Adding Transitivity and Counting to the Fluted Fragment

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

Lidia Tendera
Institute of Computer Science, University of Opole, Poland

Abstract

We study the impact of adding both counting quantifiers and a single transitive relation to the fluted fragment – a fragment of first-order logic originating in the work of W.V.O. Quine. The resulting formalism can be viewed as a multi-variable, non-guarded extension of certain systems of description logic featuring number restrictions and transitive roles, but lacking role-inverses. We establish the finite model property for our logic, and show that the satisfiability problem for its k-variable sub-fragment is in $(k+1)$-NExpTime. We also derive ExpSpace-hardness of the satisfiability problem for the two-variable, fluted fragment with one transitive relation (but without counting quantifiers), and prove that, when a second transitive relation is allowed, both the satisfiability and the finite satisfiability problems for the two-variable fluted fragment with counting quantifiers become undecidable.

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

Keywords and phrases fluted logic, transitivity, counting, satisfiability, decidability, complexity

Digital Object Identifier 10.4230/LIPIcs.CSL.2023.32

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

1 Introduction

The *fluted fragment*, or $\mathcal{FL}$, is a fragment of first-order logic in which, roughly speaking, the sequence of quantification of variables coincides with the order in which those variables appear as arguments of predicates. The *fluted fragment with counting*, or $\mathcal{FLC}$, is the extension of $\mathcal{FL}$ with the standard counting quantifiers $\exists_{[\leq M]}$, $\exists_{[\geq M]}$ and $\exists_{[=M]}$, where M is a (numeral denoting a) positive integer. The following sentence is in $\mathcal{FLC}$:

At most three lecturers introduce a professor to at least five students
$$\exists_{[\leq 3]} x_1 \Big(\text{lectr}(x_1) \wedge \exists x_2 \big(\text{prof}(x_2) \wedge \exists_{[\geq 5]} x_3 (\text{std}(x_3) \wedge \text{intro}(x_1, x_2, x_3))\big)\Big) \tag{1}$$

It was shown in [31] that $\mathcal{FL}$ has the finite model property, whence its satisfiability problem is decidable. The complexity bound of NExpTime claimed in [32] is incorrect, and as shown later in [28], the problem is non-elementary; however, the satisfiability problem for the k-variable sub-fragment of $\mathcal{FL}$ is known to be in $(k-2)$-NExpTime for $k \geq 3$ [29]. This result extends to the fragment $\mathcal{FLC}$, though with a best-known upper complexity bound of $(k-1)$-NExpTime for k in the same range [27].

It is impossible, within the fluted fragment, to express the property of transitivity: in particular, the formula $\forall x_1 \forall x_2 \big(r(x_1, x_2) \rightarrow \forall x_3 \big(r(x_2, x_3) \rightarrow r(x_1, x_3)\big)\big)$ is not fluted, because the variable sequence in the atom $r(x_1, x_3)$ omits x_2. The question therefore arises as to whether the fragments $\mathcal{FL}$ or even $\mathcal{FLC}$ can be extended by declaring certain distinguished binary predicates to be interpreted as transitive relations. For $\mathcal{FL}$, this question was largely resolved in [30]. It was shown that $\mathcal{FL}$ with equality and one distinguished transitive

31st EACSL Annual Conference on Computer Science Logic (CSL 2023).
Editors: Bartek Klin and Elaine Pimentel; Article No. 32; pp. 32:1–32:22
Leibniz International Proceedings in Informatics
Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

relation lacks the finite model property, but has decidable satisfiability and finite satisfiability problems; the former problem, restricted to the k-variable sub-fragment is shown to be in k-NExpTime, and the latter in $(k+1)$-NExpTime. In the presence of two transitive relations but without equality, the fluted fragment loses the finite model property, with the decidability of satisfiability and finite satisfiability both unknown. With either two transitive relations and equality or three transitive relations, satisfiability and finite satisfiability are both undecidable. In the present paper, we consider the combination of transitivity and counting. We show that, in the absence of equality, we may add a single transitive relation to the fragment $\mathcal{FLC}$ without losing the finite model property, and we establish an upper complexity bound of $(k+1)$-NExpTime on the satisfiability problem for the k-variable sub-fragment. In the presence of two transitive relations, however, the satisfiability and finite satisfiability problems for $\mathcal{FLC}$ are undecidable, even in the absence of equality.

The impact of counting quantifiers and transitivity on decidability and complexity of satisfiability has been widely studied in the context of other fragments of first-order logic. Of particular interest in this regard are *propositional modal logics* and *description logics*, as these are mapped to the fluted fragment via the standard translation. The simplest of these logics is the propositional modal logic K, in which the modalities $\Box$ and $\Diamond$ are interpreted with respect to a single accessibility relation. By allowing counting modalities of the form $\Diamond_{\leq n}$, $\Diamond_{=n}$ and $\Diamond_{\geq n}$, we obtain propositional graded modal logic GrK; and by imposing on these systems the requirement that the accessibility relation be interpreted as a *transitive* relation, we obtain, respectively, the logics known as K4 and GrK4. Thus, GrK4 is a sub-fragment of $\mathcal{FLC}$ with one transitive relation. The satisfiability problems for K, GrK and K4 are all PSpace-complete [25, 36]; the corresponding problem for GrK4 is NExpTime-complete [17]. It is instructive to consider the further extension of these logics with the *converse* modalities $\boxminus$ and $\diamondminus$. Denoting the extension of K with converse modalities by K^c, and similarly for the other fragments, we find that the satisfiability problems for K^c, GrK^c and $\mathrm{K4}^c$ all remain in PSpace; but the corresponding problem for $\mathrm{GrK4}^c$, is undecidable (see [2, 37]). Thus, transitivity and counting can be combined in the context of modal logics; but decidability of satisfiability is lost when converse modalities are added. Yet converse modalities are precisely what we cannot express in fluted logic. It is therefore natural to ask whether it is not in fact fluting that is the critical factor here.

