Logics with Counting and Equivalence

Ian Pratt-Hartmann

School of Computer Science, University of Manchester
Institute of Mathematics and Computer Science, Opole University
ipratt@cs.man.ac.uk

Abstract

We consider the two-variable fragment of first-order logic with counting, subject to the stipulation that a single distinguished binary predicate be interpreted as an equivalence. We show that the satisfiability and finite satisfiability problems for this logic are both NExpTime-complete. We further show that the corresponding problems for two-variable first-order logic with counting and *two* equivalences are both undecidable.

Categories and Subject Descriptors F.4.1 [*Computational logic*]

General Terms Theory

Keywords Equivalence relation, Satisfiability, Complexity

1. Introduction

The two-variable fragment of first-order logic, denoted $\mathcal{L}^2$, is the set of function-free, first-order formulas (with equality) featuring at most two variables. The two-variable fragment with counting, denoted $\mathcal{C}^2$, is the set of function-free, first-order formulas featuring at most two variables, but with the counting quantifiers $\exists_{[\leq M]}$, $\exists_{[\geq M]}$ and $\exists_{[=M]}$ ($M \geq 0$) allowed. It is impossible, in either logic, to express the fact that a given binary relation is an equivalence (i.e. is reflexive, symmetric and transitive). This suggests the possibility of adding such a facility. We denote by $\mathcal{L}^2k$E the extension of $\mathcal{L}^2$ in which $k \geq 1$ distinguished binary predicates are required to be interpreted as equivalences; and we denote by $\mathcal{C}^2k$E the analogous extension of $\mathcal{C}^2$.

For any logic $\mathcal{L}$, the *satisfiability problem* is the problem of determining whether, given a formula φ of $\mathcal{L}$, there exists a structure in which φ is satisfied; the *finite satisfiability problem* is the problem of determining whether there exists a finite structure in which φ is satisfied. We say $\mathcal{L}$ has the *finite model property* if these problems coincide. The following facts are known: $\mathcal{L}^2$ has the finite model property, and its satisfiability (= finite satisfiability) problem is NExpTime-complete [3, 12]; $\mathcal{C}^2$ lacks the finite model property, and its satisfiability and finite satisfiability problems are both NExpTime-complete [4, 13, 14]; $\mathcal{L}^2$1E retains the finite model property, and its satisfiability problem remains NExpTime-complete [9]; $\mathcal{L}^2$2E lacks the finite model property, and its satisfiability and finite satisfiability problems are both 2-NExpTime-complete [11]; the satisfiability and finite satisfiability problems for $\mathcal{L}^2k$E ($k \geq 3$) are both undecidable [9]. In this paper, we investigate $\mathcal{C}^2$1E—the two variable fragment with counting and one equivalence, and $\mathcal{C}^2$2E—the two variable fragment with counting and two equivalences. We show that the satisfiability and finite satisfiability problems for $\mathcal{C}^2$1E are both NExpTime-complete. We also show that the satisfiability and finite satisfiability problems for $\mathcal{C}^2$2E are both undecidable. Note that the undecidability of the corresponding problems for $\mathcal{C}^2k$E where $k \geq 3$ follows anyway from the above-mentioned results on $\mathcal{L}^2k$E.

A related family of logics is obtained by considering *transitive relations* in place of equivalences. We denote by $\mathcal{L}^2k$T the extension of $\mathcal{L}^2$ in which k distinguished binary predicates are required to be interpreted as transitive relations, and similarly for $\mathcal{C}^2k$T. It is easy to show that $\mathcal{L}^2$1T lacks the finite model property; and it is known (but not easy to show) that its satisfiability problem is in 2-NExpTime-time [17]. (The best known lower bound is 2-ExpTime-hard, and decidability of finite satisfiability remains open.) The satisfiability and finite satisfiability problems for $\mathcal{L}^2k$T ($k \geq 2$) are undecidable [6, 8]. (In fact, the corresponding problems for the weaker two-variable fragment with one equivalence and one transitive relation are also both undecidable [10].) However, the satisfiability and finite satisfiability problems for $\mathcal{C}^2k$T are both undecidable for all $k \geq 1$ [7, 18]. Decidability can be restored—even in the presence of an arbitrary number of transitive relations—by restricting the underlying logical syntax, and a great variety of such languages have been studied under the rubric of *description logics*; these logics will not be investigated here.

Of some historical interest in this connection is the first-order theory of k equivalence relations. Here, we have full-first-order logic at our disposal (not just $\mathcal{C}^2$), but no non-logical predicates other than those denoting equivalences. It is reported in [5] that membership of a sentence in the first-order theory of one equivalence is decidable (even with equality); however, the first-order theory of two equivalences is undecidable (even without equality).

The structure of the paper is as follows. Section 2 establishes basic concepts and notation. Section 3 shows how, given a formula of $\mathcal{C}^2$1E, a *certificate* can be constructed which, on the assumption that φ is finitely satisfiable, is guaranteed to satisfy a collection of algorithmically checkable properties. Section 4 establishes the converse: if a certificate for φ satisfies these properties, then φ is finitely satisfiable. Section 5 establishes that the required properties of certificates can be checked in nondeterministic exponential time, thus proving that the finite satisfiability problem for $\mathcal{C}^2$1E is NExpTime-complete; in addition, we outline how our proof can be adapted to deal with the satisfiability problem for $\mathcal{C}^2$1E. Section 6 shows that the satisfiability and finite satisfiability problems for $\mathcal{C}^2$2E are undecidable.

CSL-LICS 2014, July 14–18, 2014, Vienna, Austria.
Copyright © 2014 ACM 978-1-4503-2886-9/14/07…$15.00.
http://dx.doi.org/10.1145/2603088.2603117

2. Preliminaries

We employ standard logical notation and terminology throughout (see, e.g. [1]); however, we allow structures to be empty. The two-variable fragment with counting, denoted $\mathcal{C}^2$, is the set of function-free, first-order formulas featuring only the variables x and y, but with the counting quantifiers $\exists_{[\leq M]}$, $\exists_{[\geq M]}$ and $\exists_{[=M]}$ allowed. Formally, the subscripts M are bit-strings; however, we equivocate in the natural way between these bit-strings and the non-negative integers they encode. We read $\exists_{[\leq M]}x.\varphi$ as "There exist at most M x such that φ", and similarly for the other counting quantifiers. The formal semantics are as expected. The two-variable fragment with counting and one equivalence, $\mathcal{C}^2$1E, employs the same syntax and semantics as $\mathcal{C}^2$, but with the restriction that, in any structure $\mathfrak{A}$, the distinguished binary predicate E be interpreted as an equivalence. Where $\mathfrak{A}$ is clear from context, we refer to the cliques of $E^{\mathfrak{A}}$ as *equivalence classes*. The two-variable fragment with counting and two equivalences, $\mathcal{C}^2$2E, employs the same syntax and semantics as $\mathcal{C}^2$, but with the restriction that, in any structure $\mathfrak{A}$, the distinguished binary predicates E_1 and E_2 be interpreted as equivalences.

We allow equality in formulas; this represents no increase in expressive power, since identity is anyway definable by the formula $\forall x.r(x,x) \wedge \forall x \exists_{[=1]} y.r(x,y)$. We do not allow individual constants; this represents no effective decrease in expressive power, since we can always declare a unary predicate p to be uniquely instantiated by writing $\exists_{[=1]}x.p(x)$. Likewise, the use of predicates of arity greater than two adds no effective increase in expressive power, and we therefore assume all predicates are unary or binary.

We write $\mathbb{N}$ for the non-negative integers, and $[m,n] = \{k \mid m \leq k \leq n\}$. We assume all entries in matrices to be integers (possibly negative), denoting the (i,j)th entry of a matrix A by $A[i,j]$; similarly for vectors. A matrix or vector is (*absolutely*) *bounded by* a number M if (the modulus of) each of its entries is at most M. If U and V are sets of vectors, $U \oplus V = \{\underline{u} + \underline{v} \mid \underline{u} \in U,\ \underline{v} \in V\}$. Matrices and vectors (occasionally scalars) that it is helpful to think of as constants are typically in bold type. We write systems of linear inequalities in matrix form: $\mathbf{A}\underline{w} \leq \underline{\mathbf{b}}$, with solutions sought over $\mathbb{N}$. If $\mathcal{E}$ is such a system, we refer to the elements of $\mathbf{A}$ as *variable coefficients* of $\mathcal{E}$ and the elements of $\mathbf{b}$ as *constant coefficients* of $\mathcal{E}$; a *coefficient* is an element of either $\mathbf{A}$ or $\mathbf{b}$. We write $\|\mathcal{E}\|$ to denote the *size of* $\mathcal{E}$, i.e. the total number of bits required to write all its coefficients; and we write $|\mathcal{E}|$ to denote the *cardinality* of $\mathcal{E}$, i.e. the number or rows in $\mathbf{A}$. It is a standard result of integer linear programming (see, e.g. [15, Ch. 16]), that the set of solutions of $\mathcal{E}$ has the form $\mathbf{W} \oplus \left\{\sum_{\ell=1}^{L} \zeta_\ell \underline{\mathbf{w}}_\ell \mid \zeta_1, \ldots, \zeta_L \in \mathbb{N}\right\}$, where $\mathbf{W}$ is a finite set of vectors and $\underline{\mathbf{w}}_1, \ldots, \underline{\mathbf{w}}_L$ a list of vectors. The ensemble $\mathbf{W}$, $\underline{\mathbf{w}}_1, \ldots, \underline{\mathbf{w}}_L$ is known as a *Hilbert basis*. It is routine to check that all of the vectors involved in this Hilbert basis may be assumed to be absolutely bounded by an exponential function of $\|\mathcal{E}\|$.

If $\underline{w}$ is a solution of $\mathcal{E}$, the *footprint* of $\underline{w}$ is the set of variables taking non-zero values. The following result on footprints of solutions of systems of linear inequalities will be used on several occasions in the sequel.

Proposition 1 ([2], Theorem 2)**.** *Let $\mathcal{E}$ be a system of n inequalities with integer coefficients such that the absolute value of any variable coefficient of $\mathcal{E}$ is bounded by $N > 0$. If $\mathcal{E}$ has a solution over $\mathbb{N}$, then it has a solution over $\mathbb{N}$ with footprint of size at most $2n\log(4nN)$.*

Note that this bound is independent of the constant coefficients of $\mathcal{E}$.

2.1 Normal forms

A formula of $\mathcal{C}^2$1E is in *normal form* if it conforms to the pattern

$$\forall x \forall y (x = y \vee \alpha) \wedge \bigwedge_{\ell=1}^{m} \forall x \exists_{[=M_\ell]} y (\beta_\ell \wedge x \neq y) \qquad (1)$$

where α and the β_ℓ are quantifier-free, equality-free $\mathcal{L}^2$-formulas, $m > 0$, and the M_ℓ are (bit-strings representing) positive integers. We call $M = \max\{M_\ell \mid 1 \leq \ell \leq m\}$ the *ceiling* of φ. The following lemma uses a technique originally employed by [16] in the context of $\mathcal{L}^2$.

Lemma 2. *Given a $\mathcal{C}^2$1E-formula φ, we can compute, in polynomial time, a formula ψ, with ceiling M, such that, for any set A of cardinality greater than M, ψ is satisfiable over A if and only if φ is.*

Note that the formula ψ in Lemma 2 may in general feature a larger signature than φ.

In the sequel, we fix a normal-form $\mathcal{C}^2$1E-formula φ, and consider the problem of determining the existence of (finite) models of φ. We use the symbols α, m, M_ℓ, β_ℓ throughout to refer to the parts of φ as indicated in (1), and we additionally define M to be the ceiling of φ. We denote the number of symbols occurring in φ by $\|\varphi\|$, it being understood that a counting subscript M_ℓ contributes $\lceil \log M_\ell \rceil$ symbols. Henceforth, let

$$Z = \max(3mM + 1, (mM + 1)^2 + 1), \qquad (2)$$

and fix Σ to be the signature of φ together with $(\|\varphi\| + 5\lceil \log Z \rceil)$ fresh unary predicates. (We also assume Σ features at least two binary predicates other than E.) Since φ is fixed in the sequel, we refer to any quantity bounded by $p(\|\varphi\|)$, where p is a fixed polynomial, as *polynomially bounded*, or simply *polynomial*. Similar for *singly exponentially bounded* ($2^{p(\|\varphi\|)}$) and *doubly exponentially bounded* ($2^{2^{p(\|\varphi\|)}}$). Thus, $|\Sigma|$ is polynomial, while M and Z are singly exponential.

