcal{C}^2[\downarrow]$ to its finite satisfiability. The main idea is to explore regularity in infinite models of $\mathcal{C}^2[\downarrow]$ formulas. Elements realizing star-types that occur only finitely often in the model, all their ancestors in the forest and elements absorbing rays from them are members of a finite "irregular" part called the *initial segment*. The remainder of the model is represented as a finite, cyclic graph of star-types, the *star chart*. The initial segment and rays connecting its elements with the remainder constitute a finite structure which is described by a translated $\mathcal{C}^2[\downarrow]$ formula. The translation involves reasoning about 1- and 2-types and its size is exponential in the size of an input formula. However, its effective size is again polynomial, which allows us to use the decision procedure for finite satisfiability of $\mathcal{C}^2[\downarrow]$ in the previous section as a black box.

3.1 Star charts

We begin with some technical machinery for reasoning about the star-types realized infinitely often in models of $\mathcal{C}^2[\downarrow]$-formulas. Recall that an arboreal ray-type is one containing either of the atoms $\mathfrak{t}(x,y)$ or $\mathfrak{t}(y,x)$. Because we will be particularly concerned in the sequel with the directions of $\mathfrak{t}$-edges, the following terminology will be useful. If ρ is a ray-type containing the atom $\mathfrak{t}(x,y)$, we call ρ *Boreal*, and if ρ is a non-invertible ray-type containing the atom $\mathfrak{t}(y,x)$, then we call ρ *Austral*. We use the terms *Boreal ray* and *Austral ray* in the obvious sense. Note that, since $\models \mathfrak{t}(y,x) \to \theta$, Boreal ray-types are necessarily invertible.

If σ is a star-type and ρ an invertible ray-type, we write $\sigma \rightsquigarrow \rho$ if σ emits a ray of type ρ, and $\rho \rightsquigarrow \sigma$ if σ emits a ray of type ρ^{-1} (or, as we might say: if σ *absorbs* a ray of type ρ). If σ and σ' are chromatic star-types, then there can be at most one invertible ray-type ρ such that $\sigma \rightsquigarrow \rho$ and $\rho \rightsquigarrow \sigma'$. In that case we write $\sigma \xrightarrow{\rho} \sigma'$. If σ, σ', σ'' are chromatic star-types and ρ, ρ' are distinct invertible ray-types such that $\sigma'' \xrightarrow{\rho} \sigma$ and $\sigma'' \xrightarrow{\rho'} \sigma'$, we write $\sigma'' \to (\sigma;\sigma')$. Notice that, in this case, σ, σ' and σ'' must be distinct, by the chromaticity of σ''. A *star chart* is a set Ω of chromatic star-types with the property that, if $\sigma \in \Omega$ and $\sigma \rightsquigarrow \rho$ for some invertible ray-type ρ, then there exists $\sigma' \in \Omega$ such that $\sigma \xrightarrow{\rho} \sigma'$. As we might say: star-charts absorb all the invertible ray-types they emit. A star-chart may be regarded as a directed graph in the following way: the vertices are star-types, and the edges are the *Boreal* ray-types which those star-types emit or absorb. Notice that the edges of this graph are directed by the predicate $\mathfrak{t}$, not by the directedness of the ray-types: all ray-types labelling the edges of this graph are Boreal and hence by definition invertible.

Where a star chart Ω is clear from context, we write $\sigma \Rightarrow \sigma'$ if, for some $m \geq 0$, there exists a sequence $\sigma_0,\ldots,\sigma_m$ of star-types from Ω and a sequence $\rho_0,\ldots,\rho_{m-1}$ of Boreal ray-types such that $\sigma = \sigma_0$, $\sigma' = \sigma_m$ and $\sigma_i \xrightarrow{\rho_i} \sigma_{i+1}$ for all i $(0 \leq i < m)$. Likewise, we write $\sigma \Rrightarrow \sigma'$ if either: (i) $\sigma \Rightarrow \sigma'$ and $\sigma' \Rightarrow \sigma$; or (ii) there exist $\sigma'', \sigma_a, \sigma_b \in \Sigma$ such that $\sigma \Rightarrow \sigma''$, $\sigma'' \to (\sigma_a;\sigma_b)$, $\sigma_a \Rightarrow \sigma$ and $\sigma_b \Rightarrow \sigma'$. Thus, while $\sigma \Rightarrow \sigma'$ states that there is a path in Ω from σ to σ', $\sigma \Rrightarrow \sigma'$ states that there is a path in Ω from σ to itself which proceeds either via σ', or via some star-type σ'' where the path branches to σ' (Fig. 4).

