Fluted Logic with Counting

Ian Pratt-Hartmann
Department of Computer Science, University of Manchester, UK
Institute of Computer Science, University of Opole, Poland

Abstract

The fluted fragment is a fragment of first-order logic in which the order of quantification of variables coincides with the order in which those variables appear as arguments of predicates. It is known that the fluted fragment possesses the finite model property. In this paper, we extend the fluted fragment by the addition of counting quantifiers. We show that the resulting logic retains the finite model property, and that the satisfiability problem for its $(m+1)$-variable sub-fragment is in m-NExpTime for all positive m. We also consider the satisfiability and finite satisfiability problems for the extension of any of these fragments in which the fluting requirement applies only to sub-formulas having at least three free variables.

2012 ACM Subject Classification Theory of computation → Complexity theory and logic

Keywords and phrases Fluted fragment, counting quantifiers, satisfiability, complexity

Digital Object Identifier 10.4230/LIPIcs.ICALP.2021.141

Category Track B: Automata, Logic, Semantics, and Theory of Programming

Funding This work was supported by the Polish NCN, grant number 2018/31/B/ST6/03662.

Acknowledgements The author wishes to thank Prof. L. Tendera for valuable discussions.

1 Introduction

The *fluted fragment*, $\mathcal{FL}$, is a fragment of first-order logic in which, *very* roughly, the order of quantification of variables coincides with the order in which those variables appear as arguments of predicates, for example:

$$\begin{array}{l}\text{No student admires every professor}\\ \forall x_1(\text{std}(x_1) \rightarrow \neg\forall x_2(\text{prof}(x_2) \rightarrow \text{admr}(x_1, x_2)))\end{array} \tag{1}$$

$$\begin{array}{l}\text{No lecturer introduces any professor to every student}\\ \forall x_1(\text{lectr}(x_1) \rightarrow \neg\exists x_2(\text{prof}(x_2) \wedge \forall x_3(\text{std}(x_3) \rightarrow \text{intro}(x_1, x_2, x_3)))).\end{array} \tag{2}$$

More precisely, in fluted formulas, all atoms are of the form $p(x_\ell, \ldots, x_m)$, with a contiguous sequence of variables as their arguments, Boolean combinations can only be formed from formulas whose last free variable is the same, and only the last free variable in a formula may be quantified. Equality is not present. It is known that $\mathcal{FL}$ has the finite model property, and that its m-variable sub-fragment, $\mathcal{FL}^m$, is $\lfloor m/2 \rfloor$-NExpTime-hard for all $m \geq 2$ and in $(m-2)$-NExpTime for all $m \geq 3$ [17]. Hence, the satisfiability problem for $\mathcal{FL}$ is Tower-complete in the system of trans-elementary complexity classes of [24]. (It was incorrectly claimed in [20] that this problem is in NExpTime.)

Counting quantifiers are expressions of the form $\exists_{[\leq M]}$, $\exists_{[\geq M]}$ and $\exists_{[=M]}$, where M is a positive integer, with the interpretations "there exist at most/at least/exactly M ...". We investigate the addition of counting quantifiers and equality to the fluted fragment:

$$\begin{array}{l}\text{At most three lecturers introduce a professor to at least five students}\\ \exists_{[\leq 3]}x_1(\text{lectr}(x_1) \wedge \exists x_2(\text{prof}(x_2) \wedge \exists_{[\geq 5]}x_3(\text{std}(x_3) \wedge \text{intro}(x_1, x_2, x_3))))\end{array} \tag{3}$$

$$\begin{array}{l}\text{Every absent-minded professor introduces some student to himself}\\ \forall x_1(\text{abs-mnd}(x_1) \wedge \text{prof}(x_1) \rightarrow \exists x_2 \exists x_3(\text{std}(x_3) \wedge x_2 = x_3 \wedge \text{intro}(x_1, x_2, x_3))).\end{array} \tag{4}$$

48th International Colloquium on Automata, Languages, and Programming (ICALP 2021).
Editors: Nikhil Bansal, Emanuela Merelli, and James Worrell; Article No. 141; pp. 141:1–141:17
Leibniz International Proceedings in Informatics
LIPICS Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

We denote this extension of $\mathcal{FL}$ by $\mathcal{FLC}$, and its m-variable sub-fragment by $\mathcal{FLC}^m$. We also consider the corresponding fragments $\mathcal{SFC}$ and $\mathcal{SFC}^m$, in which the fluting restriction is waived for sub-formulas with at most two variables. (Formal definitions are given in Sec. 2.)

The definition of $\mathcal{FL}$ employed here is that given by Purdy [19], who traces its origins to Quine [22]. (The term "fluted" is actually Quine's.) While it is unclear whether the quantification patterns specified by Purdy are really those that Quine had in mind, it is Purdy's definition which has established itself, and indeed which is – from the point of view of recent work in computational logic – of greater interest. In particular, two-variable fluted logic with counting, $\mathcal{FLC}^2$ includes (under a simple translation) the description logic $\mathcal{ALCHOQ}$, whose sub-fragments have been the focus of intensive investigation over recent decades; and its semi-fluted extension, $\mathcal{SFC}^2$ coincides with $\mathcal{C}^2$, the two-variable fragment of first-order logic with counting quantifiers, whose satisfiability and finite satisfiability problems are known to be NExpTime-complete [6, 12, 15]. Of course, $\mathcal{FL}$ is not limited in the number of variables formulas can contain, a property it shares with the *guarded fragment* [1], which also has the finite model property, and whose satisfiability problem is in 2-ExpTime [5]. In fact, our logic $\mathcal{SFC}$ extends $\mathcal{C}^2$ with fluted formulas in much the same way as the so-called *triguarded fragment* [23] extends the two-variable fragment of first-order logic (without equality) with guarded formulas. The triguarded fragment has recently been shown to have the finite model property [10]; its satisfiability problem is 2-NExpTime-complete, but becomes undecidable in the presence of equality. We note that negation can be applied freely in $\mathcal{FLC}$ and $\mathcal{SFC}$. Thus, these fragments are not subject to any type of guardedness restrictions: for example, (2) is not guarded or even negation-guarded [2].

In this extended abstract, we show that $\mathcal{FLC}$ has the finite model property, and that the satisfiability problem for $\mathcal{FLC}^{m+1}$ is in m-NExpTime for all $m \geq 1$. We also show that the satisfiability and finite satisfiability problems for $\mathcal{SFC}$ remain decidable.

2 Preliminaries

In the context of fluted formulas, logical variables are taken from the sequence $\bar{x}_\omega = x_1, x_2, \ldots$, and all signatures are purely relational, i.e., there are no individual constants or function symbols; however, we allow 0-ary relations (proposition letters). We employ the syntax of counting quantifiers $\exists_{[\leq M]}$, $\exists_{[\geq M]}$ and $\exists_{[=M]}$, where M is a (numeral denoting a) positive integer, under the expected semantics. A *multiset* over some carrier set X is a function f from X to cardinal numbers, where, for each $x \in X$, $f(x)$ is the *multiplicity* with which x occurs in f. Informally, we identify multisets differing only by elements of multiplicity 0. Almost all mulitplicities we encounter will be finite.

We begin with the syntax of the logics considered here. Define the sets of formulas $\mathcal{FLC}^{[m]}$, for all $m \geq 0$, by simultaneous structural recursion as follows:

- **(i)** any atom $p(x_\ell, \ldots, x_m)$, where $x_\ell, \ldots, x_m$ is a contiguous subsequence of $\bar{x}_\omega$ and p a predicate of arity $m - \ell + 1$, is in $\mathcal{FLC}^{[m]}$;
- **(ii)** $\mathcal{FLC}^{[m]}$ is closed under boolean combinations;
- **(iii)** if φ is in $\mathcal{FLC}^{[m+1]}$, then $\exists x_{m+1}.\varphi$ and $\forall x_{m+1}.\varphi$ are in $\mathcal{FLC}^{[m]}$,
- **(iv)** if φ is in $\mathcal{FLC}^{[m+1]}$ and M a non-negative integer, then $\exists_{[\leq M]} x_{m+1}.\varphi$, $\exists_{[\geq M]} x_{m+1}.\varphi$ and $\exists_{[=M]} x_{m+1}.\varphi$ are in $\mathcal{FLC}^{[m]}$.

It is intended that Clause (i) allows the case $\ell = m + 1$ (empty sequence of arguments), so that the atoms in question are proposition letters; and when $m = \ell + 1$ (exactly two arguments), we allow p to be the equality predicate. We define the sets of formulas $\mathcal{FL}^{[m]}$ similarly, except that we do not allow the equality predicate in Clause (i), and Clause (iv) is dropped altogether. The *fluted fragment* is the set of formulas $\mathcal{FL} = \bigcup_{m \geq 0} \mathcal{FL}^{[m]}$, and the *fluted fragment with counting and equality*, the set of formulas $\mathcal{FLC} = \bigcup_{m \geq 0} \mathcal{FLC}^{[m]}$. Finally, we define $\mathcal{FL}^m$ to be the fragment of $\mathcal{FL}$ in which at most the variables $x_1, \ldots, x_m$ appear (free or bound); and similarly for $\mathcal{FLC}^m$. Thus, (1) is in $\mathcal{FL}^2$, and (2) in $\mathcal{FL}^3$, while (3) and (4) are in $\mathcal{FLC}^3$. Do not confuse $\mathcal{FLC}^m$ with $\mathcal{FLC}^{[m]}$: all of (1)–(4) are in $\mathcal{FLC}^{[0]}$. By *sentence*, we mean a formula with no free variables.

Denote by $\mathcal{C}^2$ the *two-variable fragment of first-order logic with counting*, i.e. the set of first-order formulas with (equality and) counting quantifiers over a purely relational signature, and featuring only two logical variables. (We may without loss of useful expressive power assume that all predicates have arity at most 2.) Formulas of $\mathcal{C}^2$ are not required to be fluted. For example, $\forall x_1 \exists x_2 . r(x_1, x_2) \wedge \forall x_2 \exists_{[\leq 1]} x_1 . r(x_1, x_2) \wedge \exists x_2 \forall x_1 \neg r(x_1, x_2)$ is a formula of $\mathcal{C}^2$, but not of $\mathcal{FLC}^2$. It is straightforward to see that this formula is satisfiable, but only in infinite models. Thus, $\mathcal{C}^2$ lacks the finite model property. It is well-known that the satisfiability and finite satisfiability problems for the three-variable fragment of first-order logic (even without with counting) are undecidable.

In the context of $\mathcal{FLC}$, the possibility arises of waiving the fluting restrictions on sub-formulas featuring at most two free variables. Define the sets of formulas $\mathcal{SFC}^{[m]}$ in the same way as $\mathcal{FLC}^{[m]}$, but with the additional clauses (for $m \geq 2$):

(v) Any $\mathcal{C}^2$-formula ψ, whose set of free variables is equal to one of $\{x_{m-1}, x_m\}$, $\{x_m\}$ or $\emptyset$, is in $\mathcal{SFC}^{[m]}$.

We then take the *semi-fluted fragment with counting* to be the set of formulas $\mathcal{SFC} = \bigcup_{m \geq 0} \mathcal{SFC}^{[m]}$, denoting its m-variable sub-fragment by $\mathcal{SFC}^m$. If φ is a formula of any of the above fragments, we take its *size*, $\|\varphi\|$, to be the number of bits required to write it, on the understanding that numerical subscripts are encoded as binary strings. Since $\mathcal{SFC}$ contains $\mathcal{C}^2$, it lacks the finite model property. We can now state our main results.