The basic description logic $\mathcal{ALC}$ is a notational variant of propositional multi-modal logic. Extensions of $\mathcal{ALC}$ are defined by allowing additional constructs, in particular: *number restrictions* (denoted $\mathcal{Q}$) corresponding to counting quantifiers as defined in this paper, *transitivity* of roles (denoted $\mathcal{S}$), role hierarchies ($\mathcal{H}$) corresponding to inclusions of binary relations, *nominals* ($\mathcal{O}$) and role "inverses" ($\mathcal{I}$).[1] The logic $\mathcal{SHOQ}$ constitutes a maximal description logic that can be embedded into $\mathcal{FLC}$ with transitive relations. Now it is known [15] that an unrestricted combination of number restrictions and transitivity leads to undecidability of concept satisfiability even in the smaller logic $\mathcal{SHQ}$. Indeed, it was shown in [18] that just three roles (two of them transitive) are sufficient for undecidability. In response to these negative results, description logics standardly impose the syntactic restriction that transitive roles or their super-roles do not appear in number restrictions. With these restrictions, decidability is restored: concept satisfiability for $\mathcal{SHOQ}$ is ExpTime-complete, and for $\mathcal{SHOIQ}$ it is NExpTime-complete [35]. On the other hand, there is no problem with allowing transitive relations to appear under number restrictions in description logics too weak to allow these roles to interact with each other. Thus, for example, concept

[1] In description logic, the term "inverse" has become established instead of the more normal "converse".

satisfiability for $\mathcal{SOQ}$ is decidable [16]: hardness is inherited from GrK4 [17] and the optimal upper bound has been recently shown in [13]. In the logic considered in the present paper, only one transitive relation is available, but it is allowed to combine freely with other relations. We show that this comparative freedom is, from the point of view of decidability of satisfiability, unproblematic as long as we remain within the confines of fluted logic – in effect, provided we have no access to role inverses.

In knowledge representation and database theory an important reasoning problem is the query entailment problem over incomplete databases enriched by ontologies described in some logic $\mathcal{L}$. This problem reduces to unsatisfiability of the conjunction $\varphi \wedge \neg q$, where φ describes the ontology and q is a Boolean query. So, (finite) query entailment under $\mathcal{L}$ is computationally no easier than the corresponding (finite) unsatisfiability problem. Decidability of query entailment for various fragments of restricted $\mathcal{SHOIQ}$ has already been shown [9, 10, 5, 11]. Finite entailment for $\mathcal{SHOIQ}$ was shown undecidable in [33], leaving decidability of the general entailment problem open. Since in real-life ontologies it is natural to find properties where transitivity interacts with counting (see e.g. [14]) the search for logics with decidable (finite) query entailment allowing such an interaction has a strong practical motivation. This quest has brought some positive results: decidability in 2-ExpTime of entailment of *regular path queries* for $\mathcal{SQ}$ [14], a corresponding lower bound from [12] as well as decidability of both general and finite entailment for $\mathcal{SQO}$ and for a restricted fragment of $\mathcal{SIQ}$, where inverse roles are not used under number restrictions. In the fluted fragment query entailment is decidable provided some form of guardedness is added [1].

The quest for decidable logics with counting and/or transitivity is not, of course, limited to description or modal logics. We mention in this context the *guarded fragment*, the *unary negation fragment* and the recently introduced *triguarded fragment*, each of them being as a different generalization of propositional modal logic (not subsuming fluted logic), enjoying the finite model property, and not being able to express transitivity or related properties. Considerable work has been done to understand the limits of decidability for extensions of these fragments where (restricted) transitivity is allowed. The general picture that emerges is that, when considering more than one transitive relation, the interaction between them must be restricted in order to secure decidability of satisfiability. For instance, satisfiability is undecidable for the the two-variable guarded fragment with two transitive relations [19], but is restored even in the presence of any number of transitive relations if the transitive relations are allowed to appear only in guard positions [34]. Similarly, satisfiability is undecidable for the two-variable guarded fragment over signatures with two *linear* orders, but is restored if we restrict the linear orders to guard positions and in addition forbid any other binary relations in the signature (a property analogous to the restriction of number restrictions to simple roles in description logics) [21]. Technically, the crucial property behind the decision procedures for the satisfiability problem in these cases is that satisfiable formulas exhibit some kind of tree-like models in which the transitive relations are defined independently of each other. For a wider picture we refer the reader to the recent papers on extensions of: the guarded fragment [24], the unary negation fragment [7, 8], the triguarded fragment [22, 23] and the references therein.