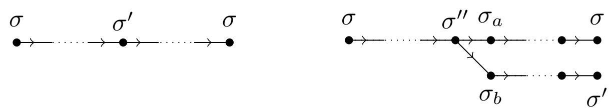

Figure 4: The relation $\sigma \Rrightarrow \sigma'$ in a star chart.

Let Ω be a star chart. A *trek* (through Ω) is a (possibly infinite) Ω-labelled tree $T = (V, E, L)$, where $L : V \to \Omega$ is the labelling function, such that, for every vertex v of T: (i) if $(v, w) \in E$, then there exists a Boreal ray-type ρ such that $L(v) \xrightarrow{\rho} L(w)$, and (ii) for every Boreal ray-type ρ emitted by $L(v)$, there exists $(v, w) \in E$ such that $L(v) \xrightarrow{\rho} L(w)$. If T is a trek and σ is the label of the root of T, we say that the *origin of* T is σ. Intuitively, a trek through Ω with origin σ is a tree of the possible routes that can be taken through the star-chart Ω starting at σ. A labelled tree satisfying only condition (i) is called a *partial trek*. A union of disjoint treks is called a *multi-trek*. Note that, for a star-chart Ω, there may be many treks with origin $\sigma \in \Omega$. This is because, if ρ is a Boreal ray-type such that $\sigma \rightsquigarrow \rho$, there may be more than one $\sigma' \in \Omega$ such that $\rho \rightsquigarrow \sigma'$. Thus, we have more than one choice as to how to unfold the trek at this point, and similarly for the Boreal ray-types emitted by whichever σ' we choose.

Lemma 13. *If Ω is a star chart and $\sigma \in \Omega$, then there exists a trek through Ω with origin σ. In fact, any partial trek through Ω can be extended to a trek.*

Lemma 14. *If Ω is a star chart, $\Xi' \subseteq \Omega$ and $\sigma \in \Omega$ such that, for all $\sigma' \in \Xi'$, $\sigma \Rrightarrow \sigma'$, then there exists a trek T through Ω in which, for every $\sigma' \in \Xi'$, there are infinitely many vertices labelled with σ'.*

Lemma 15. *Let Ω be a star chart, $\sigma, \sigma' \in \Omega$ and T a trek through Ω. Suppose B is an infinite branch of T such that (i) every vertex of B has a descendant in T labelled σ', and (ii) σ occurs infinitely often in B. Then $\sigma \Rrightarrow \sigma'$.*

Lemma 16. *Let Ω be a star chart and ρ a Boreal ray-type. Let $\Omega^\rho = \{\sigma \in \Omega \mid \rho \rightsquigarrow \sigma\}$, and $\Omega^\star_\rho = \{\sigma \in \Omega \mid$ for some $\sigma_0 \in \Omega^\rho$, $\sigma_0 \Rrightarrow \sigma\}$. If $\Omega^\rho \neq \emptyset$, then there is a trek T through Ω with origin $\sigma^* \in \Omega^\rho$ such that, for all $\sigma \in \Omega^\star_\rho$, T has infinitely many vertices labelled with σ.*

3.2 The reduction

With these preliminaries behind us, we present our promised non-deterministic exponential time Turing reduction of the satisfiability problem for $\mathcal{C}^2[\downarrow]$ to the finite the satisfiability problem for $\mathcal{C}^2[\downarrow]$.