▶ **Theorem 1.** *The logic $\mathcal{FLC}$ has the finite model property. The satisfiability problem for $\mathcal{FLC}^{m+1}$ is in m-NExpTime for all $m \geq 1$.*

▶ **Theorem 2.** *The satisfiability and finite satisfiability problems for $\mathcal{SFC}$ are decidable.*

Assuming, as we shall, that the arity of any predicate is fixed in advance, variables in fluted logic convey no information at all, and therefore can be omitted. (This may have been part of the motivation for Quine [21].) The same applies to fluted formulas with counting quantifiers. For example, (3) and (4) can be written, respectively, as:

$$\exists_{[\leq 3]}(\text{lectr} \wedge \exists(\text{prof} \wedge \exists_{[\geq 5]}(\text{std} \wedge \text{intro}))) \tag{5}$$

$$\forall(\text{abs-mnd} \wedge \text{prof} \rightarrow \exists\exists(\text{std} \wedge = \wedge \text{intro})). \tag{6}$$

It is straightforward to reconstruct (3) and (4) (up to a shift of variable indices) from (5) and (6). Consequently, we employ variable-free notation for $\mathcal{FLC}$ in the sequel, as it is more compact, though formulas such as "std $\wedge = \wedge$ intro" admittedly take some getting used to. It is important to realize that, with variable-free notation, any formula of $\mathcal{FLC}^{[m]}$ is, *without lexical change*, also a formula of $\mathcal{FLC}^{[m+1]}$. For example, the sub-formula $\exists(\text{prof} \wedge \exists_{[\geq 5]}(\text{std} \wedge \text{intro}))$ of (5) may be reconstructed as the $\mathcal{FLC}^{[1]}$-formula $\varphi(x_1) := \exists x_2(\text{prof}(x_2) \wedge \exists_{[\geq 5]} x_3(\text{std}(x_3) \wedge \text{intro}(x_1, x_2, x_3)))$, or alternatively as the $\mathcal{FLC}^{[2]}$-formula $\varphi'(x_1, x_2) := \exists x_3(\text{prof}(x_3) \wedge \exists_{[\geq 5]} x_4(\text{std}(x_4) \wedge \text{intro}(x_2, x_3, x_4)))$, and so on.

Thus, using variable-free notation, the sets $\mathcal{FLC}^{[m]}$ are the minimal family of sets of formulas satisfying:

(i) any predicate p of arity less than or equal to m is in $\mathcal{FLC}^{[m]}$;
(ii) $\mathcal{FLC}^{[m]}$ is closed under boolean combinations;
(iii) if φ is in $\mathcal{FLC}^{[m+1]}$, then $\exists\varphi$ and $\forall\varphi$ are in $\mathcal{FLC}^{[m]}$,
(iv) if φ is in $\mathcal{FLC}^{[m+1]}$ and M a non-negative integer, then $\exists_{[\leq M]}\varphi$, $\exists_{[\geq M]}\varphi$ and $\exists_{[=M]}\varphi$ are in $\mathcal{FLC}^{[m]}$.

We also use the term *sentence* in this context to mean a formula of $\mathcal{FLC}^{[0]}$, and we continue to use the notation $\mathcal{FLC}^m$ to denote the same set of formulas as defined above, but written without variables. We remark that there is no difference between a *quantifier-free* formula of $\mathcal{FLC}^m$ and a *quantifier-free* formula of $\mathcal{FLC}^{[m]}$; it is just a Boolean combination of predicates (including equality) of arity at most m.

If $\varphi \in \mathcal{FLC}^{[m]}$, and $\bar{a} = a_1, \ldots, a_m$ is an m-tuple of elements from some structure $\mathfrak{A}$ interpreting the signature of φ, we write $\mathfrak{A} \models \varphi[\bar{a}]$ as usual to indicate that $\bar{a}$ satisfies φ in $\mathfrak{A}$, under the assignments $x_1 \mapsto a_1, \ldots, x_m \mapsto a_m$. Observe that this notation remains meaningful even with variable-free notation. Indeed, if $\bar{a} \in A^m$ and $a \in A$, then $\mathfrak{A} \models \varphi[\bar{a}]$ if and only if $\mathfrak{A} \models \varphi[a\bar{a}]$. (This trick is important, because it will be used at various points in Sec. 4.) The notation $\forall^m$ stands for a block $\forall \cdots \forall$ of m universal quantifiers.

Of course, we cannot – without further ado – use variable-free notation for $\mathcal{SFC}$.

The analysis of decidable fragments is often simplified by the use of normal forms in the style of [25]. Here, we adapt the normal forms for $\mathcal{FL}$ from [17, Lemma 4.1].

▶ **Lemma 3.** *Let φ be a formula of $\mathcal{FLC}^{m+1}$ ($m \geq 1$). Then we may compute, in time bounded by a polynomial function of $\|\varphi\|$, an $\mathcal{FLC}^{m+1}$-sentence satisfiable over the same domains as φ, and having the form*

$$\bigwedge_{s \in S} \forall^m(\mu_s \rightarrow \exists_{[=M_s]}\zeta_s) \wedge \bigwedge_{t \in T} \forall^m(\nu_t \rightarrow \forall\eta_t) \wedge \forall^{m+1}\theta, \tag{7}$$

where S and T are index sets, the μ_s and ν_t are quantifier-free $\mathcal{FLC}^m$-formulas, the ζ_s, η_t and θ are quantifier-free $\mathcal{FLC}^{m+1}$-formulas, and the M_s are positive integers.

Proof. By prepending existential quantifiers if necessary, we may assume that φ is a sentence. Call a quantifier of the form $\exists_{[=M]}$ an *equality quantifier*. Somewhat counter-intuitively, we first remove all equality quantifiers from φ. Let $\varphi_0 := \varphi$, and suppose φ_0 possesses a sub-formula $\theta = \exists_{[=M]}\chi$, such that χ contains no equality quantifiers. Let ℓ be the smallest number such that $\theta \in \mathcal{FLC}^{[\ell]}$. Let p, q be fresh predicates or arity ℓ, and define $\varphi_1 := \varphi_0[\theta/(p \wedge q)]$ and

$$\psi_1 := \forall^\ell(p \leftrightarrow \exists_{[\geq M]}\chi) \wedge \forall^\ell(q \leftrightarrow \exists_{[\leq M]}\chi).$$

It is obvious that φ_0 and $\varphi_1 \wedge \psi_1$ are satisfiable over the same domains Now process φ_1 in the same way to obtain φ_2, until we eventually obtain a sentence φ_n containing no equality quantifiers. This process evidently terminates in polynomial time. Since φ and $\varphi_n \wedge \psi_n \wedge \cdots \wedge \psi_1$ are satisfiable over the same domains, we may simply assume, henceforth, that φ contains no equality quantifiers.

Since there are no equality quantifiers in φ we may move negations inward in the usual way, so that they apply to atomic formulas. At this point, we may follow the proof of the analogous theorem for $\mathcal{FL}$ presented in [17, pp. 174], obtaining, in polynomial time, a sentence satisfiable over the same domains as φ, and having the form

$$\bigwedge_{s \in S} \forall^m(\mu_s \rightarrow \exists_{[\bowtie_s M_s]}\zeta_s) \wedge \bigwedge_{t \in T} \forall^m(\nu_t \rightarrow \forall\eta_t) \wedge \forall^{m+1}\theta,$$

where each $\bowtie_s$ is one of the symbols $\leq$ or $\geq$. We now eliminate all occurrences of $\leq$ and $\geq$ with $=$. For any conjunct $\forall^m(\mu_s \rightarrow \exists_{[\leq M_s]}\zeta_s)$, we let q_s be a new predicate of arity $m+1$ and replace this conjunct with

$$\forall^m(\mu_s \rightarrow \exists_{[=M_s]}q_s) \wedge \forall^{m+1}(\zeta_s \rightarrow q_s).$$

The case $\geq$ is treated similarly. ◀

We refer to formulas of the forms (7) as *normal-form* formulas of $\mathcal{FLC}^{m+1}$.

The following notions are useful for analysing $\mathcal{FLC}$-formulas. Let Σ be a finite relational signature. A *fluted literal over* Σ is an expression p or $\neg p$, where $p \in \Sigma \cup \{=\}$. (Remember that, under variable-free notation, we think of p as an atom $p(x_\ell, \ldots, x_m)$ of $\mathcal{FLC}^{[m]}$.) The *arity* of the literal is the arity of the underlying predicate. A *fluted m-type* τ (*over* Σ) is a maximal consistent set of fluted literals over Σ having arity at most m. We call τ *reflexive* if it contains the literal $=$. We silently identify fluted m-types with their conjunctions, thus regarding them as quantifier-free $\mathcal{FL}^m$-formulas. Obviously, if τ is a fluted m-type and ψ a quantifier-free $\mathcal{FLC}^m$-formula over the same signature, then either $\models \tau \rightarrow \psi$ or $\models \tau \rightarrow \neg\psi$. Suppose $\mathfrak{A}$ is a structure over domain A interpreting Σ, and $\bar{a}$ an m-tuple of elements from A ($m \geq 1$). Then there is a unique fluted m-type τ such that $\mathfrak{A} \models \tau[\bar{a}]$. We denote τ by $\text{ftp}^{\mathfrak{A}}[\bar{a}]$, and call it the fluted m-type of $\bar{a}$.

A *fluted star-type* σ of *dimension* m *over* Σ is a multiset of fluted $(m+1)$-types over Σ at most one of which (counting multiplicities) is reflexive. The term "star-type" comes from [15]; in the present paper,we concern ourselves either with *fluted* star-types or *semi-fluted* star-types (defined below). Since we may list the fluted $(m+1)$-types over Σ as $\tau_1, \ldots, \tau_J$ in some fixed order ($J \leq 2^{|\Sigma|+1}$), we can regard any fluted star-type σ over Σ as a vector $(\sigma(\tau_1), \ldots, \sigma(\tau_J))$ of cardinal numbers. We say that σ is *M-bounded* if $|\sigma| = \sigma(\tau_1) + \cdots + \sigma(\tau_J) \leq M$. If ζ is a quantifier-free $\mathcal{FLC}^{m+1}$-formula over Σ, the *retract* of σ to ζ is the fluted star-type $\sigma \restriction \zeta$ given by:

$$(\sigma \restriction \zeta)(\tau) = \begin{cases} \sigma(\tau) & \text{if} \models \tau \rightarrow \zeta \\ 0 & \text{otherwise.} \end{cases}$$

Thus, when performing a retract to ζ, we remove from σ all occurrences of those fluted $(m+1)$-types inconsistent with ζ (i.e. set their multiplicities to 0). We say that σ is a *fluted* ζ*-star-type* if $\sigma = \sigma \restriction \zeta$. Suppose $\mathfrak{A}$ is a structure over domain A interpreting Σ, and $\bar{a}$ an m-tuple of elements from A ($m \geq 1$). If ζ is any quantifier-free formula of $\mathcal{FLC}^{m+1}$ over Σ, we may define a fluted ζ-star-type σ of dimension m by setting, for each fluted $(m+1)$-type τ over Σ, $\sigma(\tau) = |\{b \in A : \mathfrak{A} \models \tau[\bar{a}b] \text{ and } \mathfrak{A} \models \zeta[\bar{a}b]\}|$. We denote σ by $\text{fst}^{\mathfrak{A}}_{\zeta}[\bar{a}]$, and call it the *fluted* ζ*-star-type of* $\bar{a}$ *in* $\mathfrak{A}$. As an aide to intuition, think of $\bar{a}$ as emitting various "ζ-rays", each absorbed by some element $b \in A$ such that $\mathfrak{A} \models \zeta[\bar{a}b]$. Every ζ-ray has a "colour" specified by some fluted $(m+1)$-type τ such that $\models \tau \rightarrow \zeta$. The fluted star-type $\text{fst}^{\mathfrak{A}}_{\zeta}[\bar{a}]$ simply counts how many rays of each colour arise in this way. To grasp the utility of these notions, let φ be a normal-form formula (7), and suppose $\mathfrak{A} \models \varphi$. Now let $\mathfrak{B}$ be a structure such that, for every m-tuple $\bar{b}$ from B, there exists an m-tuple $\bar{a}$ from A satisfying: (i) $\text{ftp}^{\mathfrak{B}}[\bar{b}] = \text{ftp}^{\mathfrak{A}}[\bar{a}]$; (ii) $\text{fst}^{\mathfrak{B}}_{\zeta_s}[\bar{b}] = \text{fst}^{\mathfrak{A}}_{\zeta_s}[\bar{a}]$ for every s such that $\mathfrak{B} \models \mu_s[\bar{b}]$, and (iii) for every $b' \in B$, there exists $a' \in A$ such that $\text{ftp}^{\mathfrak{B}}[\bar{b}b'] = \text{ftp}^{\mathfrak{A}}[\bar{a}a']$. Then $\mathfrak{B} \models \varphi$.