The rest of the paper is organised as follows. After defining the main notions in Section 2 we begin in Section 3 by showing the ExpSpace-hardness of the satisfiability problem for the two-variable fluted logic with one transitive relation. In Section 4 we discuss *Presburger quantifiers*, which are a generalization of counting quantifiers. In Section 5 we show that the two-variable fragment of $\mathcal{FLC}$ with one transitive relation has the finite model property when restricted to signatures consisting of unary predicates and one distinguished transitive relation. In Section 6 we first lift this result to arbitrary signatures, and then drop the restriction to two-variables. In Section 7 we show that the satisfiability and finite satisfiability problems for $\mathcal{FLC}$ with two transitive relations are undecidable even in the two-variable case.

2 Preliminaries

In the sequel, we use Fraktur letters for structures and the corresponding Roman letters for their domains: thus, A is the domain of $\mathfrak{A}$ *etc.* 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. We begin by establishing the syntax of the fragment $\mathcal{FLC}$, *the fluted fragment with counting.* Define the sets of formulas $\mathcal{FLC}^{[k]}$, for all $k \geq 0$, by simultaneous structural recursion as follows:

(i) any atom $p(x_\ell, \ldots, x_k)$, where $x_\ell, \ldots, x_k$ is a contiguous subsequence of $\bar{x}_\omega$ and p a (non-equality) predicate of arity $k - \ell + 1$, is in $\mathcal{FLC}^{[k]}$;

(ii) $\mathcal{FLC}^{[k]}$ is closed under boolean combinations;

(iii) if φ is in $\mathcal{FLC}^{[k+1]}$, then $\exists x_{k+1}\, \varphi$ and $\forall x_{k+1}\, \varphi$ are in $\mathcal{FLC}^{[k]}$,

(iv) if φ is in $\mathcal{FLC}^{[k+1]}$ and M a non-negative integer, then $\exists_{[\leq M]} x_{k+1}\, \varphi$, $\exists_{[\geq M]} x_{k+1}\, \varphi$ and $\exists_{[=M]} x_{k+1}\, \varphi$ are in $\mathcal{FLC}^{[k]}$.

It is intended that Clause (i) allows the case $\ell = k+1$ (empty sequence of arguments), so that the atoms in question are proposition letters. Clause (iv) speaks of *non-negative integers* occurring in formulas; when calculating the size $\|\varphi\|$ of an $\mathcal{FLC}^k$-formula φ, we assume these to be represented as bit-strings (i.e. binary coding of counting subscripts). Define $\mathcal{FL}^{[k]}$ to be the fragment of $\mathcal{FLC}^{[k]}$ in which no counting quantifiers occur, i.e. Clause (iv) is dropped. Define the fragment $\mathcal{FLC}$ to be the union $\bigcup_{k\geq 0} \mathcal{FLC}^{[k]}$; similarly $\mathcal{FL} = \bigcup_{k\geq 0} \mathcal{FL}^{[k]}$. By $\mathcal{FLC}$+1Tr, we understand the same set of formulas as $\mathcal{FLC}$, but with a distinguished binary predicate $\mathfrak{t}$, which is required to be interpreted as a transitive relation; similarly for $\mathcal{FL}+1$Tr. Finally, define $\mathcal{FLC}^k$ to be the fragment of $\mathcal{FLC}$ in which at most the variables $x_1, \ldots, x_k$ appear (free or bound); and similarly for $\mathcal{FL}^k$, $\mathcal{FLC}^k$+1Tr and $\mathcal{FL}^k + 1$Tr. Incidentally, do not confuse $\mathcal{FLC}^k$ with $\mathcal{FLC}^{[k]}$: for example, formula (1) is in $\mathcal{FLC}^k$ for all $k \geq 3$ but in $\mathcal{FLC}^{[k]}$ only for $k = 0$. By *sentence*, we mean a formula with no free variables – that is, up to a shift of variable indices, a formula of $\mathcal{FLC}^{[0]}$.

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. For example, formula (1) can be unambiguously reconstructed – up to a shift of variable indices – from

$$\exists_{[\leq 3]}\Big(\text{lectr} \wedge \exists\big(\text{prof} \wedge \exists_{[\geq 5]}(\text{std} \wedge \text{intro})\big)\Big). \tag{2}$$

Consequently, we employ this variable-free notation for $\mathcal{FLC}$ in the sequel, as it is more compact. The issue of ambiguity with respect to shifts in variable indexing bears emphasis, however. Under variable-free notation, any formula of $\mathcal{FLC}^{[k]}$ is, *without lexical change*, also a formula of $\mathcal{FLC}^{[k+1]}$. For example, the sub-formula $\exists(\text{prof} \wedge \exists_{[\geq 5]}(\text{std} \wedge \text{intro}))$ of (2) may be reconstructed not only as the $\mathcal{FLC}^{[1]}$-formula $\varphi(x_1)$ given by $\exists x_2(\text{prof}(x_2) \wedge \exists_{[\geq 5]} x_3(\text{std}(x_3) \wedge \text{intro}(x_1, x_2, x_3)))$, but also as the (logically equivalent) $\mathcal{FLC}^{[2]}$-formula $\varphi'(x_1, x_2)$ given by $\exists x_3(\text{prof}(x_3) \wedge \exists_{[\geq 5]} x_4(\text{std}(x_4) \wedge \text{intro}(x_2, x_3, x_4)))$, and so on. This feature will actually be useful in Sec. 6, when we come to reduce the number of variables in $\mathcal{FLC}$+1Tr-formulas.