2.2 Rays, chromaticity and differentiation

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 $x \neq y$. Here, consistency is understood to take account of the requirement that E is interpreted as an equivalence: every 1-type contains $E(x,x)$; every 2-type contains $E(x,x)$ and $E(y,y)$; and every 2-type contains $E(x,y)$ if and only if it contains $E(y,x)$. 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 $\text{tp}_1(\tau)$ the 1-type obtained by removing from τ any literals containing y; and we write $\text{tp}_2(\tau) = \text{tp}_1(\tau^{-1})$. We equivocate freely between finite sets of formulas and their conjunctions.

Let $\mathfrak{A}$ be any structure interpreting Σ. If $a \in A$, there exists a unique 1-type $\pi(x)$ such that $\mathfrak{A} \models \pi[a]$; we denote π by $\text{tp}^{\mathfrak{A}}[a]$. If, in addition, $b \in A \setminus \{a\}$, there exists a unique 2-type $\tau(x,y)$ such that $\mathfrak{A} \models \tau[a,b]$; we denote τ by $\text{tp}^{\mathfrak{A}}[a,b]$. Evidently, $\tau^{-1} = \text{tp}^{\mathfrak{A}}[b,a]$; $\text{tp}_1(\tau) = \text{tp}^{\mathfrak{A}}[a]$; and $\text{tp}_2(\tau) = \text{tp}^{\mathfrak{A}}[b]$. If π is a 1-type, we say that π is *realized* in $\mathfrak{A}$ if there exists $a \in A$ with $\text{tp}^{\mathfrak{A}}[a] = \pi$; similarly for 2-types.

Recalling the form (1) of φ, we say that a 2-type τ is *compatible with* φ if $\models \tau \rightarrow (\alpha(x,y) \wedge \alpha(y,x))$. Thus, in any model of φ, all realized 2-types are compatible with φ. We call the 2-type τ *galactic* if it contains $E(x,y)$, and *cosmic* otherwise, i.e. if it contains $\neg E(x,y)$. For any 2-type τ, τ is galactic (cosmic) if and only if τ^{-1} is. We call the 2-type τ a *ray-type* if $\models \tau \rightarrow \beta_\ell$ for some ℓ ($1 \leq \ell \leq m$). If ρ is a ray-type such that ρ^{-1} is also a ray-type, we say that ρ is *invertible*. A ray-type ρ is *polarized* if it is either non-invertible or $\mathrm{tp}_1(\rho) \neq \mathrm{tp}_2(\rho)$. If ρ is a polarized, invertible cosmic ray-type, we refer to the unordered pair (ρ, ρ^{-1}) as a *symmetrized* cosmic ray-type. (We do not require a corresponding notion for other sorts of ray-types.) If τ is a 2-type such that neither τ nor τ^{-1} is a ray-type, we say that τ is *dark*.

The above terminology is supposed to suggest the following imagery. If $\mathrm{tp}^{\mathfrak{A}}[a, b]$ is a ray-type ρ, then we may imagine that a emits a ray, of type ρ, that is absorbed by b. If ρ is invertible, then b reciprocates (with a ray of type ρ^{-1}). Accordingly, we refer to the 1-types $\mathrm{tp}_1(\rho)$ and $\mathrm{tp}_2(\rho)$ as the *emission-type* and *absorption-type* of ρ, respectively. If $\mathrm{tp}^{\mathfrak{A}}[a, b]$ is dark, then neither element emits a ray that is absorbed by the other.

We say that $\mathfrak{A}$ is *polarized* if every ray realized in $\mathfrak{A}$ is polarized. We say that $\mathfrak{A}$ is *chromatic* if: (i) $\mathfrak{A}$ is polarized; and (ii) for all 1-types π and all $a \in A$, a emits at most one invertible ray with absorption-type π. It is easy to see that $\mathfrak{A}$ is chromatic if and only if no two distinct elements with the same 1-type are joined by a chain of at most two invertible rays. We say $\mathfrak{A}$ is *differentiated* if, for every 1-type π: (i) π is realized either in at most one or in at least Z equivalence classes; and (ii) π is realized in any equivalence class either at most once or at least Z times. Using the $(\|\varphi\| + 5\lceil \log Z \rceil)$ unary predicates of Σ that do not appear in φ, we can show:

Lemma 3. *If φ has a model interpreting Σ, then φ has a chromatic, differentiated model interpreting Σ over the same domain.*

Let π and π' be 1-types, not necessarily distinct. Recalling the form (1) of φ, define γ to be the formula $(\alpha(x, y) \wedge \alpha(y, x) \wedge \pi(x) \wedge \pi'(y)) \rightarrow \bigvee_{\ell=1}^{m} (\beta_\ell(x, y) \vee \beta_\ell(y, x))$. We say that π and π' are *galactically coupled*, and write $\pi \overset{g}{\sim} \pi'$, if $\models E(x, y) \wedge x \neq y \rightarrow \gamma$; and we say that π and π' are *cosmically coupled*, and write $\pi \overset{c}{\sim} \pi'$, if $\models \neg E(x, y) \rightarrow \gamma$. Galactic and cosmic coupling are important for the following reason. Suppose $\mathfrak{A} \models \varphi$, and a, b are distinct but equivalent elements of A such that $\mathrm{tp}^{\mathfrak{A}}[a] = \pi$ and $\mathrm{tp}^{\mathfrak{A}}[b] = \pi'$. If $\pi \overset{g}{\sim} \pi'$, then either $\mathrm{tp}^{\mathfrak{A}}[a, b]$ or $\mathrm{tp}^{\mathfrak{A}}[b, a]$ (possibly both) is a galactic ray-type. Similarly, suppose a, b are non-equivalent elements of A such that $\mathrm{tp}^{\mathfrak{A}}[a] = \pi$ and $\mathrm{tp}^{\mathfrak{A}}[b] = \pi'$. If $\pi \overset{c}{\sim} \pi'$, then either $\mathrm{tp}^{\mathfrak{A}}[a, b]$ or $\mathrm{tp}^{\mathfrak{A}}[b, a]$ (possibly both) is a cosmic ray-type.

2.3 Star-types and enumerations

Let us enumerate the 1-types as $\pi_1, \ldots, \pi_I$. We fix this enumeration for the remainder of this paper. Let us enumerate the polarized ray-types as $\rho_1, \ldots, \rho_{8J}$. We may choose the enumeration so that $\rho_1, \ldots, \rho_{2J}$ are all galactic and invertible, $\rho_{2J+1}, \ldots, \rho_{4J}$ are all galactic and non-invertible, $\rho_{4J+1}, \ldots, \rho_{6J}$ are all cosmic and invertible, and $\rho_{6J+1}, \ldots, \rho_{8J}$ are all cosmic and non-invertible. We need not worry that there are more invertible than non-invertible polarized ray-types: just 'pad out' the latter with unrealized dummies. Since these rays are all polarized, we may unproblematically stipulate that, for $j \in [1, J] \cup [4J + 1, 5J]$, $\rho_j^{-1} = \rho_{J+j}$. Thus, invertible ray-types (galactic or cosmic) and their inverses are enumerated in parallel. We fix this enumeration for the remainder of this paper.

A *star-type* is a pair $\sigma = \langle \pi, (v_1, \ldots, v_{8J}) \rangle$ where π is a 1-type and the v_j are non-negative integers such that $v_j > 0$ implies $\mathrm{tp}_1(\rho_j) = \pi$. We write, $\mathrm{tp}(\sigma) = \pi$ and, abusing vector notation slightly, $\sigma[j] = v_j$. Informally, we think of a star-type σ as a finite multiset over the list of polarized ray-types $\rho_1, \ldots, \rho_{8J}$; and we speak of its elements as *rays emitted by* σ.

We say σ is *chromatic* if it emits no two invertible rays with the same absorption-type: i.e. if for every 1-type π, $\sum\{v_j : 1 \leq j \leq 2J,\ \mathrm{tp}_2(\rho_j) = \pi\} + \sum\{v_j : 4J < j \leq 6J,\ \mathrm{tp}_2(\rho_j) = \pi\} \leq 1$. We say σ is *compatible with* φ if: (i) for all j $(1 \leq j \leq 8J)$, $\sigma[j] > 0 \Rightarrow \rho_j$ is compatible with φ; and (ii) for all ℓ $(1 \leq \ell \leq m)$, $\sum\{\sigma[j] \mid 1 \leq j \leq 8J,\ \models \rho_j \rightarrow \beta_\ell\} = M_\ell$. Thus, σ is compatible with φ if it emits no rays forbidden by $\forall x \forall y(x = y \vee \alpha)$ and the right numbers of rays required by the $\forall x \exists_{[=M_\ell]} y(\beta \wedge x \neq y)$. If $\mathfrak{A}$ is a finite, polarized structure interpreting Σ, and $a \in A$, define $\mathrm{st}^{\mathfrak{A}}[a] = \langle \mathrm{tp}^{\mathfrak{A}}[a], (v_1, \ldots, v_{8J}) \rangle$, where $v_j = |\{b \in A : b \neq a \text{ and } \mathrm{tp}^{\mathfrak{A}}[a, b] = \rho_j\}|$. Evidently, $\mathrm{st}^{\mathfrak{A}}[a]$ is a star-type, and we call it the *star-type of* a in $\mathfrak{A}$. If $\sigma = \mathrm{st}^{\mathfrak{A}}[a]$ for some $a \in A$, we say σ is realized in $\mathfrak{A}$. It is routine to check: $\mathfrak{A}$ is chromatic if and only if every star-type realized in $\mathfrak{A}$ is chromatic; and $\mathfrak{A} \models \varphi$ if and only if every star-type and every dark 2-type realized in $\mathfrak{A}$ is compatible with φ.

3. From models to certificates

Suppose the $\mathcal{C}^2$1E-formula φ given in (1) has a finite, chromatic, differentiated model $\mathfrak{A}$ interpreting Σ. Any $a \in A$ thus emits at most M rays of any given type. Enumerate the star-types realized in $\mathfrak{A}$ as $\sigma_1, \ldots, \sigma_K$; we fix this enumeration for the remainder of Sec. 3. Although this list of star-types depends on $\mathfrak{A}$, K is bounded as a fixed doubly exponential function of $\|(\varphi)\|$. We proceed to describe the construction of a *certificate* for φ.

3.1 Special and ordinary equivalence classes

For all i $(1 \leq i \leq I)$, execute the following procedure. If π_i is realized in at least Z equivalence classes, select Z of those equivalence classes. If, on the other hand, π_i is realized in just one equivalence class B, select B, and if, in addition, π is realized by exactly one element a of B, also select every equivalence class B' containing any b such that $\mathrm{tp}^{\mathfrak{A}}[a, b]$ is a cosmic ray-type. Call an equivalence class *special* if it is selected in this process. An equivalence-class that is not special is *ordinary*, and an element is *special* (*ordinary*) if its equivalence class is. Let $A^\dagger$ be the set of special elements, and A^* the set of ordinary elements. Thus, $A = A^\dagger \cup A^*$, and $A^\dagger \neq \emptyset$.

Enumerate the special equivalence classes as $B^1, \ldots, B^G$. Thus, G is (positive and) singly exponentially bounded. Define $\mathcal{I} = \{i \mid \pi_i \text{ is realized exactly once in } \mathfrak{A}\}$. For all i $(1 \leq i \leq I)$ define $\mathcal{G}_i = \{g \in [1, G] \mid \pi_i \text{ is realized at least once in } B^g\}$.