We mentioned earlier that the logic $\mathcal{C}^2$ lacks the finite model property, but that its satisfiability and finite satisfiability problems are in NExpTime – these being relatively non-trivial results. The following lemma illustrates the comparative weakness of its fluted sub-fragment, $\mathcal{FLC}^2$ by establishing that this has the finite model property. In fact, we adapt a well-known proof (see [8, pp. 77 ff.], [9]) of the corresponding statement for the monadic fragment of first-order logic [11].

▶ **Lemma 4.** *If φ is a satisfiable formula of $\mathcal{FLC}^2$, then φ has a model of size bounded by an exponential function of $\|\varphi\|$. The satisfiablity problem for $\mathcal{FLC}^2$ is in NExpTime.*

Proof. By Lemma 3, we may assume that φ is of the form (7), over a signature Σ, where $m = 1$. We may further assume that Σ features no predicates of arity 0, since their truth-values can be guessed. Let $M = \sum_{s \in S} M_s$, and suppose $\mathfrak{A} \models \varphi$. For each fluted 1-type π (over Σ) realized in $\mathfrak{A}$, let B_π be a set of cardinality $\min(M+1, |\{a \in A \mid \text{ftp}^{\mathfrak{A}}[a] = \pi\}|)$ and let $B = \bigcup_\pi B_\pi$. Thus, $|B| \leq (M+1)2^{|\Sigma|+1}$. We define a structure $\mathfrak{B}$ with domain B and show that $\mathfrak{B} \models \varphi$. For each $b \in B_\pi$, set $\text{ftp}^{\mathfrak{B}}[b] = \pi$. These fluted 1-types involve only unary predicates, and so may be assigned independently of each other. To complete the definition of $\mathfrak{B}$ we fix the extensions of binary predicates so as to determine $\text{ftp}^{\mathfrak{B}}[bb']$ for any ordered pair of elements $\langle b, b' \rangle \in B^2$.

Pick any $b \in B$, and let $a \in A$ be such that $\text{ftp}^{\mathfrak{A}}[a] = \text{ftp}^{\mathfrak{B}}[b] = \pi$, say. Write $S' = \{s \in S : \mathfrak{A} \models \mu_s[a]\}$ and $\zeta' = \bigvee_{s \in S'} \zeta_s$, and let A' be the set of elements $a' \in A \setminus \{a\}$ such that $\mathfrak{A} \models \zeta'[aa']$, say $A' = \{a_1, \ldots, a_k\}$, with $k \leq M$. Thus we may choose a subset $B' = \{b_1, \ldots, b_k\} \subseteq B \setminus \{b\}$ such that, for all i $(1 \leq i \leq k)$, $\text{ftp}^{\mathfrak{B}}[b_i] = \text{ftp}^{\mathfrak{A}}[a_i]$. However this is done, we are guaranteed that, for every $b' \in B \setminus (B' \cup \{b\})$, we can find some $a' \in A \setminus (A' \cup \{a\})$ such that $\text{ftp}^{\mathfrak{A}}[a'] = \text{ftp}^{\mathfrak{B}}[b']$. Now set $\text{ftp}^{\mathfrak{B}}[bb] = \text{ftp}^{\mathfrak{A}}[aa]$ and $\text{ftp}^{\mathfrak{B}}[bb_i] = \text{ftp}^{\mathfrak{A}}[aa_i]$ for all i $(1 \leq i \leq k)$. Further, for all $b' \in B \setminus (B' \cup \{b\})$, pick some $a' \in A \setminus (A' \cup \{a\})$ such that $\text{ftp}^{\mathfrak{A}}[a'] = \text{ftp}^{\mathfrak{B}}[b']$, and set $\text{ftp}^{\mathfrak{B}}[bb'] = \text{ftp}^{\mathfrak{A}}[aa']$. Observe that, in the latter case, $\mathfrak{A} \not\models \zeta'[aa']$, and therefore $\mathfrak{B} \not\models \zeta'[bb']$. Hence, $\text{fst}^{\mathfrak{B}}_{\zeta'}[b] = \text{fst}^{\mathfrak{A}}_{\zeta'}[a]$, so that $\text{fst}^{\mathfrak{B}}_{\zeta_s}[b] = \big(\text{fst}^{\mathfrak{B}}_{\zeta'}[b]\big)\restriction\zeta_s = \big(\text{fst}^{\mathfrak{A}}_{\zeta'}[a]\big)\restriction\zeta_s = \text{fst}^{\mathfrak{A}}_{\zeta_s}[a]$ for every $s \in S'$. By carrying out this construction for every $b \in B$, we fully define $\mathfrak{B}$. Note that the *fluted* 2-types assigned in this process never clash with the fluted 1-types already assigned, and never clash with each other. Thus, for every element $b \in B$ there exists $a \in A$ such that: (i) $\text{ftp}^{\mathfrak{B}}[b] = \text{ftp}^{\mathfrak{A}}[a]$; (ii) $\text{fst}^{\mathfrak{B}}_{\zeta_s}[b] = \text{fst}^{\mathfrak{A}}_{\zeta_s}[a]$ for every s such that $\mathfrak{B} \models \mu_s[b]$, and (iii) for every element $b' \in B$, there exists $a' \in A$ such that $\text{ftp}^{\mathfrak{B}}[bb'] = \text{ftp}^{\mathfrak{A}}[aa']$. Hence $\mathfrak{B} \models \varphi$. ◀

At various points in the ensuing argument, we need to vary the signatures interpreted by structures. The following notation and terminology is standard. If $\mathfrak{A}^+$ is any structure interpreting a signature Σ^+, and $\Sigma \subseteq \Sigma^+$, we denote by $\mathfrak{A}^+\restriction\Sigma$ the structure obtained by forgetting the predicates in $\Sigma^+ \setminus \Sigma$. We call $\mathfrak{A} = \mathfrak{A}^+\restriction\Sigma$ the *reduct* of $\mathfrak{A}^+$ to Σ, and say that $\mathfrak{A}^+$ is an *expansion* of $\mathfrak{A}$.

3 Existential Presburger quantifiers

In view of Lemma 4, a natural strategy for proving Theorem 1 suggests itself: reduce the satisfiability problem for $\mathcal{FLC}^{m+1}$ to that for $\mathcal{FLC}^m$. This is *nearly* the strategy we follow. To make it work, however, we must generalize the notion of counting quantifiers. Denote the natural numbers $\{0, 1, 2, \ldots\}$ by $\mathbb{N}$. A *linear Diophantine inequality* is an expression $a_1 v_1 + \cdots + a_n v_n + b \leq c_1 v_1 + \cdots + c_n v_n + d$, with coeffecients in $\mathbb{N}$. If $\mathcal{E}(\bar{v})$ is a system of linear Diophantine inequalities in variables $\bar{v}$, a *solution* of $\mathcal{E}$ is an assignment of natural numbers $\bar{a}$ to the variables $\bar{v}$ which make all inequalities of $\mathcal{E}$ true. It was shown in [4] that one can bound the values occurring in the solutions of such systems. (Various such bounds are available: see, e.g. [13].) The following is adequate for our purposes:

▶ **Theorem 5** (from [14], Corollary 1)**.** *Let $\mathcal{E}$ be a system of m linear Diophantine inequalities in n variables, with maximum coefficient M. If $\mathcal{E}$ has a solution, then it has one in which all values are bounded by $(2 + (n+1)M)^{n+m}$.*

By *Presburger arithmetic*, we understand the set of first-order formulas (with equality) over the signature $\{\mathbb{N}, +, \leq, \cdot\}$ whose atomic sub-formulas are linear Diophantine inequalities. We interpret these symbols over the domain $\mathbb{N}$ in the standard way (with the constants $\mathbb{N}$ interpreted as themselves), and say that a tuple of natural numbers $\bar{a}$ *satisfies* a formula of Presburger arithmetic $\Theta(\bar{v})$ if $\mathbb{N} \models \Theta[\bar{a}]$. The *size* of Θ, denoted $\|\Theta\|$, is the number of bits required to write it in the usual way, under the assumption that the individual constants (i.e. coefficients of the linear Diophantine inequalities) are encoded as bit-strings. By *existential Presburger arithmetic*, we mean the set of formulas of Presburger arithmetic of the form $\Theta(\bar{v}) = \exists \bar{u}.\Xi(\bar{v}, \bar{u})$, where Ξ is quantifier-free. Theorem 5 immediately yields (see also, e.g. [7, Table 1]):

▶ **Corollary 6.** *There is a non-deterministic procedure which, given a formula $\Theta(\bar{v})$ of existential Presburger arithmetic and tuple $\bar{a}$ of natural numbers bounded by M with the same arity as $\bar{v}$, has a successful run if and only if $\bar{a}$ satisfies $\Theta(\bar{v})$. This procedure runs in time bounded by a polynomial function of $\|\Theta\| + \log M$.*

Now for our generalization of counting quantifiers.

Fix some $m \geq 1$, and let Σ be a purely relational signature, M a positive integer and Θ a formula of existential Presburger arithmetic in variables $\bar{v} = v_1, \ldots, v_J$ corresponding to the fluted $(m+1)$-types $\tau_1, \ldots, \tau_J$ over Σ. We call an expression $\mathsf{Q}\langle \Sigma, M, \Theta \rangle$ a *fluted existential Presburger quantifier* (or: *fluted ep-quantifier*). If ζ is a quantifier-free formula of $\mathcal{FL}^{m+1}$, we allow formulas φ of the form $\mathsf{Q}\langle \Sigma, M, \Theta \rangle \zeta$, with semantics given by declaring, for any structure $\mathfrak{A}$ interpreting a signature $\Sigma' \supseteq \Sigma$ and any m-tuple $\bar{a}$ of elements from A:

$$\mathfrak{A} \models \varphi[\bar{a}] \text{ if and only if } \mathrm{fst}_\zeta^{(\mathfrak{A} \restriction \Sigma)}[\bar{a}] \text{ is } M\text{-bounded and satisfies } \Theta(\bar{v}). \tag{8}$$

Recall in this connection that we regard a star-type over Σ as a vector with entries in $\mathbb{N}$; the star-type in question is M-bounded if the sum of those entries is at most M. By way of maintaining some contact with familiar territory, the $\mathcal{FLC}$-formula $\exists_{[=M]}\zeta$ can be written in the new syntax as $\mathsf{Q}\langle \Sigma, M, \Theta \rangle \zeta$, where Θ is the single equation $v_1 + \cdots + v_J = M$. (Of course: this equation is just a conjunction of two linear inequalities, and so counts as a formula of Presburger arithmetic). Note that there is no existential quantification in this case.