Using variable-free notation, we say that an $\mathcal{FLC}^k$-formula ($k \geq 2$) is in *normal-form* if it conforms to the pattern

$$\bigwedge_{h=1}^{m} \forall^{k-1}(\alpha_h \rightarrow \exists_{[\lesseqgtr_h M_h]} \beta_h), \tag{3}$$

where the α_h are quantifier-free $\mathcal{FLC}^{k-1}$-formulas, the β_h are quantifier-free $\mathcal{FLC}^k$-formulas, the M_h are non-negative integers and the symbols $\lesseqgtr_h$ are chosen from the set $\{\leq, \geq, =\}$.

▶ **Lemma 1.** *Let φ be a sentence of $\mathcal{FLC}^k$ ($k \geq 2$). Then we may compute, in time bounded by a polynomial function of $\|\varphi\|$, a normal-form $\mathcal{FLC}^k$-sentence ψ such that:* (i) $\models \psi \rightarrow \varphi$*; and* (ii) *any model of φ can be expanded to a model of ψ.*

Proof. The proof is essentially the same as for $\mathcal{FL}$, via standard re-writing techniques. See, e.g. [29, Lemma 4.1]. Note that a formula of the form $\forall^k\theta$, with $\theta \in \mathcal{FLC}^k$ quantifier-free, is logically equivalent to the normal-form conjunct $\forall^{k-1}(\top \rightarrow \exists_{[=0]}\neg\theta)$. ◀

If $\mathcal{L}$ is any language, we denote the satisfiability problem for $\mathcal{L}$ by Sat($\mathcal{L}$). Since all of the problems Sat($\mathcal{L}$) with which we shall be concerned have relatively high complexity (well above NPTime), we may assume henceforth that proposition letters do not occur in formulas, as their truth-values can simply be guessed.

In this paper, we shall make some limited use of Presburger arithmetic, the first-order theory of the structure $\mathfrak{N} = (\mathbb{N}, <, 0, 1, +)$, where $\mathbb{N} = \{0, 1, 2 \ldots\}$ is the domain of natural numbers, and the symbols $<$, 0, 1, + have their usual interpretations. A *Presburger formula* is any formula Θ in this language. We call Θ *existential* if it is of the form $\exists\bar{v}\,\Xi$, where $\bar{v}$ is a (possibly empty) tuple of variables and Ξ is quantifier-free. If $\bar{a} = a_1, \ldots, a_k$ is a tuple of numbers and M a number, we write $\bar{a} \leq M$, to mean $a_i \leq M$ for all i ($1 \leq i \leq k$). If Θ is a Presburger formula featuring only the variables $\bar{w}$, where $\bar{w}$ and $\bar{a}$ have the same arity, we say $\bar{a}$ *satisfies* $\Theta(\bar{w})$ if $\mathfrak{N} \models \Theta[\bar{a}]$. We allow constants for all natural numbers in Presburger formulas, since these can be succinctly defined using 0, 1 and +. We also allow ourselves to use *relativized* quantification $\exists(\bar{v} \leq M)$, again defined in the obvious way.

3 ExpSpace-hardness

In this section, we show that Sat($\mathcal{FL}^2 + 1\text{Tr}$) is ExpSpace-hard. This improves the already known NExpTime-lower bounds from GrK4 [17] and $\mathcal{FL}^2$ [29]. Note that $\mathcal{FL}^2 + 1\text{Tr}$, the two-variable fluted fragment with one transitive relation, features neither equality nor counting quantifiers. The material also serves as an opportunity for familiarization with the variable-free syntax for fluted logic just introduced. By the well-known correspondence between deterministic and alternating complexity classes [6], it suffices to show a polynomial time reduction from the acceptance problem for an alternating Turing machine working in exponential time. The reduction can be effected using similar ideas as in, e.g. [20], where the complexity of satisfiability for the two-variable guarded fragment with one-way transitive guards was studied. However, special care needs to be taken to encode the required properties using only fluted formulas. In particular, it is essential for us to consider signatures featuring several binary predicates, in contrast to the non-fluted case [20], where ExpSpace-hardness is shown even when the distinguished transitive relation is the only binary predicate.

We employ a natural representation of integers as bit strings. Let $\bar{s} = s_{z-1}, \ldots, s_0$ be a bit-string representing an integer $n = \sum_{i=0}^{z-1} s_i \cdot 2^i$ (s_0 is the least significant bit). Within a structure $\mathfrak{A}$ interpreting unary predicates $p_0, \ldots, p_{z-1}$, each element b can be associated with the integer value between 0 and $2^z - 1$ represented by the bit-string $s_{z-1}, \ldots, s_0$, where, for all i ($0 \leq i < z$), $s_i = 1$ if $\mathfrak{A} \models p_i[b]$, and 0 otherwise.