Recall that φ is a $\mathcal{C}^2[\downarrow]$-formula of the form (1) over a signature Σ' with multiplicity m and ceiling C, that θ is the formula $\bigvee_{h=1}^{m} \beta_h \vee \mathfrak{t}(y, x)$, and that Σ is the signature of φ together with $\lceil \log(2(mC+1)^2+1) \rceil + \lceil \log(3(mC+1)) \rceil$ additional unary predicates. Recall also the formula χ_θ stating that a Σ-structure is (θ-) chromatic. Let Ω be a set of star-types and ρ an invertible ray-type. As in Lemma 16, we write Ω^ρ for the set of star-types in Ω that absorb a ray of type ρ. We write Ω° for the set of star-types in Ω that absorb no *Boreal* ray. We write $\mathrm{Inv}(\Omega)$ for the set of invertible ray-types absorbed by some star-type in Ω, and $\mathrm{Bor}(\Omega)$ for the set of Boreal ray-types absorbed by some star-type in Ω. Thus, $\mathrm{Bor}(\Omega) \subseteq \mathrm{Inv}(\Omega)$. Finally, we write $\mathrm{tp}(\Omega)$ to denote the set $\{\mathrm{tp}(\sigma) \mid \sigma \in \Omega\}$ of 1-types of the star-types in Ω. A *parameter set* for φ is a tuple $X = \langle \Omega, \Xi, \Pi_S, \Pi_N \rangle$, where Ω is a star chart, $\Xi \subseteq \Omega \setminus \Omega^\circ$, and Π_S, Π_N, disjoint sets of 1-types, satisfying the following conditions:

(I1) for all $\sigma \in \Omega$, either there exists $\sigma_0 \in \Omega^\circ$ such that $\sigma_0 \Rightarrow \sigma$, or there exists $\sigma_0 \in \Xi$ such that $\sigma_0 \Rrightarrow \sigma$;

(I2) if $\sigma \in \Omega$ and $\sigma \rightsquigarrow \rho$, then $\mathrm{tp}_2(\rho) \in \Pi_N \cup \Pi_S$;

(I3) $\mathrm{tp}(\Omega) \subseteq \Pi_N$;

(I4) if $\sigma \in \Omega$ and $\pi \in \Pi_S$, then σ emits at most one ray with absorption type π;

(I5) for all $\sigma \in \Omega$ and all $\pi \in \Pi_S$, either σ emits a non-invertible ray with absorption type π or there exists a dark 2-type τ, compatible with φ, such that τ includes both $\mathrm{tp}(\sigma)$ and $\pi(y)$;

(I6) Every star-type in Ω is compatible with $\varphi \wedge \Delta_0$ (and in particular emits at most one ray containing the atom $\mathfrak{t}(y, x)$).

By way of motivation, it helps to think of the various components of a parameter set in terms of a putative model $\mathfrak{A} \models \varphi$. Here, Ω is the star-chart consisting of those star-types that are realized infinitely often in $\mathfrak{A}$, while Π_S and Π_N are the sets of 1-types realized, respectively, uniquely and more than once, in $\mathfrak{A}$. The construction of Ξ is more complicated; however, the idea can be explained roughly as follows. Consider the graph $G = (A, \mathfrak{t}^{\mathfrak{A}})$. Now consider the subgraph H of G obtained by removing a certain finite initial segment containing all elements whose star-types are realized only finitely many times, and then dropping all Austral rays (i.e. retaining only the Boreal rays). Thus, H is also a forest—which we might call the *Boreal forest* of $\mathfrak{A}$—every component of which is therefore a tree. Note that H may have infinitely many components. The initial segment is chosen in such a way that all star-types realized by the elements of H are realized infinitely often in $\mathfrak{A}$. Of particular interest, however, are those star-types which are realized infinitely often in $\mathfrak{A}$, but in only finitely many components of H: these will need special treatment in our construction. The star-types in Ξ are the roots of these finitely many components of H.

We will be working with the signature Σ of φ together with a fresh unary predicate init. If $\mathfrak{A}$ is a structure interpreting $\Sigma \cup \{\text{init}\}$, we call the set $\text{init}^{\mathfrak{A}} \subseteq A$ the *initial segment*, and we call any ray from an element in the initial segment to an element outside the initial segment a *frontier* ray. Let $X = \langle \Omega, \Xi, \Pi_S, \Pi_N \rangle$ be a parameter set for φ, then. We define a formula $\varphi_X := \chi_\theta \wedge \psi_0 \wedge \cdots \wedge \psi_7$ over the signature $\Sigma \cup \{\text{init}\}$, where χ_θ is as in Lemma 4, and ψ_0–ψ_7, may be glossed in English as follows.