We define $\mathcal{FLC}^{m+1}_{\mathrm{ep}}$ to be the set of formulas φ given by

$$\bigwedge_{s \in S} \forall^m (\mu_s \rightarrow \mathsf{Q}\langle \Sigma, M_s, \Theta_s \rangle \zeta_s) \wedge \bigwedge_{t \in T} \forall^m (\nu_t \rightarrow \forall \eta_t) \wedge \forall^{m+1} \theta, \tag{9}$$

where Σ is a relational signature, the μ_s and ν_t are quantifier-free $\mathcal{FLC}^m$-formulas over Σ, the ζ_s, η_t and θ are quantifier-free $\mathcal{FLC}^{m+1}$-formulas over Σ, the M_s are positive integers and the Θ_s are formulas of existential Presburger arithmetic with free variables corresponding to the fluted $(m+1)$-types over Σ. We remark that we have *defined* $\mathcal{FLC}^{m+1}_{\mathrm{ep}}$ directly interms of fomulas of the form (9), rather than establishing an analogue of Lemma 3 for a language extending $\mathcal{FLC}^{m+1}$. This is intentional: formulas in which $\mathsf{Q}\langle \Sigma, M_s, \Theta_s \rangle \zeta$ appears with negative polarity might not be *succinctly* expressible in the form (9). We define the *effective size* of φ, denoted $\#(\varphi)$, to be the quantity $\log(\|\varphi\|) + |S| + |T| + \log M + |\Sigma|$, where $M = \sum_{s \in S} M_s$. Informally: when measuring the effective size of φ, we do not mind if $\|\varphi\|$ becomes exponentially large, as long as Σ, S, T and the number of bits required to write the various M_s do not.

We stress that fluted ep-quantifiers give us no additional expressive power beyond the standard counting quantifiers. Indeed, if ζ is a quantifier-free formula of $\mathcal{FLC}^{m+1}$ over a signature Σ, and supposing the fluted $(m+1)$-types over Σ to be enumerated as $\tau_1, \ldots, \tau_J$, any formula $\mathsf{Q}\langle \Sigma, M, \Theta \rangle \zeta$ is logically equivalent to the huge disjunction

$$\bigvee \left\{ \bigwedge_{j=1}^{J} \exists_{[=\sigma(\tau_j)]}\tau_j \;\middle|\; \sigma \text{ is an } M\text{-bounded fluted } \zeta\text{-star-type over } \Sigma \text{ satisfying } \Theta \right\}. \quad (10)$$

However, fluted ep-quantifiers can be more compact, and we require the added strength of the following routine extension of Lemma 4.

▶ **Lemma 7.** *If φ is a satisfiable formula of $\mathcal{FLC}^2_{\text{ep}}$, then φ has a model of size bounded by an exponential function of $\#(\varphi)$. The satisfiability problem for $\mathcal{FLC}^2_{\text{ep}}$ is in NExpTime, measured in terms of the effective size of the input.*

Proof. For the first statement, let a formula φ of $\mathcal{FLC}^2_{\text{ep}}$ be given. By definition, φ is in the form (9) with $m = 1$. Now construct $\mathfrak{B}$ as in the proof of Lemma 4. Still we have $\mathfrak{B} \models \varphi$, since it does not matter whether the permitted star-types are specified by means of standard counting quantifiers or fluted ep-quantifiers. Further, $|B| \leq (M + 1)2^{|\Sigma|+1}$ and thus is bounded by an exponential function of $\#(\varphi)$. For the second statement, we may simply guess a structure $\mathfrak{B}$ subject to this size bound and check that it satisfies φ. It follows from Lemma 5 that, for any $b \in B$ and $s \in S$ such that $\mathfrak{B} \models \mu_s[b]$, we may check, in *non-deterministic* time bounded by an exponential function of $\#(\varphi)$, that $\text{fst}^{\mathfrak{B}}_{\zeta_s}[b]$, satisfies Θ_s. Checking the remaining conditions of φ involves standard model-checking, and requires only (deterministic) time bounded by an exponential function of $\#(\varphi)$. ◀

4 Proof of main result

In this section we prove Theorems 1 and 2, proceeding via the logics $\mathcal{FLC}^m_{\text{ep}}$. Fix some $\mathcal{FLC}^{m+1}_{\text{ep}}$-formula φ as given in (9), with $m \geq 2$. We show how to construct an $\mathcal{FLC}^m_{\text{ep}}$-formula φ', such that φ and φ' are satisfiable over the same domains. The formula φ' employs a signature Σ' formed by removing from Σ all $(m + 1)$-ary predicates, while adding a fresh $(m - 1)$-ary predicate $p_{S',T'}$ and a fresh m-ary predicate $q_{S',T'}$ for each $S' \subseteq S$ and each $T' \subseteq T$. It is obvious that $|\Sigma'| \leq |\Sigma| + 2^{|S|+|T|+1}$. The maximal arity of predicates in Σ is $m + 1$, and in Σ' is m. In addition to Σ and Σ', we consider the signatures $\Sigma^+ = \Sigma \cup \Sigma'$ and $\Sigma^- = \Sigma \cap \Sigma'$. As explained in Sec. 2, we assume some fixed enumeration $\tau_1, \ldots, \tau_J$ of the $J = 2^{|\Sigma|+1}$ fluted $(m + 1)$-types over Σ. We write $\tau^+_1, \ldots, \tau^+_{J^+}$ for the corresponding enumeration of fluted $(m + 1)$-types over Σ^+; likewise we enumerate the fluted m-types over Σ' as $\tau'_1, \ldots, \tau'_{J'}$, and over Σ^- as $\tau^-_1, \ldots, \tau^-_{J^-}$.

Our essential problem is to get rid of the $(m + 1)$-ary predicates appearing in φ without affecting satisfiability. The following device will help. Let ψ be any quantifier-free $\mathcal{FLC}^{m+1}$-formula over Σ. Clearly, there is, up to logical equivalence, a unique strongest quantifier-free $\mathcal{FLC}^m$-formula over Σ^- entailed by ψ, i.e. a quantifier-free $\mathcal{FLC}^m$-formula ψ° over Σ^- satisfying: (i) $\models \psi \to \psi^\circ$; and (ii) for all quantifier-free $\mathcal{FLC}^m$-formulas χ over Σ^- such that $\models \psi \to \chi$, we have $\models \psi^\circ \to \chi$. Indeed, since there are only finitely many such χ (ignoring logical equivalents), we can take ψ° to be their conjunction. Strictly, of course, ψ° is only defined up to logical equivalence, and in fact there are various ways to construct it. Thus, for example, [16, 17] employ resolution theorem-proving; however, the procedure in the following proof requires only basic propositional logic.

▶ **Lemma 8.** *Let ψ be a quantifier-free $\mathcal{FLC}^{m+1}$-formula over Σ. Then we can compute ψ° in time bounded by an exponential function of $\|\psi\| + |\Sigma|$. Moreover, if τ^- is a fluted m-type over Σ^- such that $\models \tau^- \to \psi^\circ$, then there exists a fluted $(m + 1)$-type τ over Σ extending τ^- such that $\models \tau \to \psi$.*

Proof. We begin by writing ψ, equivalently, in disjunctive normal form. Thus, $\psi \equiv \bigvee\{\tau \mid \tau \in D\}$, where D is a set of fluted $(m+1)$-types (over Σ). For each $\tau \in D$, let τ° be the fluted m-type over Σ^- obtained by deleting from τ all conjuncts involving predicates not in Σ^-, and define ψ° to be $\bigvee\{\tau^\circ \mid \tau \in D\}$. Note that, since by assumption $m \geq 2$, we will never delete equality-literals. It is evident that this construction can be carried out in time bounded by an exponential function of $\|\psi\| + |\Sigma|$. Since $\models \tau \to \tau^\circ$ for any $(m+1)$-type τ, it is immediate that $\models \psi \to \psi^\circ$. On the other hand, suppose χ is a quantifier-free $\mathcal{FLC}^m$-formula over Σ^- entailed by ψ; without loss of generality, χ is in disjunctive normal form. We claim that τ° is a disjunct of χ for all $\tau \in D$. For if not, we have $\models \psi \to \neg\tau^\circ$, whence $\models \psi \to \neg\tau$, contradicting the supposition that $\psi \equiv \bigvee\{\tau \mid \tau \in D\}$. Hence, $\models \psi^\circ \to \chi$. Given the above construction of ψ°, the second statement of the lemma is completely trivial. ◀

Also in this connection, we define a further operation on fluted star-types. Suppose that σ is an m-dimensional fluted star-type over Σ. The *reduct of* σ *to* Σ^-, denoted σ/Σ^-, is the $(m-1)$-dimensional fluted star-type over Σ^- given by

$$(\sigma/\Sigma^-)(\tau^-) = \sum\{\sigma(\tau) : \tau \text{ a fluted } (m+1)\text{-type over } \Sigma \text{ such that } \models \tau \to \tau^-\},$$

where τ^- is any fluted m-type over Σ^-. Thus, when forming a reduct to Σ^-, we merge together fluted $(m+1)$-types which look identical in the smaller signature.