Consider the following sets of statements.

$$\{(|\mathcal{G}_i| \leq 1) \text{ or } (|\mathcal{G}_i| \geq Z) \mid 1 \leq i \leq I\} \qquad (\mathcal{B}_1)$$

$$\{\mathcal{G}_i \text{ is a singleton} \mid i \in \mathcal{I}\} \qquad (\mathcal{B}_2)$$

$$\begin{aligned}&\{(\mathcal{G}_i = \emptyset) \text{ or } (\mathcal{G}_{i'} = \emptyset) \text{ or} \\ &\quad (\mathcal{G}_i = \mathcal{G}_{i'} \text{ and } |\mathcal{G}_i| = 1) \mid i, i' \in [1, I] \setminus \mathcal{I},\ \pi_i \overset{c}{\sim} \pi_{i'}\}. \qquad (\mathcal{B}_3)\end{aligned}$$

We write $\mathcal{B} = \mathcal{B}_1 \cup \mathcal{B}_2 \cup \mathcal{B}_3$.

Lemma 4. *All the statements in $\mathcal{B}$ are true.*

Proof. Statement $\mathcal{B}_1$ follows from the selection of the special equivalence classes together with the fact that $\mathfrak{A}$ is differentiated. Statement $\mathcal{B}_2$ follows from the definition of $\mathcal{I}$. Statement $\mathcal{B}_3$ follows—via a straightforward combinatorial argument—from the definition of cosmic coupling together with the fact that $\mathfrak{A}$ is differentiated. □

3.2 Profiles of equivalence classes

Let $A' \subseteq A$. The *profile* of A' is the vector $\mathrm{pr}^{\mathfrak{A}}[A'] = (w_1, \ldots, w_K)$, where $w_k = |\{a \in A' : \mathrm{st}^{\mathfrak{A}}[a] = \sigma_k\}|$. Thus, the profile of A' gives the numbers of elements in A' realizing each of the star-types $\sigma_1, \ldots, \sigma_K$. We observed above that the length K of this vector is doubly exponential in $\|(\varphi)\|$. Intuitively, we may think of $\mathrm{pr}^{\mathfrak{A}}[A']$ as a *histogram* of A', summarizing the set in statistical terms. Of primary interest in the sequel will be the case where A' is an equivalence class or a union of equivalence classes.

Recall the enumerations $\{\pi_i\}_1^I$, $\{\rho_j\}_1^{8J}$ and $\{\sigma_k\}_1^K$. For all i, k ($1 \leq i \leq I, 1 \leq k \leq K$), define the constant $\mathbf{p}_{i,k}$ to be 1 if $\mathrm{tp}(\sigma_k) = \pi_i$, and 0 otherwise. Thus, the equation $\mathbf{p}_{i,k} = 1$ states that any element with star-type σ_k has 1-type π_i. For all j, k ($1 \leq j \leq 4J, 1 \leq k \leq K$), let $\mathbf{u}_{j,k} = \sigma_k[j]$. Thus, $\mathbf{u}_{j,k}$ gives the number of rays of (galactic) type ρ_j emitted by any element having star-type σ_k. Finally, for all i, i', k, c ($1 \leq i, i' \leq I$, $1 \leq k \leq K, 1 \leq c \leq mM+1$), let the constants $\mathbf{o}^c_{i,i',k}$ and $\underline{\mathbf{o}}^*_{i,i',k}$ with values in $\{0,1\}$ be defined in such a way that: $\mathbf{o}^c_{i,i',k} = 1$ if and only if $\mathrm{tp}(\sigma_k) = \pi_i$ and σ_k emits at least c galactic rays with absorption-type $\pi_{i'}$; and $\mathbf{o}^*_{i,i',k} = 1$ if and only if $\mathrm{tp}(\sigma_k) = \pi_i$ and σ_k emits no *non-invertible* galactic rays with absorption-type $\pi_{i'}$. Let $\mathbf{q}^c_{i,i',k}$ and $\mathbf{q}^*_{i,i',k}$ be defined analogously, but with "galactic ray" replaced by "cosmic ray". For convenience, we collect those constants whose indices differ only in the value k into vectors of length K, thus: $\underline{\mathbf{p}}_i = (\mathbf{p}_{i,1}, \ldots, \mathbf{p}_{i,K})$, $\underline{\mathbf{u}}_j = (\mathbf{u}_{j,1}, \ldots, \mathbf{u}_{j,K})$, $\underline{\mathbf{o}}^c_{i,i'} = (\mathbf{o}^c_{i,i',1}, \ldots, \mathbf{o}^c_{i,i',K})$, and similarly for $\underline{\mathbf{q}}^d_{i,i'}$, $\underline{\mathbf{o}}^*_{i,i'}$ and $\underline{\mathbf{q}}^*_{i,i'}$. These constants allow us to compute various statistics concerning any subset of $A' \subseteq A$ from its profile. For example, suppose $\mathrm{pr}^{\mathfrak{A}}[A'] = \underline{w}$. Then the number of elements of A' having 1-type π_i ($1 \leq i \leq I$) is $\underline{\mathbf{p}}_i \cdot \underline{w}$. If, in addition, A' is the union of some collection of equivalence classes, then the total number of rays of galactic type ρ_j ($1 \leq j \leq 4J$) emitted (and therefore absorbed) by elements of A' is $\underline{\mathbf{u}}_j \cdot \underline{w}$.

Now let $\underline{w} = (w_1, \ldots, w_K)$ be a tuple of variables, and consider the following sets of statements regarding $\underline{w}$.

$$\{\underline{\mathbf{u}}_j \cdot \underline{w} = \underline{\mathbf{u}}_{J+j} \cdot \underline{w} \mid 1 \leq j \leq J\} \qquad (\mathcal{C}_1^0)$$

$$\{(\underline{\mathbf{p}}_i \cdot \underline{w} \leq 1) \vee (\underline{\mathbf{p}}_i \cdot \underline{w} \geq Z) \mid 1 \leq i \leq I\} \qquad (C_2^0)$$

$$\{(\underline{\mathbf{p}}_i \cdot \underline{w} \geq c) \vee (\underline{\mathbf{o}}^c_{i',i} \cdot \underline{w} = 0) \mid 1 \leq i, i' \leq I,\ c = 1, 2\} \quad (\mathcal{C}_3^0)$$

$$\{(\underline{\mathbf{p}}_i \cdot \underline{w} > 1) \vee (\underline{\mathbf{o}}^c_{i,i'} \cdot \underline{w} = 0) \vee (\underline{\mathbf{o}}^*_{i',i} \cdot \underline{w} \geq c) \mid \\ 1 \leq i, i' \leq I,\ 1 \leq c \leq mM\} \quad (\mathcal{C}_4^0)$$

$$\{(\underline{\mathbf{p}}_i \cdot \underline{w} = 0) \vee (\underline{\mathbf{o}}^*_{i',i} \cdot \underline{w} < c) \vee (\underline{\mathbf{o}}^c_{i,i'} \cdot \underline{w} \geq 1) \mid \\ 1 \leq i, i' \leq I,\ \pi_i \overset{g}{\sim} \pi_{i'},\ c \leq mM + 1\} \quad (\mathcal{C}_5^0)$$

$$\{\underline{\mathbf{p}}_i \cdot \underline{w} \leq 1) \mid i \in \mathcal{I}\} \qquad (\mathcal{C}_6^0)$$

$$\{\underline{\mathbf{q}}^2_{i',i} \cdot \underline{w} = 0) \mid i \in \mathcal{I},\ 1 \leq i' \leq I\} \qquad (\mathcal{C}_7^0)$$

$$\{\underline{\mathbf{q}}^1_{i',i} \cdot \underline{w} = 0) \mid 1 \leq i, i' \leq I, \mathcal{G}_i = \emptyset\} \qquad (\mathcal{C}_8^0)$$

$$\{(\underline{\mathbf{p}}_i \cdot \underline{w} \leq 1) \vee (\underline{\mathbf{p}}_{i'} \cdot \underline{w} \leq 1) \mid 1 \leq i, i' \leq I,\ \pi_i \overset{g}{\sim} \pi_{i'}\}. \quad (\mathcal{C}_9^0)$$

We write $\mathcal{C}^0 = \mathcal{C}_1^0 \cup \cdots \cup \mathcal{C}_9^0$.

Lemma 5. *Suppose B is an equivalence class. Then $\mathrm{pr}^{\mathfrak{A}}[B]$ satisfies $\mathcal{C}^0(\underline{w})$.*

Proof. For $\mathcal{C}_1^0$, observe that, since B is an equivalence class, the total number of rays of (invertible, galactic) type ρ_j ($1 \leq j \leq J$) emitted by elements of B equals the total number of rays of type $\rho_j^{-1} = \rho_{J+j}$ emitted by B. For $\mathcal{C}_2^0$, observe that, since $\mathfrak{A}$ is differentiated, the number of elements of 1-type π_i is either at most 1 or at least Z. The rest are similar. □

Turning our attention now to the *ordinary* equivalence classes, consider the following sets of equations in $\underline{w}$.

$$\{\underline{\mathbf{p}}_i \cdot \underline{w} = 0 \mid 1 \leq i \leq I,\ |\mathcal{G}_i| \leq 1\} \qquad (\mathcal{C}_1^*)$$

$$\{\underline{\mathbf{q}}^*_{i',i} \cdot \underline{w} = 0 \mid i \in \mathcal{I},\ 1 \leq i' \leq I,\ \pi_i \overset{c}{\sim} \pi_{i'}\} \qquad (\mathcal{C}_2^*)$$

We write $\mathcal{C}^* = \mathcal{C}^0 \cup \mathcal{C}_1^* \cup \mathcal{C}_2^*$. We see that $|\mathcal{C}^*|$ —that is, the *number* of statements in $\mathcal{C}^*$—is singly exponentially bounded.

Lemma 6. *Suppose B is an ordinary equivalence class. Then $\mathrm{pr}^{\mathfrak{A}}[B]$ satisfies $\mathcal{C}^*(\underline{w})$.*

Proof. For $\mathcal{C}_1^*$, fix i ($1 \leq i \leq I$). By construction of the special equivalence classes, if π_i is realized in exactly one special equivalence class, then it is realized in no ordinary equivalence classes. For $\mathcal{C}_2^*$, fix $i \in \mathcal{I}$ and i' ($1 \leq i' \leq I$). By construction of $\mathcal{I}$, π_i is realized exactly once in $\mathfrak{A}$, and so let a be the element realizing π_i. By definition, a and all elements to which a sends any cosmic rays lie in special equivalence classes. Hence, if b, realizing 1-type $\pi_{i'}$, lies in an ordinary equivalence class, and $\mathrm{tp}^{\mathfrak{A}}[a, b]$ is not dark, then b sends at least one non-invertible cosmic ray to an element of type π_i. □

For all g ($1 \leq g \leq G$), let $\underline{w}^g$ be a K-tuple of fresh variables. Write $\underline{w}^\dagger$ to denote the (KG)-tuple $\underline{w}^1, \ldots, \underline{w}^G$, and write $\underline{1}$ for the vector $(1, \ldots, 1)$ of length K. For any g, consider the following sets of statements regarding $\underline{w}^\dagger$:

$$\{\underline{\mathbf{p}}_i \cdot \underline{w}^g = 0 \mid 1 \leq i \leq I \text{ and } g \notin \mathcal{G}_i\} \qquad (\mathcal{C}_1^g)$$

$$\{\underline{\mathbf{p}}_i \cdot \underline{w}^g \geq 1 \mid 1 \leq i \leq I \text{ and } g \in \mathcal{G}_i\} \qquad (\mathcal{C}_2^g)$$

$$\{\underline{\mathbf{p}}_i \cdot \underline{w}^g \geq 2 \mid 1 \leq i \leq I,\ \mathcal{G}_i = \{g\} \text{ and } i \notin \mathcal{I}\} \qquad (\mathcal{C}_3^g)$$

$$\{\underline{1} \cdot \underline{w}^g \geq 1\} \qquad (\mathcal{C}_4^g)$$

$$\{\underline{\mathbf{q}}^1_{i',i} \cdot \underline{w}^g = 0 \mid 1 \leq i, i' \leq I,\ \mathcal{G}_i = \{g\}\} \qquad (\mathcal{C}_5^g)$$

$$\{(\underline{\mathbf{q}}^d_{i,i'} \cdot \underline{w}^g = 0) \vee \sum_{1 \leq h \leq G}^{h \neq g} \underline{\mathbf{q}}^*_{i',i} \cdot \underline{w}^h \geq d \mid \\ i \in \mathcal{I},\ \mathcal{G}_i = \{g\},\ 1 \leq i' \leq I,\ d \leq mM\} \quad (\mathcal{C}_6^g)$$

$$\{(\underline{\mathbf{q}}^d_{i,i'} \cdot \underline{w}^g \geq 1) \vee \sum_{1 \leq h \leq G}^{h \neq g} \underline{\mathbf{q}}^*_{i',i} \cdot \underline{w}^h < d \mid \\ 1 \leq i, i' \leq I,\ g \in \mathcal{G}_i, \pi_i \overset{c}{\sim} \pi_{i'},\ d \leq mM + 1\}. \quad (\mathcal{C}_7^g)$$

Observe that $\mathcal{C}_6^g$ and $\mathcal{C}_7^g$ involve variables of $\underline{w}^\dagger$ other than those of $\underline{w}^g$. For all g ($1 \leq g \leq G$), we write $\mathcal{C}^g = \mathcal{C}_1^g \cup \cdots \cup \mathcal{C}_7^g$.