Turning to the actual reduction, let M be an alternating Turing machine working in time 2^{n^k}, for some $k \geq 1$. Let Σ be the tape alphabet of M and Q the set of states; we write $Q = Q_\exists \cup Q_\forall$, where $Q_\exists$ ($Q_\forall$) is the set of existential (universal) states. Without loss of generality, we assume that for each combination of a state and alphabet symbol M has exactly two transitions, which we think of as leading to "left" and "right" successor-configurations. This notion refers only to the organization of the computation tree of M and not to the directions of the head defined by the transitions of M. We also assume that M never moves left from the initial tape cell and that it accepts or rejects in exactly the 2^{n^k}-th step. The idea of the reduction is to write a formula Φ such that each of its models encodes a binary tree whose nodes correspond to configurations of M on input w: the root of the tree encodes the initial configuration, and successors of a node encode successor configurations corresponding to transitions on M. An extra unary predicate lft is used to distinguish the "left" successor-configuration of a node. A configuration is encoded by 2^{n^k} elements, each of them corresponding to a single cell of the tape of M.

Fix some input w for M of length n, and, letting $z = n^k$, take two sets of unary predicates $p_0, \ldots, p_{z-1}$, $c_0, \ldots, c_{z-1}$ that will encode, for each element b in any structure $\mathfrak{A}$ interpreting them, two integer values $\mathrm{val}_P^{\mathfrak{A}}(b)$ and $\mathrm{val}_C^{\mathfrak{A}}(b)$ in the range $[0, 2^z - 1]$, as described above. Thinking of b as a tape cell in some node of the computation tree of M on input w, we are invited to read $\mathrm{val}_P(b)$ as the *position* of the tape cell in question and $\mathrm{val}_C(b)$ as the *time step* of the node in question (i.e. *depth* in the computation tree). For each $D \in \{C, P\}$, we add to the signature the unary predicates zero_D and max_D and the binary predicates pred_D, succ_D and eq_D. We then write a satisfiable formula Φ_{count} with the property that any model $\mathfrak{A} \models \Phi_{count}$ satisfies the following conditions:

(c1) $\mathfrak{A} \models \mathrm{zero}_D[b] \Leftrightarrow (\mathrm{val}_D^{\mathfrak{A}}(b) = 0)$, and $\mathfrak{A} \models \mathrm{max}_D[b] \Leftrightarrow (\mathrm{val}_D^{\mathfrak{A}}(b) = 2^z - 1)$,
(c2) $\mathfrak{A} \models \mathrm{eq}_D[b, b'] \Leftrightarrow \mathrm{val}_D^{\mathfrak{A}}(b) = \mathrm{val}_D^{\mathfrak{A}}(b')$,
(c3) $\mathfrak{A} \models \mathrm{pred}_D[b, b'] \Leftrightarrow \mathrm{val}_D^{\mathfrak{A}}(b') = \mathrm{val}_D^{\mathfrak{A}}(b) - 1$ modulo 2^z,
(c4) $\mathfrak{A} \models \mathrm{succ}_D[b, b'] \Leftrightarrow \mathrm{val}_D^{\mathfrak{A}}(b') = \mathrm{val}_D^{\mathfrak{A}}(b) + 1$ modulo 2^z.

Thus, we may informally read $\mathrm{zero}_C(x)$ as "x corresponds to some tape cell in the *root* node (i.e. time=0); $\mathrm{succ}_P(x, y)$ as "x corresponds to a tape cell in some node, and y to the tape cell immediately to its *right* (in some possibly different node); and so on. The construction of Φ_{count} is routine; for details, see [29, pp. 1026-27]).

The predicate $\mathfrak{t}$ is used to encode the structure of the computation tree. Intuitively, we take $\mathfrak{t}[b, b']$ to mean that b and b' are tape cells in the same node with b lying to the left of b', or that b and b' are tape cells in different nodes, with the latter being a proper descendant of the former in the computation tree. We denote by Φ_{tree} the conjunction $\Phi_1 \wedge \Phi_2 \wedge \Phi_3$ establishing that any model of Φ_{tree} includes a tree substructure, where

$$\Phi_1 \equiv \exists(\mathrm{zero}_C \wedge \mathrm{zero}_P \wedge \mathrm{lft})$$
$$\Phi_2 \equiv \forall\big(\neg\mathrm{max}_P \rightarrow (\mathrm{lft} \rightarrow \exists(\mathfrak{t} \wedge \mathrm{eq}_C \wedge \mathrm{succ}_P \wedge \mathrm{lft})\big) \wedge \big(\neg\mathrm{lft} \rightarrow \exists(\mathfrak{t} \wedge \mathrm{eq}_C \wedge \mathrm{succ}_P \wedge \neg\mathrm{lft})\big)$$
$$\Phi_3 \equiv \forall\big(\mathrm{max}_P \wedge \neg\mathrm{max}_C \rightarrow \exists(\mathfrak{t} \wedge \mathrm{lft} \wedge \mathrm{zero}_P \wedge \mathrm{succ}_C) \wedge \exists(\mathfrak{t} \wedge \neg\mathrm{lft} \wedge \mathrm{zero}_P \wedge \mathrm{succ}_C)\big).$$