To explain the intuition behind the construction of φ', we first describe how a putative model $\mathfrak{A} \models \varphi$ can be expanded to a structure, $\mathfrak{A}^+$, interpreting $\Sigma^+ = \Sigma \cup \Sigma'$. Taking $\mathfrak{A}'$ to be the reduct of $\mathfrak{A}^+$ to Σ' (i.e. with all the $(m+1)$-ary predicates removed), we observe in (13), (14) and (15) that the formulas which will eventually form the conjuncts of φ' are all true in $\mathfrak{A}'$. We begin with the predicates $p_{S',T'}$ for $S' \subseteq S$ and $T' \subseteq T$. Let $\bar{a} \in A^{m-1}$ be any $(m-1)$-tuple of elements, and consider what φ tells us about the relationship of $\bar{a}$ to other elements in the structure. Since φ is fluted, and bearing in mind the form (9), what really matters here are the different subsets $S' \subseteq S$ and $T' \subseteq T$ for which there exists an element a such that $\mathfrak{A} \models \mu_s[a\bar{a}]$ for all $s \in S'$ and $\mathfrak{A} \models \nu_t[a\bar{a}]$ for all $t \in T'$. The $(m-1)$-ary predicates $p_{S',T'}$ will simply record which pairs S', T' are realized in this way. That is, we set, for every $S' \subseteq S$ and $T' \subseteq T$,

$$p^{\mathfrak{A}^+}_{S',T'} = p^{\mathfrak{A}'}_{S',T'} = \{\bar{a} \in A^{m-1} : \text{for some } a \in A,\ \mathfrak{A} \models \mu_s[a\bar{a}] \text{ for all } s \in S' \text{ and } \mathfrak{A} \models \nu_t[a\bar{a}] \text{ for all } t \in T'\}. \quad (11)$$

The predicates $q_{S',T'}$ are only slightly more complicated. Fix $S' \subseteq S$ and $T' \subseteq T$ for the moment, and suppose that $\mathfrak{A}^+ \models p_{S',T'}[\bar{a}]$. By construction, there exists $a \in A$ such that $\mathfrak{A} \models \mu_s[a\bar{a}]$ for all $s \in S'$ and $\mathfrak{A} \models \nu_t[a\bar{a}]$, for all $t \in T'$. So pick *any* such a and denote it by $\dot{a}$. The formula φ then guarantees that, for each $s \in S'$, the m-tuple $\dot{a}\bar{a}$ satisfies the formula $\mathsf{Q}\langle\Sigma, M_s, \Theta_s\rangle\zeta_s$. Defining $B_{\bar{a}} = \{b \in A : \mathfrak{A} \models \zeta_s[\dot{a}\bar{a}b] \text{ for some } s \in S'\}$, we set

$$q^{\mathfrak{A}^+}_{S',T'} = q^{\mathfrak{A}'}_{S',T'} = \{\bar{a}b \in A^m : \mathfrak{A}^+ \models p_{S',T'}[\bar{a}] \text{ and } b \in B_{\bar{a}}\}. \quad (12)$$

Thus, $B_{\bar{a}}$ serves to pick out the witnesses required by the various fluted ep-quantifiers $\mathsf{Q}\langle\Sigma, M_s, \Theta_s\rangle$ for the tuple $\dot{a}\bar{a}$, as s varies over S'. Letting $\zeta_{S'} = \bigvee_{s \in S'} \zeta_s$, we see that $B_{\bar{a}}$ is the set of elements absorbing $\zeta_{S'}$-rays emitted by $\dot{a}\bar{a}$. Observe, in particular, that $|B_{\bar{a}}| \leq \sum_{S'} M_s$. The whole construction is illustrated in Fig. 1. Here we see, arranged in a horizontal strip, the m-tuple $\dot{a}\bar{a}$, which satisfies μ_s for all $s \in S'$ and ν_t for all $t \in T'$. The elements $b \in B_{\bar{a}}$ absorbing the $\zeta_{S'}$-rays emitted by $\dot{a}\bar{a}$, are taken to lie on the periphery of the fan-shaped region. Each of these elements b absorbs a ζ_s-ray, for at least one (in general

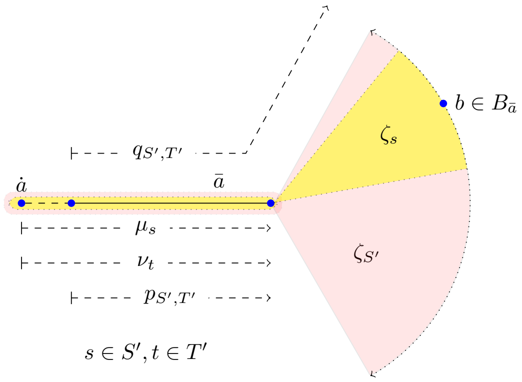

Figure 1 Intended interpretations of the predicates $p_{S',T'}$ and $q_{S',T'}$ in $\mathfrak{A}^+$ (and $\mathfrak{A}'$).

several) values of $s \in S'$. Discarding the element $\dot{a}$, we take the new predicate $p_{S',T'}$ to apply to the $(m-1)$-tuple $\bar{a}$ and the new predicate $q_{S',T'}$ to apply to the m-tuples $\bar{a}b$, where $b \in B_{\bar{a}}$.

We now construct the promised formula φ' in three steps, guided by the structure $\mathfrak{A}'$ just described. The first step is to write formulas reflecting the intended interpretations of the predicates $p_{S',T'}$. Indeed, we see immediately that, for all $S' \subseteq S$ and $T' \subseteq T$,

$$\mathfrak{A}' \models \forall^m \Big(\bigwedge_{s \in S'} \mu_s \wedge \bigwedge_{t \in T'} \nu_t \rightarrow p_{S',T'} \Big). \tag{13}$$

The second step is to write formulas reflecting the intended interpretations of the predicates $q_{S',T'}$, and specifically, the fact that they identify *all* $\zeta_{S'}$-witnesses for m-tuples of interest. Concretely, select an $(m-1)$-tuple $\bar{a}$ from A and suppose $\mathfrak{A} \models p_{S',T}[\bar{a}]$. Now let $\dot{a}$ and $B_{\bar{a}}$ be as chosen in the definition of $q^{\mathfrak{A}'}_{S',T'}$. Then $b \notin B_{\bar{a}}$ implies that $\dot{a}\bar{a}b$ does not satisfy ζ_s for any $s \in S'$. On the other hand, $\dot{a}\bar{a}b$ satisfies η_t for every $t \in T'$ as well as θ. Thus, writing, say, ψ for $(\bigwedge_{s \in S'} \neg\zeta_s \wedge \bigwedge_{t \in T'} \eta_t \wedge \theta)$, we see that $\mathfrak{A} \models \psi[\dot{a}\bar{a}b]$. Recalling that ψ° denotes the strongest quantifier-free formula over Σ^- entailed by ψ, we obviously have $\mathfrak{A}^- \models \psi^\circ[\dot{a}\bar{a}b]$, where $\mathfrak{A}^-$ denotes the reduct $\mathfrak{A}{\restriction}\Sigma^- = \mathfrak{A}'{\restriction}\Sigma^-$. But ψ° involves no predicates of arity $m+1$, whence $\mathfrak{A}^- \models \psi^\circ[\bar{a}b]$. (Observe how we are here exploiting variable-free notation: while ψ° is indeed a formula of $\mathcal{FLC}^{[m+1]}$, it involves only predicates in Σ^- , and therefore is simultaneously a formula of $\mathcal{FLC}^{[m]}$, which cannot "see" the element $\dot{a}$ in the tuple $\dot{a}\bar{a}b$.) Thus, since ψ° is also a formula over the signature Σ', we have shown that, for all $S' \subseteq S$ and $T' \subseteq T$,

$$\mathfrak{A}' \models \forall^{m-1} \big(p_{S',T'} \rightarrow \forall \big(\neg q_{S',T'} \rightarrow \big(\bigwedge_{s \in S'} \neg\zeta_s \wedge \bigwedge_{t \in T'} \eta_t \wedge \theta \big)^\circ \big) \big). \tag{14}$$

The final step in the construction of φ' requires us to define, for all subsets $S' \subseteq S$, a fluted ep-quantifier $\mathsf{Q}\langle \Sigma', M_{S'}, \Theta_{S'} \rangle$. To motivate the definition, let $\bar{a}$ again be any $(m-1)$-tuple from A, let S', T' be such that $\mathfrak{A}^+ \models p_{S',T'}[\bar{a}]$, and let $\dot{a}$ be the element selected in the definition of $q^{\mathfrak{A}^+}_{S',T'}$, so that we have

$$B_{\bar{a}} = \{ b \in A : \mathfrak{A} \models \zeta_s[\dot{a}\bar{a}b] \text{ for some } s \in S' \} = \{ b \in A : \mathfrak{A}' \models q_{S',T'}[\bar{a}b] \}.$$

Remembering that $\zeta_{S'} = \bigvee_{s \in S'} \zeta_s$, define the fluted star-types $\sigma = \mathrm{fst}^{\mathfrak{A}}_{\zeta_{S'}}[\dot{a}\bar{a}]$ (m-dimensional, over Σ) and $\sigma' = \mathrm{fst}^{\mathfrak{A}'}_{q_{S',T'}}[\bar{a}]$ ($(m-1)$-dimensional, over Σ'). Define in addition the $(m-1)$-dimensional fluted star-type over Σ^- by setting, for any fluted m-type τ^- over Σ^-,

$$\sigma^-(\tau^-) = |\{ b \in B_{\bar{a}} : \mathfrak{A}^- \models \tau^-[\bar{a}b] \}|.$$

We see immediately by consideration of the set $B_{\bar{a}}$ that $\sigma/\Sigma^- = \sigma^- = \sigma'/\Sigma^-$. Warning: it will not in general be the case that σ^- arises as $\text{fst}^{\mathfrak{A}^-}_{\zeta^-}[\bar{a}]$ for any quantifier-free formula ζ^- over the signature Σ^-. Indeed, $\mathfrak{A}^-$ interprets neither the predicate $q_{S',T'}$ nor the predicates of $\zeta_{S'}$, and is therefore insensitive to the extension of $B_{\bar{a}}$ used to define σ^-. Nevertheless, we have established:

$$\sigma/\Sigma^- = \sigma'/\Sigma^-. \tag{L1}$$

A little thought shows that σ satisfies various further properties. Fix some $s \in S'$. Since $\models \zeta_s \to \zeta_{S'}$, we see that the retract $\sigma\restriction\zeta_s$ is equal to the fluted star-type $\text{fst}^{\mathfrak{A}}_{\zeta_s}[\dot{a}\bar{a}]$ and hence (from the fact that $\mathfrak{A} \models \varphi$), is M_s-bounded and satisfies Θ_s. Thus, we have:

$$\text{for all } s \in S',\ \sigma\restriction\zeta_s \text{ is } M_s\text{-bounded and satisfies } \Theta_s. \tag{L2}$$

Now fix some $t \in T'$. Since $\mathfrak{A} \models \varphi$, it follows immediately that, for every $b \in B_{\bar{a}}$, $\mathfrak{A} \models \eta_t[\dot{a}\bar{a}b]$, and indeed $\mathfrak{A} \models \theta[\dot{a}\bar{a}b]$, whence:

$$\text{for all } t \in T \text{ and all } \tau \text{ such that } \sigma(\tau) > 0,\ \models \tau \to \eta_t. \tag{L3}$$

$$\text{for all } \tau \text{ such that } \sigma(\tau) > 0,\ \models \tau \to \theta. \tag{L4}$$