Lemma 7. *Let $B^1, \ldots, B^G$ be the special equivalence classes of $\mathfrak{A}$. Then, for all g ($1 \leq g \leq G$), the (KG)-tuple $\mathrm{pr}^{\mathfrak{A}}[B^1], \ldots, \mathrm{pr}^{\mathfrak{A}}[B^G]$ satisfies $\mathcal{C}^g(\underline{w}^\dagger)$.*

Proof. Similar to Lemmas 5 and 6. □

Now define the collection of statements $\mathcal{C}^\dagger$ by

$$\mathcal{C}^\dagger(\underline{w}^\dagger) = \bigcup_{g=1}^{G} \left(\mathcal{C}^0(\underline{w}^g) \cup \mathcal{C}^g(\underline{w}^\dagger) \right).$$

Again, we see that $|\mathcal{C}^\dagger|$ is singly exponentially bounded.

3.3 Spectra and clusters

We now group the equivalence classes of $\mathfrak{A}$ into larger units, called clusters. We employ the following apparatus for this purpose. Let $A' \subseteq A$. The *cosmic spectrum* (or *c-spectrum*) of A' is the vector $\mathrm{cs}^{\mathfrak{A}}[A'] = (v_1, \ldots, v_{4J})$, where $v_j = |\{\langle a, b\rangle \in A' \times A : b \neq a \text{ and } \mathrm{tp}^{\mathfrak{A}}[a, b] = \rho_{4J+j}\}|$. The *symmetrized c-spectrum* of A' is the vector $\mathrm{ss}^{\mathfrak{A}}[A'] = (t_1, \ldots, t_J)$, where $t_j = |\{\langle a, b\rangle \in A' \times A : b \neq a \text{ and } \mathrm{tp}^{\mathfrak{A}}[a, b] = \rho_{4J+j} \text{ or } \mathrm{tp}^{\mathfrak{A}}[a, b] = \rho_{5J+j}\}|$. Thus, the c-spectrum lists the total number of rays of each cosmic type emitted by elements of A'; similarly, the symmetrized c-spectrum lists the total number of rays of each symmetrized cosmic type emitted by elements of A'. Intuitively, we may think of these vectors as more compressed versions of the profile of A'.

We can easily calculate (symmetrized) c-spectra from profiles. Define the array $\mathbf{V}$, of dimension $(4J \times K)$, by $\mathbf{V}[j, k] = \sigma_k[4J + j]$ $(1 \le j \le 4J, 1 \le k \le K)$. Taking $\mathbf{I}$ to be the identity matrix and $\mathbf{O}$ the zero-matrix, both of dimension $(J \times J)$, define the $(J \times 4J)$-matrix $\mathbf{T} = (\mathbf{I} \mid \mathbf{I} \mid \mathbf{O} \mid \mathbf{O})$. Evidently, if the profile of A' is $\underline{w}$, then its c-spectrum and symmetrized c-spectrum are, respectively, $\mathbf{V}\underline{w}$ and $\mathbf{TV}\underline{w}$.

Now we can explain how to group equivalence classes into clusters. For the special elements $A^\dagger = B^1 \cup \cdots \cup B^G$, clustering is degenerate: we define the *special clusters* to be $C^1, \ldots, C^G$, where $C^g = B^g$ for all g $(1 \le g \le G)$. Since, for all g $(1 \le g \le G)$, $\mathrm{pr}^{\mathfrak{A}}[B^g] = \underline{w}^g$, we have:

$$\mathrm{cs}^{\mathfrak{A}}[B^g] = \mathbf{V}\underline{w}^g \qquad \mathrm{ss}^{\mathfrak{A}}[B^g] = \mathbf{T}\underline{v}^g. \tag{3}$$

Let us now consider the ordinary equivalence classes. Lemma 6 states that the profile of any ordinary equivalence class satisfies $\mathcal{C}^*(\underline{w})$. Evidently, $\mathcal{C}^*(\underline{w})$ is equivalent to a disjunction of systems of inequalities $\mathbf{A}\underline{w} \le \underline{\mathbf{b}}$. And for each of these systems, we can compute a Hilbert basis $\mathbf{W}$, $\underline{\mathbf{w}}_1, \ldots, \underline{\mathbf{w}}_L$, such that its space of solutions is $\mathbf{W} \oplus \left\{ \sum_{\ell=1}^{L} \zeta_\ell \underline{\mathbf{w}}_\ell \mid \zeta_1, \ldots, \zeta_L \in \mathbb{N} \right\}$. Thus, although there is no *a priori* bound on the number of ordinary equivalence classes in $\mathfrak{A}$, we do know that the profile of any of these equivalence classes is a linear combination of one of a bounded collection of sequences of vectors obtained from these Hilbert bases. This observation plays a fundamental role in our decision procedure.

In order to obtain a tight complexity bound, we need to exercise care in choosing these Hilbert bases. Using Proposition 1, we can find integers $H \ge G$, $L \ge 1$ (both doubly exponentially bounded) and a partition of A^* into sets $C^{G+1}, \ldots, C^H$, such that, for all h $(G < h \le H)$, there exist a matrix $\mathbf{A}^h$, a vector $\mathbf{b}^h$, and an L-tuple of vectors $\mathbf{w}_0^h, \ldots, \mathbf{w}_L^h$ satisfying the following properties:

(i) C^h is a union of equivalence classes, and for each of these equivalence classes, B, there exist non-negative integers $\zeta_1, \ldots \zeta_L$ (depending on B) such that, writing $\underline{\mathbf{v}}_\ell^h = \mathbf{V}\underline{\mathbf{w}}_\ell^h$ and $\underline{\mathbf{t}}_\ell^h = \mathbf{TV}\underline{\mathbf{w}}_\ell^h$,

$$\mathrm{cs}^{\mathfrak{A}}[B] = \underline{\mathbf{v}}_0^h + \zeta_1 \underline{\mathbf{v}}_1^h + \cdots + \zeta_L \underline{\mathbf{v}}_L^h \tag{4}$$

$$\mathrm{ss}^{\mathfrak{A}}[B] = \underline{\mathbf{t}}_0^h + \zeta_1 \underline{\mathbf{t}}_1^h + \cdots + \zeta_L \underline{\mathbf{t}}_L^h; \tag{5}$$

(ii) the vectors $\mathbf{w}_\ell^h$ $(0 \le \ell \le L)$ satisfy

$$\mathbf{A}^h \underline{\mathbf{w}}_0^h \le \underline{b}^h \tag{6}$$

$$\mathbf{A}^h \underline{\mathbf{w}}_\ell^h \le \underline{0} \quad \text{for all } \ell \, (1 \le \ell \le L); \tag{7}$$

(iii) the system of linear inequalities $\mathbf{A}^h \underline{w} \le \underline{\mathbf{b}}^h$ propositionally entails $\mathcal{C}^*(\underline{w})$;

(iv) for some fixed integer N (singly exponentially bounded), the matrix $\mathbf{A}^h$ has dimension $N \times K$ and is singly exponentially absolutely bounded; the vector $\mathbf{b}^h$ has length N, and is singly exponentially absolutely bounded; the vectors $\mathbf{w}_\ell^h$ $(0 \le \ell \le L)$ are doubly exponentially absolutely bounded, of length K, but with singly exponential footprint.

We call the sets C^h $(G < h \le H)$ the *ordinary clusters* of $\mathfrak{A}$. Thus, we have partitioned $\mathfrak{A}$ into the special clusters, $C^1, \ldots, C^G$, each comprising exactly one equivalence class, together with the ordinary clusters, $C^{G+1}, \ldots, C^H$, each comprising some non-empty collection of equivalence classes. This partition is illustrated in Fig. 1.

To understand the significance of the partition of Fig. 1, fix some ordinary cluster C^h $(G < h \le H)$, and consider its c-spectrum. Evidently, $\mathrm{cs}^{\mathfrak{A}}[C^h] = \sum\{\mathrm{cs}^{\mathfrak{A}}[B] \mid B \subseteq C^h\}$, whence, from (4), $\mathrm{cs}^{\mathfrak{A}}[C^h]$ has the form

$$z_0^h \underline{\mathbf{v}}_0^h + z_1^h \underline{\mathbf{v}}_1^h + \cdots + z_L^h \underline{\mathbf{v}}_L^h.$$

where z_0^h is the number of equivalence classes in $\mathcal{C}^h$, and $z_\ell^h \in \mathbb{N}$ $(1 \le \ell \le L)$. Pictorially, we may imagine each equivalence class $B \subseteq C^h$ to be composed of various groups of elements, or 'constellations': a single 'core constellation' having c-spectrum $\underline{\mathbf{v}}_0^h$, and, for each ℓ $(1 \le \ell \le L)$, some number (possibly zero) of 'peripheral constellations' each having c-spectrum $\underline{\mathbf{v}}_\ell^h$. The numbers z_ℓ^h $(0 \le \ell \le L)$ are simply the totals obtained by summing over all B. The key to our approach is that—subject to a caveat to be discussed in Sec. 3.4—we do not particularly mind how the various peripheral constellations are distributed between the equivalence classes in C^h: all that matters is the total number of constellations of each type, as given by the parameters z_ℓ^h $(G < h \le H, 0 \le \ell \le L)$. And, while we have no *a priori* bound on the number of ordinary equivalence classes, we do have such a bound on L and H.

We conclude with some motivating remarks concerning equations (6) and (7). Consider again any equivalence class $B \subseteq C^h$, and define the vector

$$\underline{w} = \underline{\mathbf{w}}_0^h + \zeta_1 \underline{\mathbf{w}}_1^h + \cdots + \zeta_L \underline{\mathbf{w}}_L^h, \tag{8}$$

where the coefficients $\zeta_1, \ldots, \zeta_L$ are those satisfying equations (4)–(5). We know from (6) and (7) that $\underline{w}$ is a solution of $\mathbf{A}^h \underline{w} \le \underline{\mathbf{b}}^h$, and hence satisfies $\mathcal{C}^*(\underline{w})$; moreover equations (4)–(5), ensure that $\mathbf{V}\underline{w} = \mathrm{cs}^{\mathfrak{A}}[B]$ and $\mathbf{TV}\underline{w} = \mathrm{ss}^{\mathfrak{A}}[B]$. Intuitively, we may think of $\underline{w}$ as an *ersatz* profile for B: a vector which satisfies $\mathcal{C}^*(\underline{w})$ (i.e. the conditions required for a vector to be the profile of an ordinary equivalence class), and which yields the correct c-spectrum on application of $\mathbf{V}$. Note that there is no guarantee that $\underline{w}$ is equal to $\mathrm{pr}^{\mathfrak{A}}[B]$; however, as far as the rest of the structure is concerned, B behaves as if this equation held. We remark that it is crucial to the subsequent proof that every $\mathbf{w}_\ell^h$ is doubly exponentially absolutely bounded, and has singly exponential footprint.