- ψ_0: Every realized 2-type is compatible with φ, and every star-type realized by an element of the initial segment is compatible with φ.
- ψ_1: For every $\rho \in \mathrm{Bor}(\Xi)$, there is a frontier ray of this type.
- ψ_2: The star-types in Ω are able to absorb all the invertible frontier rays.
- ψ_3: If $\langle a, b \rangle$ satisfy $\mathfrak{t}$ and b is in the initial segment, then so is a.
- ψ_4: Every 1-type in Π_S is uniquely realized and is realized in the initial segment.
- ψ_5: Every 1-type in Π_N is realized at least $3(mC + 1)$ times in the initial segment.
- ψ_6: Elements realizing 1-types in Π_S do not emit frontier rays.
- ψ_7: The 1-type of any element outside the initial segment is consistent with some star-type in Ω.

We then prove a pair of matching lemmas showing that φ is satisfiable if and only if there is a parameter set X such that φ_X is finitely satisfiable.

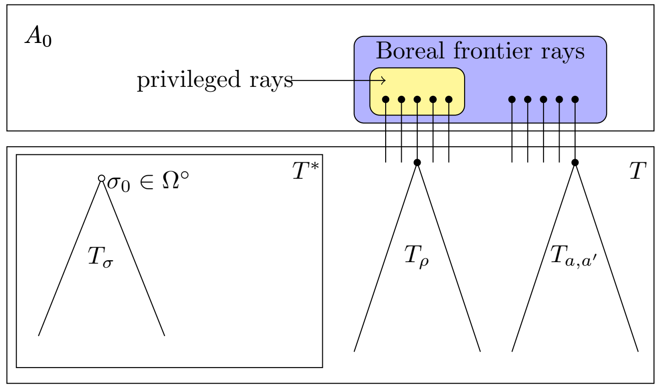

Figure 5: Constructing a model $\mathfrak{B}$ of φ from a finite model of φ_X (Lemma 17).

Lemma 17. *Let X be a parameter set for φ. If φ_X has a finite dendral model, then φ has a dendral model.*

Proof idea. Suppose $\mathfrak{A}^+$ is a finite dendral model of φ_X, where $X = \langle \Omega, \Xi, \Pi_S, \Pi_N \rangle$. We build a structure $\mathfrak{B}$ over a domain consisting of the initial segment A_0 of $\mathfrak{A}^+$ together with the vertex-set V of an infinite multi-trek T through Ω. To construct T, we take all star types $\sigma \in \Omega$ and build treks originating with star-types σ_0 (whose existence is guaranteed by condition **(I1)**) that either do not absorb Boreal rays (and thus come from Ω^0) or do absorb a Boreal ray (and thus come from Ξ). The former lead to the multi-trek T^* in Fig. 5. The latter, by formula ψ_1, are connected by Boreal frontier rays to the initial segment. We call these rays privileged; they lead to treks T_ρ in Fig. 5. The conditions **(I1)**–**(I6)** and the conjuncts ψ_0–ψ_7 ensure that the rest of the model $\mathfrak{B}$ can be built. In particular, for all remaining Boreal frontier rays we are able to take the emitting node a, add a fresh node a' to V and construct a trek $T_{a,a'}$. Over A_0, $\mathfrak{B}$ is the Σ-reduct of $\mathfrak{A}^+$, and thus 1-types of elements and 2-types of element pairs in A_0 are all defined. Over V, we define $\mathfrak{B}$ in such a way that elements realize the star-types with which they are labelled in the multi-trek. In particular, condition **(I1)** ensures that every star type $\sigma \in \Omega$ is indeed realized infinitely often in $\mathfrak{B}$, as it is realized in T^* or in T_ρ, for some Boreal ray ρ. This allows to establish the required invertible (non-arboreal) ray-types between elements of T and invertible (both Boreal and non-arboreal) ray-types between elements of A_0 and T. Austral ray-types between T and A_0 are all originating in T^* and can be assigned an absorbtion site in A_0, as every 1-type in $\Pi_S \cup \Pi_N$ is realized there. Non-invertible non-arboreal ray-types between A_0 and T can be set using the well-known cyclic construction. The remaining pairs of elements of $\mathfrak{B}$ can be assigned dark 2-types. All 2-types assigned as well as star-types of $\mathfrak{B}$ are compatible with φ.