To help the reader further with the variable-free notation we give the formula Φ_3 in standard first-order syntax:

$$\forall x_1\big(\mathrm{max}_P(x_1) \wedge \neg\mathrm{max}_C(x_1) \rightarrow \exists x_2(\mathfrak{t}(x_1, x_2) \wedge \mathrm{lft}(x_2) \wedge \mathrm{zero}_P(x_2) \wedge \mathrm{succ}_C(x_1, x_2)) \wedge \exists x_2(\mathfrak{t}(x_1, x_2) \wedge \neg\mathrm{lft}(x_2) \wedge \mathrm{zero}_P(x_2) \wedge \mathrm{succ}_C(x_1, x_2))\big).$$

The conjunct Φ_1 specifies that each model contains an element belonging to the root of the tree, for which both counters are set to 0. This element is also marked as satisfying lft and it will correspond to the first tape cell of the initial configuration of M on w. The conjunct Φ_2 provides remaining elements belonging to the same node. They will correspond to successive tape cells within the configuration encoded in this node: these elements form a chain connected by $\mathfrak{t}$ with the P-counter increasing from 0 to $2^z - 1$, and the C-counter stable. The predicate lft is uniformly true or false for all elements in such a chain, depending on whether the first element (with P-counter 0) is marked lft. The conjunct Φ_3 provides elements that belong to successors of a given node in the tree; they will correspond to the first tape cell in the configurations encoded by the successor nodes. Specifically, each element with $\mathrm{val}_P = 2^z - 1$ and $\mathrm{val}_C < 2^z - 1$ has two witnesses connected by $\mathfrak{t}$, each with the C-counter incremented by 1 and the P-counter set to 0: one satisfying lft, the other one $\neg$lft. Note that by transitivity of $\mathfrak{t}$ each element belonging to a given node is connected by $\mathfrak{t}$ to all elements in descendant nodes.

To encode a configuration of M, we employ further unary predicates: h, $a_0, \ldots, a_s$, $q_0, \ldots, q_r$. Here, h indicates the position of the head of M, $a_0, \ldots, a_s$ correspond to the symbols of the alphabet of M, and $q_0, \ldots, q_r$ correspond to the states of M. We assume that a_0 represents the blank symbol, a_s a special start-of-tape symbol, q_0 the initial state and q_r the only rejecting state. Note that we use the same letters for the alphabet symbols of M and the predicates representing them in the signature of Φ; similarly for states.

Let Φ_M be a conjunction of formulas describing the following basic properties of an accepting computation tree of M:

(m1) every tape cell contains exactly one symbol,
(m2) in each configuration the head is scanning at most one cell: $\forall(h \to \forall(\mathfrak{t} \wedge \mathrm{eq}_C \to \neg h))$,
(m3) the information about the current state is stored in elements scanned by the head,
(m4) the root of the computation tree contains the initial configuration of M on input w,
(m5) when moving from a configuration to its successor configurations only tape cells scanned by the head are allowed to change: $\bigwedge_{0 \le i \le s} \forall(a_i \wedge \neg h \to \forall(\mathfrak{t} \wedge \mathrm{eq}_P \wedge \mathrm{succ}_C \to a_i))$.

Properties (m1), (m3) and (m4) are one-variable formulas, so they clearly belong to $\mathcal{FL}^2$.

To encode the transition function of M we consider separately transitions from existential and from universal states. Suppose M has the following possible transitions for an existential state q and letter a: $\delta(q,a) = \{(q', a', R), (q'', a'', L)\}$. We want to ensure that the left child of nodes corresponding to configurations of M in state q reading a encodes one of the possible next configurations. This can be formalized by the formula $\Phi^{\exists}_{q,a}$ below:

$$\Phi^{\exists}_{q,a} \equiv \forall\Big(q \wedge a \wedge h \to \big(\forall(\mathfrak{t} \wedge \mathrm{lft} \wedge \mathrm{succ}_C \to (\mathrm{eq}_P \to a') \wedge (\mathrm{succ}_P \to q' \wedge h)) \quad \vee \\ \forall(\mathfrak{t} \wedge \mathrm{lft} \wedge \mathrm{succ}_C \to (\mathrm{pred}_P \to q'' \wedge h) \wedge (\mathrm{eq}_P \to a''))\big)\Big).$$

When encoding the computation at universal states we require that one transition is encoded by the left child and the other one by the right child. In particular, when M has the transitions $\delta(q,a) = \{(q', a', R), (q'', a'', L)\}$ and $q \in Q_\forall$, we define the formula $\Phi^{\forall}_{q,a}$ as follows:

$$\Phi^{\forall}_{q,a} \equiv \forall\Big(q \wedge a \wedge h \to \forall\big(\mathfrak{t} \wedge \mathrm{lft} \wedge \mathrm{succ}_C \to (\mathrm{eq}_P \to a') \wedge (\mathrm{succ}_P \to q' \wedge h)\big) \quad \wedge \\ \forall\big(\mathfrak{t} \wedge \neg\mathrm{lft} \wedge \mathrm{succ}_C \to (\mathrm{pred}_P \to q'' \wedge h) \wedge (\mathrm{eq}_P \to a'')\big)\Big).$$

Finally, let $\Phi \equiv \Phi_{count} \wedge \Phi_{tree} \wedge \Phi_M \wedge \bigwedge_{a \in \Sigma}(\bigwedge_{q \in Q_\exists} \Phi^{\exists}_{q,a} \wedge \bigwedge_{q \in Q_\forall} \Phi^{\forall}_{q,a}) \wedge \forall \neg q_r$. The role of the last conjunct of Φ is to ensure that M never enters a rejecting state. The length of Φ is polynomial in the length of w, regarding M as fixed.