3.4 Sectors and terminators

It turns out that clusters are not quite sufficient for our purposes, and in this section, we corral the equivalence classes in every ordinary cluster C^h into an alternating sequence of groups which we refer to as sectors and terminators. If $H = G$, there are no ordinary clusters, and so nothing to do. Therefore, we may assume $H > G \ge 1$. Denoting the number of equivalence classes in C^h by $\mathfrak{c}(h)$, let us enumerate them as $B_1^h, \ldots, B_{\mathfrak{c}(h)}^h$. From (8), for all h, s $(G < h \le H, 1 \le s \le \mathfrak{c}(h))$, B_s^h has an ersatz profile

$$\zeta_{s,0}^h \underline{\mathbf{w}}_0^h + \zeta_{s,1}^h \underline{\mathbf{w}}_1^h + \cdots + \zeta_{s,L}^h \underline{\mathbf{w}}_L^h. \tag{9}$$

where $\zeta_{s,0}^h = 1$ and $\zeta_{s,1}^h, \ldots, \zeta_{s,L}^h \in \mathbb{N}$. Indeed:

$$\mathrm{cs}^{\mathfrak{A}}[B_s^h] = \zeta_{s,0}^h \underline{\mathbf{v}}_0^h + \zeta_{s,1}^h \underline{\mathbf{v}}_1^h + \cdots + \zeta_{s,L}^h \underline{\mathbf{v}}_L^h \tag{10}$$

$$\mathrm{ss}^{\mathfrak{A}}[B_s^h] = \zeta_{s,0}^h \underline{\mathbf{t}}_0^h + \zeta_{s,1}^h \underline{\mathbf{t}}_1^h + \cdots + \zeta_{s,L}^h \underline{\mathbf{t}}_L^h. \tag{11}$$

To reduce notational clutter, for all h and s $(1 \le h \le H$, $1 \le s \le \mathfrak{c}(h))$, we write $\underline{t}_s^h$ for $\mathrm{ss}^{\mathfrak{A}}[B_s^h]$, i.e., the symmetrized c-spectrum of the sth equivalence class in the hth cluster. Now fix some j $(1 \le j \le J)$: thus, ρ_{4J+j} is an invertible cosmic ray-type, and $(\rho_{4J+j}, \rho_{4J+j}^{-1}) = (\rho_{4J+j}, \rho_{5J+j})$ is a symmetrized cosmic ray-type. Observe that, for every h $(1 \le h \le H)$, and every s $(1 \le s \le \mathfrak{c}(h))$ the elements of B_s^h cannot possibly emit more rays of symmetrized cosmic type $(\rho_{4J+j}, \rho_{5J+j})$ than all the rest of $\mathfrak{A}$ put together, that is to say:

$$\underline{t}_s^h[j] \le \sum \left\{ \underline{t}_{s'}^{h'}[j] \;\middle|\; \begin{array}{l} 1 \le h' \le H,\, 1 \le s' \le \mathfrak{c}(h'), \\ (h, s) \ne (h', s') \end{array} \right\}. \tag{12}$$

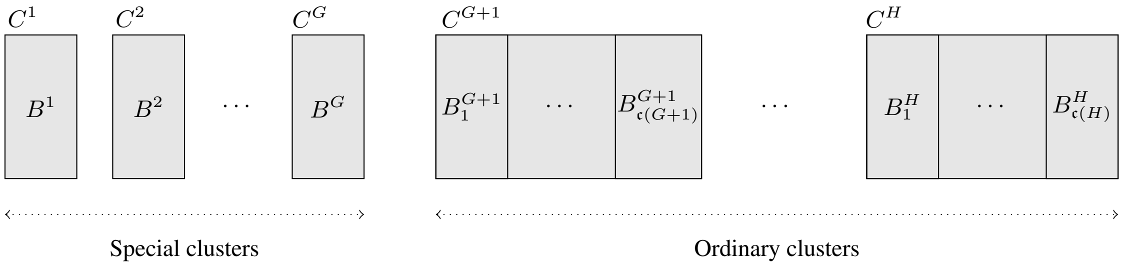

Figure 1: The partition of equivalence classes into clusters.

As we might put it: no equivalence class has an absolute majority in respect of any symmetrized cosmic ray-type.

Let us write $(h, s) \prec (h', s')$ if either $h < h'$, or both $h = h'$ and $s < s'$. Thus, $\prec$ is simply the lexicographic ordering of indices, and enumerates the equivalence classes of $\mathfrak{A}$ left-to-right as depicted in Fig. 1. We define the symbols $\preceq$, $\succ$ and $\succeq$ in the expected way, and we always compare pairs of integers (h, s) with reference to this ordering, employing terms such as *greatest*, *maximal* etc., with the obvious meaning. Continuing to fix j ($1 \leq j \leq J$), define $(\mathfrak{h}(j), \mathfrak{s}(j))$ to be the greatest index pair $(1, 1) \prec (\mathfrak{h}(j), \mathfrak{s}(j)) \preceq (H, \mathfrak{c}(H))$, satisfying

$$\sum\{\underline{t}_h^s[j] \mid (h, s) \prec (\mathfrak{h}(j), \mathfrak{s}(j))\} \leq \sum\{\underline{t}_h^s[j] \mid (h, s) \succeq (\mathfrak{h}(j), \mathfrak{s}(j))\}. \quad (13)$$

To see that $(\mathfrak{h}(j), \mathfrak{s}(j))$ exists, note that, by assumption, $H \geq 2$, and put $h = s = 1$ in (12). We claim that

$$\sum\{\underline{t}_h^s[j] \mid (h, s) \preceq (\mathfrak{h}(j), \mathfrak{s}(j))\} \geq \sum\{\underline{t}_h^s[j] \mid (h, s) \succ (\mathfrak{h}(j), \mathfrak{s}(j))\}. \quad (14)$$

If $(\mathfrak{h}(j), \mathfrak{s}(j)) = (H, \mathfrak{c}(H))$, then (14) is trivial; if $(\mathfrak{h}(j), \mathfrak{s}(j))$ is the immediate $\prec$-predecessor of $(H, \mathfrak{c}(H))$, then it follows by putting $h = H$ and $s = \mathfrak{c}(H)$ in (12); otherwise, it follows from the maximality of $(\mathfrak{h}(j), \mathfrak{s}(j))$. Furthermore, putting $h = \mathfrak{h}(j)$ and $s = \mathfrak{s}(j)$ in (12), we have

$$\underline{t}_{\mathfrak{h}(j)}^{\mathfrak{s}(j)}[j] \leq \sum\{\underline{t}_h^s[j] \mid (h, s) \prec (\mathfrak{h}(j), \mathfrak{s}(j))\} + \sum\{\underline{t}_h^s[j] \mid (h, s) \succ (\mathfrak{h}(j), \mathfrak{s}(j))\}. \quad (15)$$

Let this construction be carried out for all j ($1 \leq j \leq J$). Inequalities (13)—(15) will play a key role in constructing the certificate for φ—once we have re-organized them slightly. To this end, fix h ($1 \leq h \leq H$), and let

$$S^h = \{\mathfrak{s}(j) \mid 1 \leq j \leq J, \; \mathfrak{h}(j) = h\} \cup \{\mathfrak{c}(h)\}.$$

Thus, S^h records those indices $\mathfrak{s}(j)$ (with j varying), for which $B_{\mathfrak{s}(j)}^{\mathfrak{h}(j)}$ is included in C^h, and adds in the final index $\mathfrak{c}(h)$. We remark that $\mathfrak{c}(h)$ may be the only element of S^h. Let $\mathfrak{b}(h) = |S^h|$, and enumerate S^h as a strictly increasing sequence of integers $s_1 \leq \cdots \leq s_{\mathfrak{b}(h)}$. Thus, $1 \leq \mathfrak{b}(h) \leq J + 1$, and $s_{\mathfrak{b}(h)} = \mathfrak{c}(h)$. By definition, $\mathfrak{s}(j)$ must be one of the elements s_p in the enumeration $s_1, \ldots, s_{\mathfrak{b}(h)}$ of $S^{\mathfrak{h}(j)}$; and we write $\mathfrak{p}(j) = p$ to identify the index of this element. The functions

$$\mathfrak{h} : [1, J] \to [1, H] \qquad \mathfrak{p} : [1, J] \to [1, J + 1]$$

will form part of the certificate for φ. Notice that $\mathfrak{p}(j) \leq \mathfrak{b}(\mathfrak{h}(j))$. Keeping h fixed, and writing $s_0 = 0$, define, for all p ($1 \leq p \leq \mathfrak{b}(h)$),

$$\begin{aligned} \dot{B}_p^h &= B_{s_p}^h \\ \hat{B}_p^h &= \bigcup\{B_s^h \mid s_{p-1} < s < s_p\}. \end{aligned}$$

We refer to the $\dot{B}_p^h$ as *terminators*, and to the $\hat{B}_p^h$ as *sectors*. This internal organization of clusters is illustrated in Fig. 2.

Now consider the c-spectra and symmetrized c-spectra of these sectors and terminators. For all h, p ($1 \leq h \leq H$, $1 \leq p \leq \mathfrak{b}(h)$), let

$$\begin{aligned} \dot{\underline{v}}_p^g &= \mathrm{cs}^{\mathfrak{A}}[\dot{B}_p^h] & \qquad \dot{\underline{t}}_p^g &= \mathrm{ss}^{\mathfrak{A}}[\dot{B}_p^h] \\ \hat{\underline{v}}_p^g &= \mathrm{cs}^{\mathfrak{A}}[\hat{B}_p^h] & \qquad \hat{\underline{t}}_p^g &= \mathrm{ss}^{\mathfrak{A}}[\hat{B}_p^h]. \end{aligned}$$

We proceed to write arithmetic expressions for these quantities. Again, we adopt different strategies for the special clusters and ordinary clusters. The special clusters involve no new work. By construction, if $1 \leq h \leq G$, then $\dot{B}_1^h = B_1^h$ and $\dot{B}_1^h = \emptyset$; consulting (3), therefore, we have, for all g ($1 \leq g \leq G$):

$$\{\dot{\underline{v}}_1^g = \mathbf{V}\underline{w}^g, \; \dot{\underline{t}}_1^g = \mathbf{T}\dot{\underline{v}}_1^g \mid 1 \leq g \leq G\} \quad (\mathcal{D}_1)$$

$$\{\hat{\underline{v}}_1^g = \underline{0}, \; \hat{\underline{t}}_1^g = \underline{0} \mid 1 \leq g \leq G\}. \quad (\mathcal{D}_2)$$

The ordinary clusters require more careful treatment. Fixing h ($G < h \leq H$), recall the expressions (10) and (11) giving the c-spectrum and symmetrized c-spectrum of any equivalence class B_s^h ($1 \leq s \leq \mathfrak{s}(h)$). Recalling the enumeration $s_1 \leq \cdots \leq s_{\mathfrak{b}(h)}$ of S^h (where the s_p depend on h), we see that the c-spectrum of $\dot{B}_p^h$ is simply the c-spectrum of $B_{s_p}^h$; moreover, the c-spectrum of $\hat{B}_p^h$ is the sum of the c-spectra of the equivalence classes B_s^h such that $B_s^h \subseteq \hat{B}_p^h$. (See Fig. 2.) Corresponding remarks apply in the case of symmetrized c-spectra. Let us therefore write

$$\begin{aligned} \dot{z}_{p,\ell}^h &= \zeta_{s_p,\ell}^h \\ \hat{z}_{p,\ell}^h &= \sum\{\zeta_{s,\ell}^h \mid s_{p-1} < s < s_p\}. \end{aligned}$$

It follows that, for all h ($G < h \leq H$) and all p ($1 \leq p \leq \mathfrak{b}(h)$), the following sets of equations hold:

$$\{\dot{\underline{v}}_p^h = \Sigma_{\ell=0}^L \dot{z}_{p,\ell}^h \underline{\mathbf{v}}_\ell^h \mid G < h \leq H, \; p \leq \mathfrak{b}(h)\} \quad (\mathcal{D}_3)$$

$$\{\hat{\underline{v}}_p^h = \Sigma_{\ell=0}^L \hat{z}_{p,\ell}^h \underline{\mathbf{v}}_\ell^h \mid G < h \leq H, \; p \leq \mathfrak{b}(h)\} \quad (\mathcal{D}_4)$$

$$\{\dot{\underline{t}}_p^h = \Sigma_{\ell=0}^L \dot{z}_{p,\ell}^h \underline{\mathbf{t}}_\ell^h \mid G < h \leq H, \; p \leq \mathfrak{b}(h)\} \quad (\mathcal{D}_5)$$

$$\{\hat{\underline{t}}_p^h = \Sigma_{\ell=0}^L \hat{z}_{p,\ell}^h \underline{\mathbf{t}}_\ell^h \mid G < h \leq H, \; p \leq \mathfrak{b}(h)\}. \quad (\mathcal{D}_6)$$

Observe that, with the aid of $\mathcal{D}_1$–$\mathcal{D}_6$, we have expressed the c-spectrum and symmetrized c-spectrum of every sector and every terminator in terms of the various parameters $\underline{w}^g$, $\dot{z}_{p,\ell}^h$ and $\hat{z}_{p,\ell}^h$.