□

Lemma 18. *If φ has a dendral model, then there exists a parameter set X such that φ_X has a finite, dendral model.*

Proof idea. Suppose $\mathfrak{A}$ is a dendral model of φ. We select a finite subset A_0 consisting (roughly) of those elements of A whose star-types are realized only finitely often, together with all their ancestors in the graph $(A, \mathfrak{t}^{\mathfrak{A}})$. This subset is the initial segment of the constructed model $\mathfrak{B}$. Let Ω be the set of star-types realized infinitely often in $\mathfrak{A}$. We define Ξ to be the set of those star-types in Ω that connect elements of $A \setminus A_0$ with elements of A_0 by Boreal ray-types. If Ξ

defined in this way does not satisfy **(I1)**, we recover this property by an appropriate expansion of A_0. Let Π_S and Π_N be the sets of 1-types realized in $\mathfrak{A}$, respectively, once and more than once. We put $X = \langle \Omega, \Xi, \Pi_S, \Pi_N \rangle$. We then build a finite dendral model $\mathfrak{B}$ of φ_X over a domain consisting of A_0 together with finitely many representatives of the star-types in Ω. $\square$

Lemma 19. *Let X be a parameter set for φ. If φ_X has a finite dendral model, then there exists a parameter set X' such that $\varphi_{X'} = \varphi_X$ and X' is exponential in $|\varphi|$.*

Proof idea. Let $X = \langle \Omega, \Xi, \Pi_S, \Pi_N \rangle$. We begin with some definitions. A *junction* is any star-type λ such that 1) there are at most four θ-ray-types ρ such that $\lambda \rightsquigarrow \rho$, 2) among them there are at most two Boreal ray-types, at most one ray-type whose inverse is a Boreal ray-type and at most one other (i.e. non-arboreal) invertible ray-type. A junction is intended to represent some selected subset of information that an ordinary star-type captures. The information in question concerns: (i) an invertible ray-type connecting a node to its mother in a tree (a type whose inverse is a Boreal ray-type), provided that the mother exists, (ii) Boreal ray-types connecting the node to one or two of its daughters, and (iii) a ray-type connecting the node to another node in the structure. If σ is any star-type compatible with φ and λ is a junction, we say that σ is an *expansion of* λ if 1) for every ρ such that $\lambda \rightsquigarrow \rho$ we have $\sigma \rightsquigarrow \rho$ and 2) if $\sigma \rightsquigarrow \tau$ where τ^{-1} is a Boreal θ-ray-type then $\lambda \rightsquigarrow \tau$. Note that there are at most exponentially many junctions over Σ, in contrast to the number of ordinary star-types compatible with φ, which is doubly exponential in $|\varphi|$.

If λ is a junction and Ψ is a set of star-types, define $\Psi_\lambda = \{\sigma \in \Psi \mid \sigma \text{ is an expansion of } \lambda\}$. Recalling now the sets of star-types Ξ and Ω featured in X, from each non-empty Ξ_λ, where λ is a junction, select one representative and define Ξ' as the set of all these representatives (as λ varies over all junctions); similarly, from each non-empty Ω_λ select one representative and and define Ω' as the set consisting of all these representatives and of all elements of Ξ'. Define $X' = \langle \Omega', \Xi', \Pi_S, \Pi_N \rangle$. As there are exponentially many junctions, both Ω' and Ξ' are exponential in $|\varphi|$. Thus X' is also exponential in $|\varphi|$. It is a routine to show that X' is a parameter set and $\varphi_{X'} = \varphi_X$. $\square$

This yields the sought-after procedure, `SatF`(φ), for determining the satisfiability of a normal-form $\mathcal{C}^2[\downarrow]$-formula φ. The procedure `SatF`(φ) consists of two steps:

1. Non-deterministically guess a parameter set X of size exponential in $|\varphi|$, and compute the formula φ_X in weak normal form.