It is not difficult to show that Φ is satisfiable if and only if M accepts w. For the only-if direction, suppose that there is an accepting computation tree of M with input w; we construct a model $\mathfrak{A}$ of Φ in the form of a tree. The initial configuration of M is transformed into the root of $\mathfrak{A}$ as suggested by part (m4) of Φ_M. Then we proceed recursively. When a configuration $\mathcal{C}$ in the computation tree of M corresponding to a node b in $\mathfrak{A}$ is universal, we transform the left subtree of $\mathcal{C}$ into the left subtree of b in $\mathfrak{A}$, and the right subtree of $\mathcal{C}$ into the right subtree of b. If $\mathcal{C}$ is existential we transform the accepting subtree of $\mathcal{C}$ into both the left and the right subtree of b. Since M is accepting, the formula Φ is true in the model.

For the opposite direction, suppose $\mathfrak{A} \models \Phi$. We construct an accepting computation tree for M on input w, starting with the initial configuration, and then proceed recursively. Suppose we have constructed a configuration $\mathcal{C}$ of M on w of depth $d < 2^z - 1$. The formula Φ_{tree} ensures that $\mathfrak{A}$ contains two chains of elements $\bar{b} = b_0, \ldots, b_{2^z-1}$ and $\bar{b}' = b'_0, \ldots, b'_{2^z-1}$ (i.e. for all i $(0 \leq i < 2^z - 1)$ $\mathfrak{A} \models \mathfrak{t}[b_i, b_{i+1}] \wedge \mathfrak{t}[b'_i, b'_{i+1}]$) connected to the elements representing $\mathcal{C}$ by $\mathfrak{t}$ such that for all i $(0 \leq i \leq 2^z - 1)$ we have: $\mathfrak{A} \models \mathrm{val}_P[b_i] = \mathrm{val}_P[b'_i] = i$, $\mathfrak{A} \models \mathrm{val}_C[b_i] = \mathrm{val}_C[b'_i] = d + 1$, $\mathfrak{A} \models \mathrm{lft}[b_i] \wedge \neg\mathrm{lft}[b'_i]$. If $\mathcal{C}$ is existential we translate the information encoded by the chain $\bar{b}$ of elements $b_0, \ldots, b_{2^z-1}$ into a successor configuration of $\mathcal{C}$. If $\mathcal{C}$ is universal we translate the information encoded by the chain $\bar{b}$ into the left successor of $\mathcal{C}$, and the information encoded by the chain $\bar{b}'$ into the right successor of $\mathcal{C}$. The conjuncts $\Phi^{\exists}_{q,a}$ and $\Phi^{\forall}_{q,a}$ ensure that the tree constructed in such a way is a computation tree of M, and by the conjunct $\forall\neg q_r$ of Φ, the tree is accepting.

Hence, we have

▶ **Theorem 2.** $\mathrm{Sat}(\mathcal{FL}^2 + 1\mathrm{Tr})$ *is* ExpSpace*-hard.*

4 Presburger quantifiers

In Sec. 3, we obtained a lower complexity bound on $\mathrm{Sat}(\mathcal{FLC}^2{+}1\mathrm{Tr})$ by investigating the corresponding problem for the *sub*-fragment of that logic without counting quantifiers. In Secs. 5–6, we will obtain upper complexity bounds on $\mathrm{Sat}(\mathcal{FLC}^k{+}1\mathrm{Tr})$ $(k \geq 2)$ by investigating the corresponding problem for a *super*-fragment of that logic in which the notion of counting quantifier has been generalized. These generalized counting quantifiers were introduced in [27] in the context of $\mathcal{FLC}$ (but see [3] for a closely related idea); this section provides a more streamlined presentation.