Since, in (9), we set $\zeta_{s,0}^h = 1$, it follows that that $\dot{z}_{p,0}^h = 1$, while $\hat{z}_{p,0}^h$ equals the number of equivalence classes included in the

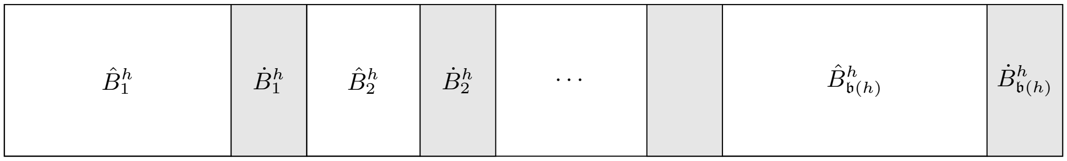

Figure 2: The division of C^h into sectors $\hat{B}^h_p$ and terminators $\dot{B}^h_p$.

sector $\hat{B}^h_s$. (The condition $\hat{z}^h_{p,0} = 0$ means that the sector $\hat{B}^h_p$ is empty.) Thus, the following sets of equations hold:

$$\{\dot{z}^h_{p,0} = 1 \mid G < h \le H,\ 1 \le p \le \mathfrak{b}(h)\} \tag{$\mathcal{E}_1$}$$

$$\{(\hat{z}^h_{p,0} = 0) \to \sum_{\ell=1}^{L} \hat{z}^h_{p,\ell} = 0 \mid G < h \le H,\ 1 \le p \le \mathfrak{b}(h)\}. \tag{$\mathcal{E}_2$}$$

Now let us express the c-spectrum and symmetrized c-spectrum of every cluster (special or ordinary) in terms of the various parameters $\underline{w}^g$, $\dot{z}^h_{p,\ell}$ and $\hat{z}^h_{p,\ell}$. For all h ($1 \le h \le H$), let

$$\underline{v}^h = \mathrm{cs}^{\mathfrak{A}}[C^h] \qquad \underline{t}^h = \mathrm{ss}^{\mathfrak{A}}[C^h].$$

Then we have

$$\{\underline{v}^h = \sum_{p=1}^{\mathfrak{b}(h)} \left(\underline{\hat{v}}^h_p + \underline{\dot{v}}^h_p\right) \mid 1 \le h \le H\} \tag{$\mathcal{D}_7$}$$

$$\{\underline{t}^h = \sum_{p=1}^{\mathfrak{b}(h)} \left(\underline{\hat{t}}^h_p + \underline{\dot{t}}^h_p\right) \mid 1 \le h \le H\}. \tag{$\mathcal{D}_8$}$$

And of course we may write analogous expressions for the c-spectrum and symmetrized c-spectrum of the entire structure. Defining

$$\underline{v} = \mathrm{cs}^{\mathfrak{A}}[A] \qquad \underline{t} = \mathrm{ss}^{\mathfrak{A}}[A].$$

we have:

$$\underline{v} = \sum_{h=1}^{H} \underline{v}^h \qquad \underline{t} = \sum_{h=1}^{H} \underline{t}^h. \tag{$\mathcal{D}_9$}$$

Since each ray of invertible cosmic type ρ_{4J+j} ($1 \le j \le J$) may be paired with a ray of (distinct) inverse type, $\rho^{-1}_{4J+j} = \rho_{5J+j}$, we have the equations

$$\{\underline{v}[j] = \underline{v}[j+J] \mid 1 \le j \le J\}. \tag{$\mathcal{E}_3$}$$

It easily follows from $\mathcal{D}_1$–$\mathcal{D}_9$ that $\underline{t} = \mathbf{T}\underline{v}$, and hence from $\mathcal{E}_3$ that every entry in $\underline{t}$ is even.

Now we can complete the required re-organization of the inequalities (13)—(15). We first define three additional scalar variables for each j ($1 \le j \le J$):

$$t^-_j = \sum\{\underline{\dot{t}}^h_p \mid (h,p) \prec (\mathfrak{h}(j), \mathfrak{p}(j))\} + \sum\{\underline{\hat{t}}^h_p[j] \mid (h,p) \preceq (\mathfrak{h}(j), \mathfrak{p}(j))\} \tag{$\mathcal{D}_{10}$}$$

$$t^\circ_j = \underline{\dot{t}}^{\mathfrak{h}(j)}_{\mathfrak{p}(j)}[j] \tag{$\mathcal{D}_{11}$}$$

$$t^+_j = \sum\{\underline{\dot{t}}^h_p + \underline{\hat{t}}^h_p[j] \mid (\mathfrak{h}(j), \mathfrak{p}(j)) \prec (h,p) \preceq (H, \mathfrak{b}(H))\}. \tag{$\mathcal{D}_{12}$}$$

Thus, $\underline{t}^-_j$ is the number of rays of symmetrized type $(\rho_{4J+j}, \rho_{5J+j})$ emitted by those equivalence classes lying strictly to the left of $B^{\mathfrak{h}(j)}_{\mathfrak{p}(j)}$ in the diagram of Fig. 1; $\underline{t}^\circ_j$ is the number of rays emitted by $B^{\mathfrak{h}(j)}_{\mathfrak{p}(j)}$ itself; and $\underline{t}^+_j$ is the the number of rays emitted by those equivalence classes lying strictly to its right. Then (13)—(15) can be written as:

$$\{\underline{t}^-_j \le \underline{t}^\circ_j + \underline{t}^+_j \mid 1 \le j \le J\} \tag{$\mathcal{E}_4$}$$

$$\{\underline{t}^+_j \le \underline{t}^-_j + \underline{t}^\circ_j \mid 1 \le j \le J\} \tag{$\mathcal{E}_5$}$$

$$\{\underline{t}^\circ_j \le \underline{t}^-_j + \underline{t}^+_j \mid 1 \le j \le J\}. \tag{$\mathcal{E}_6$}$$

The significance of $\mathcal{E}_4$–$\mathcal{E}_6$ is that they constitute a succinct guarantee that, for each j ($1 \le j \le J$), no sector or terminator—and hence certainly no equivalence class—has an absolute majority in respect of the symmetrized cosmic ray-type $(\rho_{4J+j}, \rho_{5J+j})$.

Let us gather together the above sets of equations, writing:

$$\mathcal{D} = \mathcal{D}_1 \cup \cdots \cup \mathcal{D}_{12}$$
$$\mathcal{E} = \mathcal{E}_1 \cup \cdots \cup \mathcal{E}_6.$$

Regarding the equations in $\mathcal{D}$ as definitions of the expressions on their left-hand sides, we can view the equations in $\mathcal{E}$ as constraints on the vector variable $\underline{w}^\dagger$ and the integer variables $\dot{z}^h_{p,\ell}$ and $\hat{z}^h_{p,\ell}$ ($G < h \le H$, $1 \le \ell \le L$, $1 \le p \le \mathfrak{b}(h)$). Note that $|\mathcal{E}_1|$ and $|\mathcal{E}_2|$ are doubly exponentially bounded, and $|\mathcal{E}_3|, \ldots, |\mathcal{E}_6|$ are singly exponentially bounded.

We let $\underline{\dot{z}}$ stand for the tuple of all integers $\dot{z}^h_{p,\ell}$ in some order, and similarly for $\underline{\hat{z}}$. Given a formula φ of the form (1), let $\Sigma, \pi_1, \ldots, \pi_I$ and $\rho_1, \ldots, \rho_{8J}$, be as described in Sec. 2. A *certificate* (for φ) is a tuple $\mathfrak{C} = \langle G, H, K, L, \{\sigma_k\}, \mathcal{I}, \{\mathcal{G}_i\}, \mathfrak{b}, \mathfrak{h}, \mathfrak{p}, \{\mathbf{A}^h\}, \{\underline{\mathbf{b}}^h\}, \{\underline{\mathbf{w}}^h_\ell\}, \{\underline{w}^g\}, \{\dot{z}^h_{p,\ell}\}, \{\hat{z}^h_{p,\ell}\}\rangle$, as described in this section, such that the tuple $\underline{w}^\dagger = \underline{w}^1, \ldots, \underline{w}^G$ satisfies $\mathcal{C}^\dagger$, and the tuple $(\underline{w}^\dagger, \underline{\dot{z}}, \underline{\hat{z}})$ satisfies $\mathcal{E}$ (under the definitions $\mathcal{D}$). By Lemma 3, if φ has a finite model, then φ has a chromatic, differentiated, finite model interpreting Σ. In this section, we used such a model to guide the construction of a certificate for φ. We have thus shown:

Lemma 8. *Suppose φ is a $\mathcal{C}^2 1E$-formula in the form* (1). *If φ is finitely satisfiable, then φ has a certificate.*

4. From certificates to models

Suppose $\mathfrak{C}$ is a certificate for φ, and let $\{\sigma_k\}$ be the collection of star-types in $\mathfrak{C}$. Consider any multiset defined over these star-types. It will be helpful to regard this multiset as a *set*, A, in which every element is identified as an instance of a particular star-type σ_k. We call any element $a \in A$ a *star*, and we call the star-type σ_k of which it is an instance the *intrinsic star-type* of a, denoted $\mathrm{st}(a)$. If $A' \subseteq A$, we define the *intrinsic profile* of A', denoted $\mathrm{pr}(A')$, to be the vector $(w_1, \ldots, w_K)$, where $w_k = |\{a \in A' \mid \mathrm{st}(a) = \sigma_k\}|$. Again, it helps to think of a star as object emitting a set of (polarized) rays, each having some intrinsic ray-type.

Suppose now that the set of stars A forms the domain of some finite, polarized structure $\mathfrak{A}$ interpreting Σ. In this case, for all $a \in A$, the star-type $\mathrm{st}^{\mathfrak{A}}[a]$ is defined as in Sec. 2.3. Do not confuse $\mathrm{st}^{\mathfrak{A}}[a]$ with $\mathrm{st}(a)$: the former depends on the structure $\mathfrak{A}$; the latter, only on the identity of a as an element of the set A. Define a *cosmos* to be a finite polarized structure $\mathfrak{A} \models \varphi$ over a set of stars A where, for all $a \in A$, $\mathrm{st}(a) = \mathrm{st}^{\mathfrak{A}}[a]$. We proceed to construct a cosmos.