Casting this discussion in terms of $\sigma' = \text{fst}^{\mathfrak{A}'}_{q_{S',T'}}[\bar{a}]$ and writing $M_{S'} = \sum_{s \in S'} M_s$, we see that σ' is $M_{S'}$-bounded and satisfies the property that there exists a fluted star-type σ such that (L1)–(L4) hold. Crucially, this property can be naturally formulated using a formula of existential Presburger arithmetic. Letting $\bar{v}$ be a tuple of variables corresponding to the fluted $(m+1)$-types over Σ and $\bar{v}'$ a tuple of variables corresponding to the fluted m-types over Σ', we see that (L1) is a system of equations $A\bar{v} = A'\bar{v}'$, where A, A' are matrices with entries in $\{0,1\}$ depending only on the fixed ordering of the fluted $(m+1)$-types over Σ and the fixed ordering of the fluted m-types over Σ' and Σ^-. And certainly, (L3) and (L4) can be expressed as a single equation setting certain values in $\bar{v}$ to zero. Let us write $\Lambda(\bar{v}, \bar{v}')$ for the conjunction of all the equations expressing (L1), (L3) and (L4). Considering that any retract $\sigma\restriction\zeta_s$ amounts to the zeroing of certain entries in σ, we may assume without loss of generality that the corresponding variables do not occur in Θ_s. (If they do, we may replace them by 0.) And in that case, (L2) is expressed by the conjunction $\bigwedge_{s\in S'} \Theta_s(\bar{v})$. Thus, we may formulate the above conditions on σ' as the requirement that (considered as a vector $\bar{v}'$ of length J' over $\mathbb{N}$) it satisfies the formula of Presburger arithmetic $\exists\bar{v}\big(\Lambda(\bar{v},\bar{v}') \wedge \bigwedge_{s\in S'} \Theta_s(\bar{v})\big)$. Writing $\Theta_s(\bar{v})$ as $\exists\bar{u}_s.\Xi_s(\bar{u}_s, \bar{v})$ for each $s \in S$, we obtain, by renaming variables to avoid clashes, the equivalent existential Presburger formula $\Theta_{S'}(\bar{v}') \equiv \exists\bar{v}\bar{u}\big(\Lambda(\bar{v},\bar{v}') \wedge \bigwedge_{s\in S'} \Xi_s(\bar{u}_s, \bar{v})\big)$, where $\bar{u}$ is the concatenation of the (disjoint) tuples $\bar{u}_s$ for $s \in S'$. Thus, we have shown:

$$\mathfrak{A}' \models \forall^{m-1}(p_{S',T'} \to \mathsf{Q}\langle\Sigma', M_{S'}, \Theta_{S'}\rangle q_{S',T'}). \tag{15}$$

Now we are ready to define φ' as the conjunction of the following formulas:

$$\bigwedge_{S'\subseteq S}\bigwedge_{T'\subseteq T} \forall^m \Big(\bigwedge_{s\in S'} \mu_s \wedge \bigwedge_{t\in T'} \nu_t \to p_{S',T'}\Big) \tag{16}$$

$$\bigwedge_{S'\subseteq S}\bigwedge_{T'\subseteq T} \forall^{m-1}(p_{S',T'} \to \mathsf{Q}\langle\Sigma', M_{S'}, \Theta_{S'}\rangle q_{S',T'}) \tag{17}$$

$$\bigwedge_{S'\subseteq S}\bigwedge_{T'\subseteq T} \forall^{m-1}\Big(p_{S',T'} \to \forall\big(\neg q_{S',T'} \to \big(\bigwedge_{s\in S'} \neg\zeta_s \wedge \bigwedge_{t\in T'} \eta_t \wedge \theta\big)^\circ\big)\Big). \tag{18}$$

By re-ordering of conjuncts, we see that φ' is an $\mathcal{FLC}^m_{\text{ep}}$-formula. Moreover, it follows from (13), (14) and (15) that $\mathfrak{A}' \models \varphi'$. Hence, we have proved:

▶ **Lemma 9.** *If φ is satisfiable over some domain A, then so is φ'.*

We now prove a converse to Lemma 9. Suppose $\mathfrak{B}' \models \varphi'$, where $\mathfrak{B}'$ has domain B. We proceed to construct a model $\mathfrak{B} \models \varphi$ over the same domain. Let $\mathfrak{B}^- = \mathfrak{B}' \restriction \Sigma^-$. Notice that $\mathfrak{B}^-$ features no predicates of arity $m+1$, and none of the "new" predicates $p_{S',T'}$ or $q_{S',T'}$. We shall expand $\mathfrak{B}^-$ to a structure $\mathfrak{B}$ and then show that $\mathfrak{B} \models \varphi$. It suffices to specify, for every $a, b \in B$ and $\bar{a} \in B^{m-1}$, whether the tuple $a\bar{a}b$ is in the extension of each $(m+1)$-ary predicate of Σ. Equivalently, we must specify the fluted $(m+1)$-type of every tuple $a\bar{a}b$ in $\mathfrak{B}$.

Fix $a \in B$ and $\bar{a} \in B^{m-1}$. Let $S' = \{s \in S \mid \mathfrak{B}^- \models \mu_s[a\bar{a}]\}$, $T' = \{t \in T \mid \mathfrak{B}^- \models \nu_t[a\bar{a}]\}$ and $\zeta_{S'} = \bigvee_{s \in S'} \zeta_s$. Since all the μ_s and ν_t are Σ^--formulas, we could equivalently have replaced $\mathfrak{B}^-$ by $\mathfrak{B}'$ in the definitions of S' and T'. Define $\sigma' = (v'_1, \ldots, v'_{J'}) = \mathrm{fst}^{\mathfrak{B}'}_{q_{S',T'}}[\bar{a}]$ and $B_{a\bar{a}} = \{b \in B \mid \mathfrak{B}' \models q_{S',T'}[\bar{a}b]\}$. Notice that, since the sets S' and T' depend on a as well as $\bar{a}$, then so does the set $B_{a\bar{a}}$. It follows from (16) that $\mathfrak{B}' \models p_{S',T'}[\bar{a}]$, and from (17), that σ' is $M_{S'}$-bounded and satisfies the existential Presburger formula $\Theta_{S'}(\bar{v}') \equiv \exists \bar{v} \bar{u} \big(\Lambda(\bar{v}, \bar{v}') \wedge \bigwedge_{s \in S'} \Xi_s(\bar{u}_s, \bar{v})\big)$. But $\Theta_{S'}(\bar{v}')$ asserts that $\bar{v}'$ is the vector representation of a fluted star-type σ' (over Σ') for which there exists an m-dimensional fluted star-type σ (over Σ) satisfying conditions (L1)–(L4). Letting $\sigma^- = \sigma'/\Sigma^-$, it follows from (L1) that $\sigma^- = \sigma/\Sigma^-$. We proceed to set the fluted $(m+1)$-type of all tuples $a\bar{a}b$, as b ranges over B; we shall do this in such a way that $\mathrm{fst}^{\mathfrak{B}}_{\zeta_{S'}}[a\bar{a}] = \sigma$. The plan is first to find all the required witnesses in the set $B_{a\bar{a}}$, and then to ensure that no unwanted witnesses appear outside this set.

We begin with the elements $b \in B_{a\bar{a}}$. We first partition $B_{a\bar{a}}$ into groups which are indistinguishable from the point of view of the signature Σ^-. Specifically, we write $\sigma^- = (v_1^-, \ldots, v_{J^-}^-)$, and for each j^- $(1 \leq j^- \leq J^-)$, we let $B_{j^-} = \{b \in B_{a\bar{a}} \mid \mathfrak{B}^- \models \tau^-_{j^-}[\bar{a}, b]\}$. Writing $J^-_{a\bar{a}}$ for the set of indices j^- for which B_{j^-} is non-empty, we see that $|B_{j^-}| = v_{j^-}$ for all j^- $(1 \leq j^- \leq J^-)$, $v_{j^-} = 0$ for all $j^- \notin J^-_{a\bar{a}}$, and the family of sets $\{B_{j^-} \mid j^- \in J^-_{a\bar{a}}\}$ forms a partition of $B_{a\bar{a}}$.

Now consider just one cell of this partition, say B_{j^-}. We have

$$B_{j^-} = \{b \in B_{a\bar{a}} \mid \mathfrak{B}^- \models \tau^-_{j^-}[\bar{a}, b]\} = \{b \in B_{a\bar{a}} \mid \mathfrak{B}' \models \tau_{j'}[\bar{a}, b] \text{ for some } j' \text{ s.t. } \models \tau'_{j'} \to \tau^-_{j^-}\}.$$

And since $v^-_{j^-} = |B_{j^-}|$, we obtain

$$v^-_{j^-} = \sum\{v_{j'} \mid 1 \leq j' \leq J' \text{ and } \models \tau'_{j'} \to \tau^-_{j^-}\} = \sum\{v_j \mid 1 \leq j \leq J \text{ and } \models \tau_j \to \tau^-_{j^-}\},$$

the second equality arising from the fact that $\sigma/\Sigma^- = \sigma'/\Sigma^-$. Thus, we may choose, for each j such that $\models \tau_j \to \tau^-_{j^-}$, a fresh collection $B_{j^-,j}$ of v_j elements of B_{j^-}, and for each of these elements, b, set $\mathrm{ftp}^{\mathfrak{B}}[a\bar{a}b] = \tau_j$. Because $\models \tau_j \to \tau^-_{j^-}$, the only predicates being defined afresh here have arity $(m+1)$, so that these assignments represent an expansion of $\mathfrak{B}^-$. Once these assignments are made, the set B_{j^-} will contain v_j elements b such that $\mathrm{ftp}^{\mathfrak{B}}[a\bar{a}b] = \tau_j$ for all j such that $\models \tau_j \to \tau^-_{j^-}$. Repeating this procedure for every $j^- \in J^-_{a\bar{a}}$, and writing $J_{a\bar{a}} = \{j \mid 1 \leq j \leq J \text{ and } \models \tau_j \to \tau^-_{j^-} \text{ for some } j^- \in J^-_{a\bar{a}}\}$, we see that the set $B_{a\bar{a}}$ will contain v_j elements b such that $\mathrm{ftp}^{\mathfrak{B}}[a\bar{a}b] = \tau_j$ for all $j \in J_{a\bar{a}}$. On the other hand, $\sigma' = (v'_1, \ldots, v'_{J'}) = \mathrm{fst}^{\mathfrak{B}'}_{q_{S',T'}}[\bar{a}]$, so that $j^- \notin J^-_{a\bar{a}}$ implies $v^-_{j^-} = 0$ and hence

$$\sum\{v_j \mid 1 \leq j \leq J \text{ and } \models \tau_j \to \tau^-_{j^-}\} = 0,$$

since $\sigma \restriction \Sigma^- = \sigma^-$. That is, $v_j = 0$ for all $j \notin J_{a\bar{a}}$, and we have shown that $B_{a\bar{a}}$ contains v_j elements b such that $\mathrm{ftp}^{\mathfrak{B}}[a\bar{a}b] = \tau_j$, for all j $(1 \leq j \leq J)$.

Next, we deal with the elements $b \in B \setminus B_{a\bar{a}}$. By definition, $\mathfrak{B}' \not\models q_{S',T'}[\bar{a}b]$, so that, abbreviating $\bigwedge_{s \in S'} \neg\zeta_s \wedge \bigwedge_{t \in T'} \eta_t \wedge \theta$ by ψ, (18) yields $\mathfrak{B}' \models \psi^\circ[\bar{a}b]$. Writing $\tau^- = \text{ftp}^{\mathfrak{B}^-}[\bar{a}, b]$, and noting that ψ° is a Σ^--formula, we have $\models \tau^- \to \psi^\circ$. By Lemma 8, let $\tilde{\tau}$ be a fluted $(m+1)$-type over Σ extending τ^- such that $\models \tilde{\tau} \to \psi$, and set $\text{ftp}^{\mathfrak{B}}[a\bar{a}b] = \tilde{\tau}$. In particular, $\mathfrak{B} \not\models \zeta_{S'}[a\bar{a}b]$. Thus, we have set the fluted $(m+1)$-type of all $(m+1)$-tuples $a\bar{a}b$, as b ranges over B. Repeating this process for each m-tuple $a\bar{a}$, which we may do independently, the structure $\mathfrak{B}$ will have been completely defined. In the course of this construction, we have shown that, for every $b \notin B_{a\bar{a}}$, $\mathfrak{B} \not\models \zeta_{S'}[a\bar{a}b]$. But we showed above that the number of elements $b \in B_{a\bar{a}}$ such that $\mathfrak{B} \models \tau_j[a\bar{a}b]$ is v_j, where $\sigma = (v_1, \ldots, v_J)$. It follows that $\text{fst}^{\mathfrak{B}}_{\zeta_{S'}}[a\bar{a}] = \sigma$, as required.