2. Run `FinSatF`(φ_X) and report the result.

Recall that multiplicity of φ is m and ceiling of φ is C. Examination of the construction of φ_X establishes:

Lemma 20. *We can non-deterministically compute a $\mathcal{C}^2[\downarrow]$-formula φ_X in weak normal-form with multiplicity m and ceiling C over a signature Σ_X, such that: (i) φ has a dendral model if and only if, for some run of the computation, φ_X has a finite dendral model; (ii) $|\varphi_X|$ is $O\left(|\varphi| \cdot |\Sigma| \cdot 2^{O(|\Sigma|)}\right)$; (iii) $|\Sigma_X|$ is $O(|\Sigma| + \log(mC))$. The computation of φ_X requires time polynomial in $|\varphi_X|$.*

Lemmas 1, 10, 19, 20 and Theorem 3 imply Theorem 2.

References

[1] Bartosz Bednarczyk, Witold Charatonik, and Emanuel Kieronski. Extending two-variable logic on trees. In Valentin Goranko and Mads Dam, editors, *26th EACSL Annual Conference on Computer Science Logic, CSL 2017, August 20-24, 2017, Stockholm, Sweden*, volume 82 of *LIPIcs*, pages 11:1–11:20. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 2017. URL: https://doi.org/10.4230/LIPIcs.CSL.2017.11, doi:10.4230/LIPIcs.CSL.2017.11.

[2] Saguy Benaim, Michael Benedikt, Witold Charatonik, Emanuel Kieronski, Rastislav Lenhardt, Filip Mazowiecki, and James Worrell. Complexity of two-variable logic on finite trees. *ACM Trans. Comput. Log.*, 17(4):32:1–32:38, 2016. URL: http://dl.acm.org/citation.cfm?id=2996796.

[3] Witold Charatonik and Piotr Witkowski. Two-variable logic with counting and a linear order. *Logical Methods in Computer Science*, 12(2), 2016. Extended abstract in CSL 2015. URL: https://doi.org/10.2168/LMCS-12(2:8)2016, doi:10.2168/LMCS-12(2:8)2016.

[4] Witold Charatonik and Piotr Witkowski. Two-variable logic with counting and trees. *ACM Trans. Comput. Log.*, 17(4):31:1–31:27, 2016. Extended abstract in LICS 2013. URL: http://dl.acm.org/citation.cfm?id=2983622.

[5] E. Grädel, P. Kolaitis, and M. Vardi. On the decision problem for two-variable first-order logic. *Bulletin of Symbolic Logic*, 3(1):53–69, 1997.

[6] E. Grädel, M. Otto, and E. Rosen. Two-variable logic with counting is decidable. In *Logic in Computer Science*, pages 306–317. IEEE, 1997.

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

[8] L. Pacholski, W. Szwast, and L. Tendera. Complexity of two-variable logic with counting. In *Logic in Computer Science*, pages 318–327. IEEE, 1997.

[9] I. Pratt-Hartmann. Complexity of the two-variable fragment with counting quantifiers. *Journal of Logic, Language and Information*, 14(3):369–395, 2005.

[10] Ian Pratt-Hartmann. The two-variable fragment with counting revisited. In A. Dawar and R. de Queiroz, editors, *WoLLIC*, number 6188 in LNAI. Springer, 2010.

[11] Ian Pratt-Hartmann. Logics with counting and equivalence. In Thomas A. Henzinger and Dale Miller, editors, *Joint Meeting of the Twenty-Third EACSL Annual Conference on Computer Science Logic (CSL) and the Twenty-Ninth Annual ACM/IEEE Symposium on Logic in Computer Science (LICS), CSL-LICS '14, Vienna, Austria, July 14 - 18, 2014*, pages 76:1–76:10. ACM, 2014. URL: http://doi.acm.org/10.1145/2603088.2603117, doi:10.1145/2603088.2603117.

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