4.1 Galaxies

Let $\sigma = \langle \pi, (v_1, \ldots, v_{8J}) \rangle$ be a star-type. Recall that the ray-types $(\rho_1, \ldots, \rho_{8J})$ are divided into the galactic ray-types $(\rho_1, \ldots, \rho_{4J})$ and the cosmic ray-types $(\rho_{4J+1}, \ldots, \rho_{8J})$. Accordingly, we define the *galactic part* of σ to be the pair $\sigma_\star = \langle \pi, (v_1, \ldots, v_{4J}) \rangle$; and we speak of any such object as a *galactic star-type*. Thus, if A is a set of stars and $a \in A$, then a has an intrinsic galactic star-type $\sigma_\star(a)$. If, in addition, $\mathfrak{A}$ is a polarized structure over A interpreting Σ, then a also has a galactic star-type $\sigma_\star^{\mathfrak{A}}[a]$. Recalling the formula φ given in (1), we say that a *galaxy* is a polarized structure $\mathfrak{B}$ over a set of stars B satisfying the following properties:

(i) $E^{\mathfrak{B}}$ is the total relation $B \times B$;

(ii) for all $b \in B$, $\mathrm{st}(b)$ is compatible with φ, and $\mathrm{st}_\star^{\mathfrak{B}}[b] = \mathrm{st}_\star(b)$;

(iii) every (dark, galactic) 2-type realized in $\mathfrak{B}$ is compatible with φ.

Thus, in a galaxy, there is just one equivalence class, all *galactic* rays emitted by stars have been found landing sites, and, as far as *galactic* 2-types are concerned, the requirements of φ are satisfied.

The next lemma is, in effect, the converse of Lemma 5: it allows us to construct galaxies from sets of stars whose profiles satisfy the conditions $\mathcal{C}^0$.

Lemma 9. *Suppose* $\underline{w} = (w_1, \ldots, w_K)$ *satisfies* $\mathcal{C}^0(\underline{w})$. *Then there exists a galaxy* $\mathfrak{B}$ *such that* $\mathrm{pr}(B) = \underline{w}$.

4.2 Constructing the cosmos: the stars

Consider again the certificate $\mathfrak{C}$. By assumption, we have $\mathcal{C}^\dagger(\underline{w}^\dagger)$, and hence, for all g $(1 \leq g \leq G)$, $\mathcal{C}^0(\underline{w}^g)$. By Lemma 9, then, let $\mathfrak{B}^g$ be a galaxy with intrinsic profile $\underline{w}^g$. By $\mathcal{C}_4^g$, $\mathfrak{B}^g$ is non-empty.

Now fix some h $(G < h \leq H)$. For any p $(1 \leq p \leq \mathfrak{b}(h))$, define

$$\begin{aligned}\underline{\dot{w}}_p^h =& \dot{z}_{p,0}^h \underline{\mathbf{w}}_0^h + \dot{z}_{p,1}^h \underline{\mathbf{w}}_1^h + \cdots + \dot{z}_{p,L}^h \underline{\mathbf{w}}_L^h \\ \underline{\hat{w}}_p^h =& \hat{z}_{p,0}^h \underline{\mathbf{w}}_0^h + \hat{z}_{p,1}^h \underline{\mathbf{w}}_1^h + \cdots + \hat{z}_{p,L}^h \underline{\mathbf{w}}_L^h.\end{aligned}$$

By $\mathcal{E}_1$, $\dot{z}_{p,0}^h = 1$, whence, by (6)—(7), $\underline{\dot{w}}_p^h$ is a solution of $\mathbf{A}^h \underline{w} \leq \underline{\mathbf{b}}^h$, and thus satisfies $\mathcal{C}^*(\underline{w})$. By Lemma 9, then, let $\dot{\mathfrak{B}}_p^h$ be a galaxy with intrinsic profile $\underline{\dot{w}}_p^h$ over a domain $\dot{B}_p^h$. Now consider $\underline{\hat{w}}_p^h$. We have two cases. If $\hat{z}_{p,0}^h = 0$, then, by $\mathcal{E}_2$, $\underline{\hat{w}}_p^h = \underline{0}$; in that case, we let $\hat{\mathfrak{B}}_p^h$ be the empty structure, with profile $\underline{0}$. Otherwise, $\hat{z}_{p,0}^h > 0$, and we proceed as follows. By (6)—(7), both $\underline{\mathbf{w}}_0^h$ and also the vector $\underline{\tilde{w}} = \underline{\mathbf{w}}_0^h + \dot{z}_{p,1}^h \underline{\mathbf{w}}_1^h + \cdots + \dot{z}_{p,L}^h \underline{\mathbf{w}}_L^h$ are solutions of $\mathbf{A}^h \underline{w} \leq \underline{\mathbf{b}}^h$, and thus satisfy $\mathcal{C}^*(\underline{w})$. By Lemma 9, then, there exist galaxies $\mathfrak{B}$, with intrinsic profile $\underline{\mathbf{w}}_0^h$, and $\tilde{\mathfrak{B}}$, with intrinsic profile $\underline{\tilde{w}}$. Now take $\hat{z}_{p,0}^h - 1$ copies of $\mathfrak{B}$ and a single copy of $\tilde{\mathfrak{B}}$ (all disjoint), and let $\hat{B}_p^h$ be the union of their domains, so that $\mathrm{pr}(\hat{B}_p^h) = \underline{\hat{w}}_p^h$. We remark in passing that this construction reflects the observation, made in Sec. 3.3, concerning the notional decomposition of any equivalence class (here: galaxy) into a core constellation (corresponding to some vector $\underline{\mathbf{w}}_0^h$) and peripheral constellations (corresponding to vectors $\underline{\mathbf{w}}_\ell^h$ for $1 \leq \ell \leq L$). We remarked there that we do not particularly mind how the peripheral constellations are distributed between the various equivalence classes. In the present construction, we have placed all peripheral constellations in a single galaxy $\tilde{\mathfrak{B}}$ within the relevant sector. Let

$$A = (B^1 \cup \cdots \cup B^G) \cup \bigcup_{h=G+1}^{H} \bigcup_{p=1}^{\mathfrak{b}(h)} (\dot{B}_p^h \cup \hat{B}_p^h).$$

This completes the definition of the domain, A.

We wish to define a model $\mathfrak{A} \models \varphi$ over A. If $\mathfrak{B}$ is any of the galaxies formed in the construction of A, then we set $\mathfrak{A}_{|B} = \mathfrak{B}$. Thus, all galactic rays emitted by the stars in A have been found absorption sites in the same galaxy as the star emitting them, and all remaining pairs of elements from the same galaxy of A have been assigned a dark galactic 2-type compatible with φ. It remains to specify the 2-types of pairs of elements from different galaxies in such a way that $\mathrm{st}^{\mathfrak{A}}[a] = \mathrm{st}(a)$ for all $a \in A$, and that all dark cosmic 2-types are compatible with φ.

We outline the construction only. Consider first the invertible cosmic rays emitted by the stars of A. By $\mathcal{E}_3$ we know that, for all j $(1 \leq j \leq J)$, the total number of such rays of type ρ_{4J+j} equals the total number of such rays of type $\rho_{5J+j} = \rho_{4J+j}^{-1}$. Thus these sets of rays can be put in 1–1 correspondence. Indeed, $\mathcal{E}_1$–$\mathcal{E}_3$ imply that no galaxy accounts for an absolute majority of the cosmic rays of symmetrized invertible type $(\rho_{4J+j}, \rho_{5J+j})$ emitted throughout the cosmos. This can be shown to be sufficient to ensure that the pairing of rays of type ρ_{4J+j} with those of type ρ_{5J+j} can in fact be chosen in such a way that rays emitted by stars in the same galaxy are never paired. In this way, if a emits a ray of invertible cosmic type ρ, paired with a ray of type ρ^{-1} emitted by b, we may legitimately set $\mathrm{tp}^{\mathfrak{A}}[a, b]$ to be the invertible cosmic ray-type ρ. It easily follows from the fact that every σ_k is chromatic that these assignments do not overwrite each other.

Consider next the non-invertible cosmic rays emitted by the stars of A. In this case, we use the statements $\mathcal{B} \cup \mathcal{C}^\dagger \cup \mathcal{C}_1^* \cup \mathcal{C}_2^*$ to show that there are sufficient stars of the appropriate types among $B^1 \cup \cdots \cup B^G$ to absorb all such rays emitted throughout the cosmos. To complete the construction, we need to define any remaining 2-types to be dark, cosmic 2-types compatible with φ. Specifically, let a and b be stars from different galaxies of A, with respective 1-types π and π'. Using the statements $\mathcal{B} \cup \mathcal{C}^\dagger \cup \mathcal{C}_1^* \cup \mathcal{C}_2^*$, it can be shown that *either* $\mathrm{tp}^{\mathfrak{A}}[a, b]$ has already been defined (as a ray-type or its inverse), or π and π' are not cosmically coupled. In the latter case, choose a dark cosmic 2-type τ compatible with φ and satisfying $\mathrm{tp}_1(\tau) = \pi$, $\mathrm{tp}_2(\tau) = \pi'$, and set $\mathrm{tp}^{\mathfrak{A}}[a, b] = \tau$.

We thus obtain a structure $\mathfrak{A}$ in which the galaxies are exactly the equivalence classes, and every dark 2-type realized in $\mathfrak{A}$ is compatible with φ. Since the star-types σ_k appearing in $\mathfrak{C}$ are all by assumption compatible with φ, we have $\mathfrak{A} \models \varphi$. That is:

Lemma 10. *Suppose φ is a $\mathcal{C}^2 1E$-formula in the form* (1). *If a certificate for φ exists, then φ is finitely satisfiable.*

5. Proof of main result

The proof of Lemma 8 constructs a certificate $\mathfrak{C}$ from a finite, chromatic, differentiated model φ. We were careful to note that each of the vectors $\underline{\mathbf{w}}_\ell^h$ is doubly exponentially absolutely bounded and has singly exponential footprint (and thus requires only exponentially many bits to write). On the other hand, the quantities H, K and L are doubly exponentially bounded. Using Proposition 1, however, it can be shown that, if a certificate exists for φ, then one can be found in which only a singly exponential number of the variables involved are non-zero. By eliminating any zero-variables, we can require H, K and L to be singly exponentially bounded.

Lemma 11. *Suppose φ is a $\mathcal{C}^2 1E$-formula in the form* (1). *If φ is finitely satisfiable, then φ has a certificate $\mathfrak{C}$ such that the total number of bits required to write $\mathfrak{C}$ is singly exponentially bounded as a function of $\|\varphi\|$.*

Theorem 12. *The finite satisfiability problem for $\mathcal{C}^2 1E$ is* NExpTime*-complete.*

Proof. The upper bound follows from Lemmas 2, 10 and 11. The lower bound follows from the fact that the (finite) satisfiability problem for the two-variable fragment of first-order logic is NExpTime-hard. □

Denote by $\mathbb{N}^*$ the set $\mathbb{N} \cup \{\aleph_0\}$. We interpret the arithmetic operations $+$ and $\cdot$ as well as the ordering $<$ over $\mathbb{N}^*$ as expected. By considering solutions of systems of linear inequalities over $\mathbb{N}^*$ rather than over $\mathbb{N}$, and making various minor adjustments to the above proof, we easily obtain

Theorem 13. *The satisfiability problem for $\mathcal{C}^2 1E$ is in* NExpTime.

6. Two equivalence relations

In this section, we show that the satisfiability and finite satisfiability problems for $\mathcal{C}^2$2E are both undecidable.

A *deterministic* 2*-counter machine* $\mathbf{M}$ has a finite set of states $s_0, \ldots, s_L$ and two counters, c_1 and c_2, each holding a non-negative integer. We regard s_0 as a start state and s_L as a stop state. The basic operations of $\mathbf{M}$ are: test whether c_i holds the value 0; and increment/decrement c_i (where attempting to decrement zero yields zero). The program of $\mathbf{M}$ associates with each state s_ℓ other than s_L a basic operation (i.e. an increment, a decrement or a zero-test), together with a specification of the next state of the machine (depending, in the case of of zero-tests, on the outcome). No action is specified for the stop state. A *configuration* for $\mathbf{M}$ is a triple comprising a state together with the values of c_1 and c_2. The *run* of $\mathbf{M}$ is the (finite or infinite) sequence of configurations starting with $\langle s_0, 0, 0\rangle$, where each configuration is obtained from its predecessor as specified by the program of $\mathbf{M}$, in the obvious way. We allow this sequence to stop if a configuration featuring the stop state, s_L, is encountered, in which case we say that the machine $\mathbf{M}$ *terminates*. It is well-known that deterministic Turing machines may be effectively simulated by deterministic 2-counter machines. Hence, the problem of deciding whether a given deterministic 2-counter machine terminates is r.e.-complete.