In constructing $\mathfrak{B}$, we have secured the following property. Take any m-tuple $a\bar{a}$ with $a \in B$ and $\bar{a} \in B^{m-1}$, and define $S' = \{s \in S \mid \mathfrak{B}^- \models \mu_s[a\bar{a}]\}$, $T' = \{t \in T \mid \mathfrak{B}^- \models \nu_t[a\bar{a}]\}$, $\zeta_{S'} = \bigvee_{s \in S'} \zeta_s$ and $\sigma' = \text{fst}^{\mathfrak{B}'}_{q_{S',T'}}[\bar{a}]$. Then $\text{fst}^{\mathfrak{B}}_{\zeta_{S'}}[a\bar{a}] = \sigma$, where σ is some fluted star-type over Σ satisfying (L1)–(L4), whose existence is guaranteed by the fact that σ' satisfies the existential Presburger formula $\Theta_{S'}$. We have used (L1) in the construction of $\mathfrak{B}$. We now use (L2)–(L4) to ensure that $\mathfrak{B} \models \varphi$.

We first show that, for all $s \in S$, $\mathfrak{B} \models \forall^m(\mu_s \to \mathsf{Q}\langle\Sigma, M_s, \mathcal{L}_s\rangle\zeta_s)$. For consider any m-tuple $a\bar{a}$ with $a \in B$ and $\bar{a} \in B^{m-1}$ such that $\mathfrak{B} \models \mu_s[a\bar{a}]$, and let S', $\zeta_{S'}$, and σ be as just defined. Since $s \in S'$, we have $\models \zeta_s \to \zeta_{S'}$, and hence $\text{fst}^{\mathfrak{B}}_{\zeta_s}[a\bar{a}] = \sigma\restriction\zeta_s$, which, by (L2), is M_s-bounded and satisfies $\Theta_{S'}$. We next show that, for all $t \in T$, $\mathfrak{B} \models \forall^m(\nu_t \to \forall\eta_t)$. Again, consider any m-tuple $a\bar{a}$ with $a \in B$ and $\bar{a} \in B^{m-1}$ such that $\mathfrak{B} \models \nu_t[a\bar{a}]$, and let S', T' and σ be as just defined. Pick any $b \in B$. If b is in the set $B_{a\bar{a}}$ used in the construction of $\mathfrak{B}$, then $\sigma(\text{ftp}^{\mathfrak{B}}[a\bar{a}b]) > 0$, whence by (L3), we have $\mathfrak{B} \models \eta_t[a\bar{a}b]$. On the other hand, if $b \notin B_{a\bar{a}}$, then $\text{ftp}^{\mathfrak{B}}[a\bar{a}b]$ was set to some $\tilde{\tau}$ entailing $\bigwedge_{s \in S'} \neg\zeta_s \wedge \bigwedge_{t \in T'} \eta_t \wedge \theta$, so that, again $\mathfrak{B} \models \eta_t[a\bar{a}b]$. Finally, to show that $\mathfrak{B} \models \forall^{m+1}\theta$, we proceed as in the previous case, but using (L4) instead of (L3). Thus we have proved:

▶ **Lemma 10.** *If φ' is satisfiable over some domain B, then so is φ.*

▶ **Lemma 11.** *The formula φ' can be computed in time bounded by an exponential function of $\|\varphi\|$. Moreover, $\#(\varphi')$ is bounded by an exponential function of $\#(\varphi)$.*

Proof. Recalling that φ has the form

$$\bigwedge_{s \in S} \forall^m(\mu_s \to \mathsf{Q}\langle\Sigma, M_s, \Theta_s\rangle\zeta_s) \wedge \bigwedge_{t \in T} \forall^m(\nu_t \to \forall\eta_t) \wedge \forall^{m+1}\theta$$

over signature Σ', let $M = \sum_{s \in S} M_s$. Writing φ' given by (16)–(18) in the same form, over signature Σ', we notice first of all that the sizes of the index sets certainly only increase by an exponential. Let M' be the sum of all the numbers occurring in the fluted ep-quantifiers of φ', i.e. $\sum\{2^{|T|} \cdot M_{S'} \mid S' \subseteq S\}$, whence $\log M' \leq |S| + |T| + \log M \leq \#(\varphi)$. We noted above that $|\Sigma'| \leq |\Sigma| + 2^{|S|+|T|+1} \leq \#(\varphi) + 2^{\#(\varphi)+1}$. To show that $\log\|\varphi'\|$ is also bounded by an exponential function of $\#(\varphi)$, we need only show that $\log\|\Theta_{S'}\|$ is bounded by an exponential function of $\#(\varphi)$, for each $S' \subseteq S$. Fix some $S' \subseteq S$, then, and recall that

$$\Theta_{S'}(\bar{v}') \equiv \exists\bar{v}\{\bar{u}_s\}_{s \in S'}\big(\Lambda(\bar{v}, \bar{v}') \wedge \bigwedge_{s \in S'} \Xi_s(\bar{u}_s, \bar{v})\big).$$

Let e be the maximum number of existentially quantified variables in any Θ_s as s ranges over S; certainly, $e \leq \|\varphi\|$. The number of free variables in $\Theta_{S'}$ is simply the number of fluted m-types over Σ' and hence at most $2^{|\Sigma'|+1}$. The number of existentially quantified variables in in $\Theta_{S'}$ is bounded by $|S'|.e + 2^{|\Sigma|+1}$. Moreover, Λ consists of $2^{|\Sigma^-|+1}$ equations featuring at most $2^{|\Sigma|+1} + 2^{|\Sigma'|+1}$ terms. Finally we evidently have $\|\bigwedge_{s \in S'} \Theta_s\| \leq |S'| \cdot \|\varphi\|$. Adding all of these together, we see that $\log\|\Theta_{S'}\|$ is bounded by an exponential function of $\#(\varphi)$, as required. ◀

▶ **Lemma 12.** *Let $m \geq 1$. Any satisfiable formula of $\mathcal{FLC}^{m+1}_{\text{ep}}$ has a finite model. Moreover, the satisfiability problem for $\mathcal{FLC}^{m+1}_{\text{ep}}$ is in m-NExpTime, measured in terms of the effective size of the input.*

Proof. The case $m = 1$ is Lemma 4. The inductive step is from Lemmas 9, 10 and 11. ◀

We may now prove Theorem 1. Let a formula φ of $\mathcal{FLC}^{m+1}$ be given. By Lemma 3, we may suppose without loss of generality that φ is in normal form (7). Further, we may replace any sub-formula $\exists_{[=M_s]}\zeta_s$ in φ by the equivalent sub-formula $\mathsf{Q}\langle\Sigma, M_s, \Theta_s\rangle\zeta_s$, where Θ_s is simply the equation $v_1 + \cdots + v_J = M_s$, thus obtaining, in time bounded by an exponential function of $\|\varphi\|$, an $\mathcal{FLC}^{m+1}_{\text{ep}}$-formula φ_0 equivalent to φ. Note that $\#(\varphi_0)$ is bounded by a polynomial function of $\|\varphi\|$. The result then follows from Lemma 12.

5 The semi-fluted fragment

Finally, we consider the semi-fluted fragment with counting, $\mathcal{SFC}$. The analysis proceeds largely as for $\mathcal{FLC}$. We remind the reader that variable-free notation is no longer available in this case. Of course, we could employ predicate-functor-style syntax here; however, for the little material remaining, this seems excessive.

Let Σ be a purely relational signature and m a positive integer. A *semi-fluted m-atom* over Σ is a formula of the form $p(\bar{x})$ or its negation where either: (i) $p \in \Sigma$ and $\bar{x}$ is a contiguous sequence of variables $x_\ell, \ldots, x_m$, or (ii) $p \in \Sigma \cup \{=\}$ and $\bar{x}$ is a sequence of at most two variables chosen (repeats allowed) from the set $\{x_{m-1}, x_m\}$. A *semi-fluted m-literal* is either a semi-fluted m atom or its negation. Thus, we have the same restriction on argument patterns as in fluted logic generally, except that semi-fluted m-literals of arity at most 2 may feature the variables x_{m-1} and x_m in any order we like. A *semi-fluted m-type* (over Σ) is a maximal consistent set of semi-fluted m-literals (over Σ). For these purposes, consistency takes account of the special meaning of the equality predicate: thus, $\{p(x_1), \neg p(x_2), x_1 = x_2\}$ is not consistent. As before, where convenient, we identify a semi-fluted m-type τ with the conjunction of its members and call τ *reflexive* if it contains $x_{m-1} = x_m$. We remark that semi-fluted 1- and 2-types are simply maximal consistent sets of literals (atomic formulas or their negations) in the variables x_1 and x_2, and are usually referred to in the literature simply as 1- and 2-*types*. If τ is a semi-fluted literal of arity $m \geq 2$, define $\text{tp}_1(\tau)$ to be the set of literals in τ featuring only the variable x_{m-1}. (Note that replacing the variable x_{m-1} in $\text{tp}_1(\tau)$ by x_1 would yield a 1-type.) A *semi-fluted star-type of dimension m over Σ* is a multiset of semi-fluted $(m+1)$-types over Σ at most one of which is reflexive, subject to the additional condition that the value $\text{tp}_1(\tau)$ is the same for all τ occurring (i.e. having non-zero multiplicity) in σ. A semi-fluted star-type is M-bounded if the sum of its multiplicities is M. By enumerating the semi-fluted $(m+1)$-types over Σ in some fixed order, we may regard any semi-fluted star-type as a vector of cardinal numbers.

If $\mathfrak{A}$ is a structure interpreting Σ, and $\bar{a}$ is an m-tuple of elements from A (repeats allowed), there is a unique semi-fluted m-type satisfied by $\bar{a}$; we denote this by $\text{sftp}^{\mathfrak{A}}[\bar{a}]$. If, in addition, ζ is a quantifier-free formula of $\mathcal{SFC}^{m+1}$, then we may define a semi-fluted star-type σ of dimension m by setting, for each semi-fluted $(m+1)$-type τ over Σ, $\sigma(\tau) = |\{b \in A : \mathfrak{A} \models \tau[\bar{a}b] \text{ and } \mathfrak{A} \models \zeta[\bar{a}b]\}|$. Notice incidentally that for all τ occurring in σ, we have $\mathfrak{A} \models \text{tp}_1(\tau)[a]$, where a is the last element of $\bar{a}$, whence these 1-types are indeed all the same. We call σ the *semi-fluted ζ-star-type of $\bar{a}$ in $\mathfrak{A}$*, and denote it by $\mathrm{sfst}^{\mathfrak{A}}_{\zeta}[\bar{a}]$. As with fluted ζ-star-types, so with their semi-fluted cousins, we can think of the tuple $\bar{a}$ as "emitting" various "ζ-rays", which are "absorbed" by various elements of $\mathfrak{A}$. The principal difference is that the (non-empty) semi-fluted ζ-star-type of $\bar{a}$ gives us the 1-type satisfied by the final element a of $\bar{a}$, and indeed of the full 2-types (not just the fluted 2-types) of the pairs ab, for b an element absorbing one of the ζ-rays emitted by $\bar{a}$.