We proceed to show how runs of deterministic 2-counter machines can be encoded using the logic $\mathcal{C}^2$2E. Recall that, in $\mathcal{C}^2$2E, the distinguished binary predicates E_1 and E_2 must be interpreted as equivalences. Where a structure $\mathfrak{A}$ is clear from context, we refer to the equivalence classes of $E_1^{\mathfrak{A}}$ as E_1*-classes*, and similarly for E_2. Note that the coarsest common refinement $E_1^{\mathfrak{A}} \cap E_2^{\mathfrak{A}}$ of these two equivalences is also an equivalence; to aid intuition, we refer to its equivalence classes as *configurations*. We write $E_{12}(x, y)$ as an abbreviation for the formula $E_1(x, y) \wedge E_2(x, y)$. We employ unary predicates d_1, d_2 to partition the universe, in such a way that, within any E_1- or E_2-class, the elements satisfying them form configurations:

$$\forall x((d_1(x) \vee d_2(x)) \wedge (\neg d_1(x) \vee \neg d_2(x))) \tag{16}$$

$$\bigwedge_{k=1}^{2} \forall x \forall y(E_{12}(x, y) \wedge d_k(x) \rightarrow d_k(y)) \tag{17}$$

$$\bigwedge_{k=1}^{2} \bigwedge_{j=1}^{2} \forall x \forall y(E_k(x, y) \wedge d_j(x) \wedge d_j(y) \rightarrow E_{3-k}(x, y)). \tag{18}$$

We call a configration whose elements satisfy d_k a d_k*-configuration*. It follows that each equivalence class contains at most one d_1-configuration, and at most one d_2-configuration. Where two different configurations, B and B', lie in some E_k-class ($k \in \{1, 2\}$), then we say that B' is the *successor of* B if B is a d_k-configuration and B' a d_{3-k}-configuration. Thus, for B and B' as described, one is the successor of the other. Note also that successors, where they exist, are unique.

We employ unary predicates $s_1, \ldots, s_L$, and refer to them as *states*; we also employ an additional unary predicate s to stand for their disjunction. We require that every configuration contains a unique element satisfying s, which will be in exactly one state:

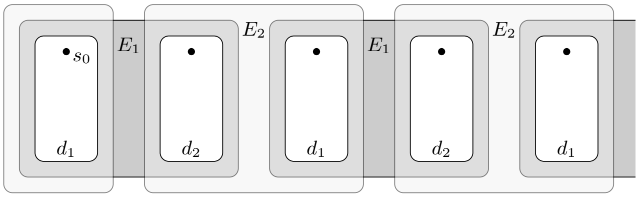

Figure 3: Initial segment of a chain of configurations: each configuration (white region) contains a unique s-element determining its state; the first configuration is in the start state, and forms an E_2-class on its own.

$$\forall x \exists_{[=1]} y(E_{12}(x, y) \wedge s(y)) \tag{19}$$

$$\forall x \left(s(x) \rightarrow \bigvee_{\ell=1}^{L} s_\ell(x) \right) \wedge \bigwedge_{1 \leq \ell < \ell' \leq L} \forall x(s_\ell(x) \rightarrow \neg s_{\ell'}(x)). \tag{20}$$

A configuration whose s-element satisfies s_ℓ will be said to be *in state* s_ℓ. We call s_0 the *start state* and s_L, the *stop state*. We employ a binary predicate t, and we require that $t(x, y)$ holds only between s-elements of configurations one of which is the successor of the other:

$$\begin{aligned} &\forall x \forall y(t(x, y) \rightarrow \\ &\quad (s(x) \wedge s(y) \wedge \\ &\quad\quad \bigvee_{k=1}^{2} (E_k(x, y) \wedge \neg E_{3-k}(x, y) \wedge d_k(x)))). \end{aligned} \tag{21}$$

We require that there exists a d_1-configuration in the start state, that this configuration is the only one in its E_2-class (i.e., is not the successor of any configuration), and that every configuration in a state other than the stop state has a successor:

$$\exists x(d_1(x) \wedge s_0(x) \wedge \forall y(E_2(x, y) \rightarrow E_1(x, y))) \tag{22}$$

$$\bigwedge_{\ell=0}^{L-1} \forall x(s_\ell(x) \rightarrow \exists y.t(x, y)). \tag{23}$$

It follows that, in any model of (16)–(23), there is a chain, $B_0, B_1, \ldots$, (possibly infinite) of distinct configurations, where B_0 is in the start state, and where each B_{i+1} is the successor of B_i. Moreover, if this chain is finite and maximal (i.e. cannot be extended), then its final configuration must be in the stop state. Notice that this condition must obtain if the model is finite. The situation is illustrated in Fig. 3.

Recall that, if B is any configuration, then B contains exactly one element satisfying s. We employ two further unary predicates c_1 and c_2: we refer to the set of elements of B satisfying c_i ($1 \leq i \leq 2$) as the *the c_i-counter in B*, and we refer to the cardinality of this set as the *value of* that counter. It helps to assume that the sets of elements of B satisfying the respective predicates s, c_1 and c_2 partition B; however, this is not formally a requirement.

We now consider any deterministic 2-register machine, $\mathbf{M}$, and proceed to describe the run of $\mathbf{M}$ using $\mathcal{C}^2$2E-formulas. We first define, for $i = 1, 2$, a 1-place formula $c_i^\circ(x)$, which, in effect, states that the c_i-register in the configuration containing x is zero:

$$\neg \exists y(E_{12}(x, y) \wedge c_i(y)).$$

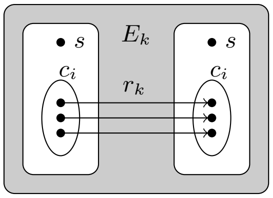

Figure 4: Successive configurations whose elements satisfy the formula $c_i^{=}(x,y)$: the c_i counters are in 1–1 correspondence under r_k, and so are equinumerous.

Using these formulas, we fix these register values for any d_1-configuration that is not a successor to be zero:

$$\forall x(d_1(x) \wedge \forall y(E_2(x,y) \to E_1(x,y)) \to c_1^{\circ}(x) \wedge c_2^{\circ}(x)).$$

We next define a formula $c_i^{=}(x,y)$ with the following property. Suppose b and b' are elements of configurations B and B', respectively, where B' is the successor of B: if the pair $\langle b, b' \rangle$ satisfies $c_i^{=}(x,y)$ then the c_i-counter of B and the c_i-counter of B' contain the same value. To construct $c_i^{=}(x,y)$, we employ a pair of binary predicates r_1, r_2, denoting relations contained within the equivalences E_1, E_2, respectively, but disjoint from the other:

$$\bigwedge_{k=1}^{2} \forall x \forall y(r_k(x,y) \to E_k(x,y) \wedge \neg E_{3-k}(x,y)). \qquad (24)$$

Recall that, under our assumptions concerning b and b', if b satisfies d_k then b and b' lie in a common E_k class. The formula $c_i^{=}(x,y)$ then simply states that, in that case, every element in the c_i-register of B is related by r_k to exactly one element in the c_i-register of B', and that every element in the c_i-register of B' is related by the inverse of r_k to exactly one element in the c_i-register of B.

$$\begin{aligned}&\bigwedge_{k=1}^{2}(d_k(x) \to \\ &\quad \forall y(E_{12}(x,y) \wedge c_i(y) \to \exists_{[=1]}x(r_k(y,x) \wedge c_i(x)))\wedge \\ &\qquad \forall x(E_{12}(y,x) \wedge c_i(x) \to \exists_{[=1]}y(r_k(y,x)) \wedge c_i(y))).\end{aligned}$$

Note how the variables x and y are 're-used' by quantifiers. This formula relies on the sentence (24) to have its advertised effect: the relation r_k holds only between elements in the same E_k-class but different E_{3-k}-classes. The situation is illustrated in Fig. 4.

Similarly, we can define a formula $c_i^{+}(x,y)$ entailing that, if the configuration B' containing y is the successor of the configuration B containing x, then the c_i-register of B' is one greater than that of B, and a formula $c_i^{-}(x,y)$ entailing that the c_i-register of B' is one less than that of B (or that both are zero).

Using the formulas $c_i^{\circ}(x)$, $c_i^{+}(x,y)$ and $c_i^{-}(x,y)$, we may then encode the program of $\mathbf{M}$ in the expected way. For example, if the basic operation of $\mathbf{M}$ associated with state s_i is to increment counter c_1 and move to state s_j, then we require:

$$\forall x \forall y(s_i(x) \wedge t(x,y) \to (c_1^{+}(x,y) \wedge c_2^{=}(x,y) \wedge s_j(y))).$$

Writing such formulas for all states s_i ($0 \leq i < L$), we can effectively construct a $\mathcal{C}^2$2E-formula $\varphi_{\mathbf{M}}$ any model of which contains a sequence of configurations $B_0, B_1, \ldots$, encoding the run of $\mathbf{M}$. Indeed, $\varphi_{\mathbf{M}}$ has a finite model if and only if $\mathbf{M}$ has a terminating run. Hence:

Theorem 14. *The finite satisfiability problem for $\mathcal{C}^2 2E$ is r.e.-complete.*

Bearing in mind that $\mathbf{M}$ terminates just in case its run encounters the stop state, we see that $\varphi_{\mathbf{M}} \wedge \forall x \neg s_L(x)$ has an (infinite) model if and only if $\mathbf{M}$ is non-terminating. Hence:

Theorem 15. *The satisfiability problem for $\mathcal{C}^2 2E$ is co-r.e.-complete.*

Acknowledgments

This work is supported by the EPSRC, grant number EP/K017438/1. The author would like to thank Mr. Y. Guskov, Dr. E. Kieroński, Prof. L. Tendera and Prof. Tinko Tinchev for valuable comments and suggestions.

References

[1] C. Chang and H. Keisler. *Model Theory*. North Holland, Amsterdam, 1973.

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

[3] 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.

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

[5] A. Janiczak. Undecidability of some simple formalized theories. *Fundamenta Mathematicae*, 40(1):131–139, 1953.

[6] Y. Kazakov. *Saturation-based decision procedures for extensions of the guarded fragment*. PhD thesis, Universität des Saarlandes, Saarbrücken, Germany, 2006.

[7] Y. Kazakov, U. Sattler, and E. Zolin. How many legs do I have? Non-simple roles in number restrictions revisited. In *Proc. of LPAR*, volume 4790 of *LNCS*, Berlin, 2007. Springer.

[8] E. Kieroński. Results on the guarded fragment with equivalence or transitive relations. In *Computer Science Logic*, volume 3634 of *LNCS*, pages 309–324. Springer, 2005.

[9] E. Kieroński and M. Otto. Small substructures and decidability issues for first-order logic with two variables. In *Logic in Computer Science*, pages 448–457. IEEE, 2005.

[10] E. Kieroński and L. Tendera. On finite satisfiability of two-variable first-order logic with equivalence relations. In *Logic in Computer Science*. IEEE, 2009.

[11] E. Kieroński, J. Michalyszyn, I. Pratt-Hartmann, and L. Tendera. Two-variable first-order logic with equivalence closure. *SIAM Journal on Computing*, 43(3):1012–1063, 2014.

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

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

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

[15] A. Schrijver. *Theory of Linear and Integer Programming*. John Wiley & sons, New York, 1998.

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

[17] W. Szwast and L. Tendera. FO^2 with one transitive relation is decidable. In *STACS*, volume 20 of *LIPIcs*, pages 317–328. Schloß Dagstuhl - Leibniz-Zentrum für Informatik, 2013.

[18] L. Tendera. Counting in the two-variable guarded fragment with transitivity. In *Proc. of STACS*, volume 3404 of *Lecture Notes in Computer Science*, pages 83–96, Berlin, 2005. Springer.