Normal forms analogous to those of Lemma 3 are easily obtainable:

▶ **Lemma 13.** *Let φ be a formula of $\mathcal{SFC}^{m+1}$ ($m \geq 1$). Then we may compute, in time bounded by a polynomial function of $\|\varphi\|$, an $\mathcal{SFC}^{m+1}$-sentence satisfiable over the same domains as φ, and having the form*

$$\bigwedge_{s\in S} \forall x_1 \dots \forall x_m(\mu_s \to \exists_{[=M_s]} x_{m+1}.\zeta_s) \wedge \bigwedge_{t\in T} \forall x_1 \dots \forall x_m(\nu_t \to \forall \eta_t) \wedge \forall x_{m+1}.\theta, \qquad (19)$$

where S and T are index sets, the μ_s and ν_t are quantifier-free $\mathcal{SFC}^m$-formulas, the ζ_s, η_t and θ are quantifier-free $\mathcal{SFC}^{m+1}$-formulas, and the M_s are positive integers.

The proof proceeds as for Lemma 3, almost verbatim.

Extending the notion of ep-quantifiers to the semi-fluted case is again routine. If ζ is a quantifier-free formula of $\mathcal{SFC}^{m+1}$, we allow formulas φ of the form $\mathsf{Q}\langle\Sigma, M, \Theta\rangle\zeta$. For any structure $\mathfrak{A}$ interpreting a signature $\Sigma' \supseteq \Sigma$ and any m-tuple $\bar{a}$ of elements from A, we declare:

$$\mathfrak{A} \models \varphi[\bar{a}] \text{ if and only if } \mathrm{sfst}^{(\mathfrak{A}\restriction\Sigma)}_{\zeta}[\bar{a}] \text{ is } M\text{-bounded and satisfies } \Theta(\bar{v}), \qquad (20)$$

just as with (8). We then define $\mathcal{SFC}^{m+1}_{\mathrm{ep}}$ to be the set of formulas φ given by

$$\bigwedge_{s\in S} \forall x_1 \dots \forall x_m(\mu_s \to \mathsf{Q}\langle\Sigma, M_s, \Theta_s\rangle x_{m+1}.\zeta_s)\wedge$$
$$\bigwedge_{t\in T} \forall x_1 \dots \forall x_m(\nu_t \to \forall x_{m+1}.\eta_t) \wedge \forall x_1 \dots \forall x_{m+1}.\theta, \qquad (21)$$

where Σ is a relational signature, the μ_s and ν_t are quantifier-free $\mathcal{SFC}^m$-formulas over Σ, the ζ_s, η_t and θ are quantifier-free $\mathcal{SFC}^{m+1}$-formulas over Σ, the M_s are positive integers and the Θ_s are formulas of existential Presburger arithmetic with free variables corresponding to the semi-fluted $(m+1)$-types over Σ. Of course, this parallels the definition of $\mathcal{FLC}^{m+1}_{\mathrm{ep}}$ given in (9). Again, semi-fluted ep-quantifiers give us no expressive power beyond the ordinary counting quantifiers, since the translation (10) (with the counting quantifier taken to bind the variable x_{m+1}) holds also when ζ is only semi-fluted. We may define the *effective size*, $\#(\varphi)$ of an $\mathcal{SFC}^m_{\mathrm{ep}}$-formula φ exactly as with $\mathcal{FLC}^m_{\mathrm{ep}}$-formulas.

At this point, we are in a position to sketch the proof of Theorem 2. Let an $\mathcal{SFC}^{m+1}$-formula φ be given ($m \geq 1$). By Lemma 13, we may assume without loss of generality that φ is in the form (19). Clearly, this may be converted, in time bounded by an exponential function of $\|\varphi\|$, and with at most a polynomial increase in $\#(\varphi)$, into an $\mathcal{SFC}^{m+1}_{\mathrm{ep}}$-formula of the form (21). The reduction described in Sec. 4 can then be repeated almost verbatim, since the transformation of φ into φ' never affects the two final variables. Lemmas 9, 10 and 11 then continue to hold. This allows us to transform any formula of $\mathcal{SFC}^{m+1}$ ($m \geq 2$) eventually into a formula of $\mathcal{SFC}^2_{\mathrm{ep}}$ satisfiable over the same domains. But we have already argued that any $\mathcal{SFC}^2_{\mathrm{ep}}$-formula can be translated into equivalent formula of $\mathcal{C}^2$. Since the satisfiability and finite satisfiability problems for $\mathcal{C}^2$ are in NExpTime, we thus obtain Theorem 2, as promised.

The complexity bound we may extract from the above argument is rather weak. The translation given above from a $\mathcal{SFC}^2_{\text{ep}}$-formula φ to an equivalent $\mathcal{C}^2$-formula ψ' runs in time bounded by a doubly exponential function of $\|\varphi\|$, and hence a triply exponential function of $\#(\varphi)$. Since the satisfiability and finite satisfiability problems for $\mathcal{C}^2$ are in NExpTime, this means that the corresponding problems for $\mathcal{SFC}^2_{\text{ep}}$ are in non-deterministic time bounded by quadruply exponential time as a function of $\#(\varphi)$. This results in an upper complexity bound of $(m+3)$-NExpTime for the satisfiability and finite satisfiability problems for $\mathcal{SFC}^{m+1}_{\text{ep}}$ $(m \geq 1)$. We claim that the satisfiability and finite satisfiability problems for $\mathcal{SFC}^{m+1}$ in fact remain in m-NExpTime. However, the proof appears to require a modified version of existing proofs of the complexity bounds for $\mathcal{C}^2$, in order to accommodate semi-fluted ep-quantifiers. Such a reconstruction is beyond the scope of the current paper.

6 Discussion

For $m \geq 2$, the upper complexity bound of m-NExpTime for $\mathcal{FLC}^{m+1}$ in Theorem 1 is laxer than the corresponding upper complexity bound of $(m-1)$-NExpTime for $\mathcal{FL}^{m+1}$ from [17]. The best known lower complexity bound on satisfiability for both logics is $\lfloor (m+1)/2 \rfloor$-NExpTime-hard, from the same source. It is currently not known how to close this gap. It is plausible that, for $m \geq 2$, the upper bound given in this extended abstract for $\mathcal{FLC}^{m+1}$ could be reduced by one exponential, by adapting the procedure of in [17] for $\mathcal{FL}^3$. The probable difficulty of doing so coupled with the fact that a complexity gap would remain for $\mathcal{FLC}^5$ and above acts, however, as a deterrent to trying.

It is shown in [18] that the satisfiability and finite satisfiability problems for the fluted fragment, $\mathcal{FL}$ remain decidable even in the presence of a distinguished binary predicate required to be interpreted as a transitive relation (equality is also permitted); with just two transitive relations (or three transitive relations without equality), however, decidability is lost. The question arises as to whether a single transitive relation can be added to $\mathcal{FLC}$ without losing decidability of satisfiability. The argumentation of Sec. 4 will reduce this problem (with blow-up given by the same towers of exponentials) to the corresponding problem for $\mathcal{FLC}^2$ with a single transitive relation. However, this latter problem appears to be open.

We noted above that the fluted ep-quantifiers introduced here do not extend the expressive power of ordinary counting quantifiers. In this regard, they do less work than the "existential Presburger formulas" of [3], which strictly extend the expressive power of $\mathcal{C}^2$, permitting, saliently, counting *modulo k* for $k \geq 2$. The objects referred to as "*behaviours*" in that paper play a role very similar to the fluted star-types considered here, except that the values they assign are not integers, but *semi-linear sets* of integers. Unifying these approaches seems therefore to be a natural line of future enquiry.

References

1 H. Andréka, J. van Benthem, and I. Németi. Modal languages and bounded fragments of predicate logic. *Journal of Philosophical Logic*, 27:217–274, 1998.
2 V. Bárány, B. ten Cate, and L. Segoufin. Guarded negation. *Journal of the ACM*, 62(3):22:1–26, 2015.
3 M. Benedikt, E. Kostylev, and T. Tan. Two-variable logic with ultimately periodic counting. In A. Dawar A. Czumaj and E. Merelli, editors, *Proc. of ICALP 2020*, LIPIcs, pages 112:1–112:16. Schloß Dagstuhl - Leibniz-Zentrum für Informatik, 2020.
4 I. Borosh and L. Treybig. Bounds on the positive integral solutions of linear Diophantine equations. *Proceedings of the American Mathematical Society*, 55(2):299–304, 1976.

5 E. Grädel. On the restraining power of guards. *Journal of Symbolic Logic*, 64:1719–1742, 1999.
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 C. Haase. A survival guide to Presburger arithmetic. *ACM SIGLOG News*, 5(3):67–82, 2018.
8 D. Hilbert and W. Ackermann. *Grundzüge der mathematischen Logik*. Springer, Berlin, 1st edition, 1928.
9 D. Hilbert and W. Ackermann. *Principles of Mathematical Logic*. Chelsea, New York, 1950.
10 E. Kieroński and S. Rudolph. Finite model theory of the triguarded fragment and related logics. To appear, *Logic in Computer Science* (LICS), 2021; arXiV preprint `arXiv:2101.08377v2`.
11 L. Löwenheim. Über Möglichkeiten im Relativkalkül. *Math. Annalen*, 76:447–470, 1915.
12 L. Pacholski, W. Szwast, and L. Tendera. Complexity results for two-variable first-order logic with counting. *SIAM Journal of Computing*, 29:1083–1117, 1999.
13 C. Papadimitriou. On the complexity of integer programming. *Journal of the ACM*, 28(4):765–768, 1981.
14 L. Pottier. Minimal solutions of linear Diophantine systems : bounds and algorithms. In R. Book, editor, *Proc. Rewriting Techniques and Applications*, volume 488 of *LNCS*, pages 162–173, Berlin, 1991. Springer.
15 I. Pratt-Hartmann. Complexity of the two-variable fragment with counting quantifiers. *Journal of Logic, Language and Information*, 14(3):369–395, 2005.
16 I. Pratt-Hartmann. The finite satisfiability problem for two-variable, first-order logic with one transitive relation is decidable. *Mathematical Logic Quarterly*, 64(3):218–248, 2018.
17 I. Pratt-Hartmann, W. Szwast, and L. Tendera. The fluted fragment revisited. *Journal of Symbolic Logic*, 84(3):1020–1048, 2019.
18 Ian Pratt-Hartmann and Lidia Tendera. The fluted fragment with transitive relations, 2020. `arXiv:2006.11169`.
19 W. Purdy. Fluted formulas and the limits of decidability. *Journal of Symbolic Logic*, 61(2):608–620, 1996.
20 W. Purdy. Complexity and nicety of fluted logic. *Studia Logica*, 71:177–198, 2002.
21 W. V. Quine. Variables explained away. *Proceedings of the American Philosophical Society*, 104(3):343–347, 1960.
22 W. V. Quine. On the limits of decision. In *Proceedings of the 14th International Congress of Philosophy*, volume III, pages 57–62. University of Vienna, 1969.
23 S. Rudolph and M. Šimkus. The triguarded fragment of first-order logic. In *Proceedings of 22nd International Conference for Programming, Artificial Intelligence and Reasoning (LPAR-22)*, volume 57, pages 604–619. EPiC Series in Computing, 2018.
24 S. Schmitz. Complexity hierarchies beyond Elementary. *ACM Transactions on Computation Theory*, 8(1:3):1–36, 2016.
25 D. Scott. A decision method for validity of sentences in two variables. *Journal Symbolic Logic*, 27:477, 1962.