Fix $k \geq 1$. A *fluted k-atom* over Σ is a predicate $p \in \Sigma$ of arity $d \leq k$. A *fluted k-literal* over Σ is a fluted k-atom over Σ or its negation. A *fluted k-type* over Σ is a maximal consistent set of fluted literals over Σ. It is worth explaining how variable-free notation is to be understood here: we think of a literal $\pm p$ in a fluted k-type as "really" being the formula $\pm p(x_{k-d+1}, \ldots, x_k)$, so that variables are aligned in such a way that all literals notionally have the same last variable. We denote the set of fluted k-types over Σ by Ftp^{Σ}_k. If $M \geq 0$, an *M-bounded, fluted k-profile* over Σ is a function $\zeta : \mathrm{Ftp}^{\Sigma}_{k+1} \to \{0, \ldots, M\}$. An *$M$-bounded, fluted k-star-type* over Σ is a pair $\langle \pi, \zeta \rangle$, where π is a fluted k-type over Σ and ζ an M-bounded, fluted k-profile over Σ. In all cases, reference to Σ is suppressed when clear from context. (The intuition behind these definitions will be explained presently.) We use π to range over fluted 1-types, ρ, σ, τ to range over fluted k-types (for various k), and ζ, η to range over M-bounded, fluted k-profiles (for various M and k). To reduce notational clutter, and assuming that Σ is finite, we silently identify fluted k-types with their conjunctions where convenient, thus allowing ourselves to write, for instance, τ for the quantifier-free $\mathcal{FLC}^k$-formula $\bigwedge \tau$.

Let $\mathfrak{A}$ be a structure interpreting a finite signature Σ and $\bar{a}$ a k-tuple of elements of A. The set of fluted k-literals over Σ satisfied by $\bar{a}$ in $\mathfrak{A}$ is evidently a fluted k-type, τ. We say that $\bar{a}$ *realizes* τ, and refer to it as *the fluted k-type of $\bar{a}$ in* $\mathfrak{A}$, denoted $\mathrm{ftp}^{\mathfrak{A}}[\bar{a}]$. Intuitively, a fluted k-type is to be thought of as a specification, for some k-tuple of elements, of which fluted literals are satisfied by that k-tuple. Now consider the $(k+1)$-tuples $\bar{a}b$ extending $\bar{a}$ by a single element b. Each such $(k+1)$-tuple has some fluted $(k+1)$-type in $\mathrm{Ftp}^{\Sigma}_{k+1}$, and we may *count*, up to some bound M, the number of times each $(k+1)$-type in $\mathrm{Ftp}^{\Sigma}_{k+1}$ arises in this way. More precisely, the *M-bounded, fluted profile of* $\bar{a}$ in $\mathfrak{A}$ is the function $\mathrm{fpr}^{\mathfrak{A}}_{M}[\bar{a}] : \mathrm{Ftp}^{\Sigma}_{k+1} \to \{0, \ldots, M\}$ defined by

$$\mathrm{fpr}^{\mathfrak{A}}_{M}[\bar{a}](\tau) = \min(|\{b \in A : \mathrm{ftp}^{\mathfrak{A}}[\bar{a}b] = \tau\}|, M).$$

Clearly, $\mathrm{fpr}^{\mathfrak{A}}_{M}[\bar{a}]$ is an M-bounded, fluted k-profile, as defined in the preceding paragraph. Intuitively, an M-bounded, fluted k-profile is to be thought of as a specification, for some k-tuple of elements, of how many different $(k+1)$-tuples it can be extended to (up to a ceiling of M) satisfying each of the possible fluted $(k+1)$-types. Finally, we say that the *M-bounded, fluted star-type of* $\bar{a}$ in $\mathfrak{A}$ is the pair $\mathrm{fst}^{\mathfrak{A}}_{M}[\bar{a}] = \langle \mathrm{ftp}^{\mathfrak{A}}[\bar{a}], \mathrm{fpr}^{\mathfrak{A}}_{M}[\bar{a}] \rangle$. Intuitively, an M-bounded, fluted k-star-type is simply a combination of a fluted k-type and an M-bounded, fluted k-profile.

We now turn to the promised generalization of counting quantifiers. By way of motivation, let β be a quantifier-free $\mathcal{FLC}^{k+1}$-formula, and consider the $\mathcal{FLC}^{k}$-formula $\exists_{\leq M}\beta$. Applied to any k-tuple of elements $\bar{a}$, this latter formula makes a statement about the fluted $(k+1)$-types satisfied by the various $(k+1)$-tuples $\bar{a}b$ as b ranges over the domain of quantification, namely, that for at most M such elements b, the fluted type of $\bar{a}b$ entails β. This formulation invites generalization. For each fluted $(k+1)$-type τ over the relevant signature, let v_{τ} be the number of elements b such that $\bar{a}b$ has fluted type τ, up to some fixed ceiling M (after which we stop counting). Then we can impose any set of conditions on the collection of integers v_{τ} (in the range $[0, M]$) thus obtained. Accordingly, we say that a *computable counting quantifier* is an expression $Q(k, \Sigma, M, \Theta)$, where $k \geq 1$, Σ is a signature, $M \geq 0$ and Θ a set of M-bounded fluted k-profiles over Σ, presented in some way which allows membership to be effectively determined. (For example, Θ could be presented as a Turing machine that terminates on all inputs.) We regard computable counting quantifiers as $\mathcal{FLC}^{k}$-formulas in their own right. If $\mathfrak{A}$ is any structure interpreting a signature including Σ, and $\bar{a}$ is a k-tuple of elements from A, then we define

$$\mathfrak{A} \models Q(k, \Sigma, M, \Theta)[\bar{a}] \quad \text{iff} \quad \mathrm{fpr}^{\mathfrak{A}}_{M}[\bar{a}] \in \Theta.$$

In particular, if β is a quantifier-free $\mathcal{FLC}^{k+1}$-formula, then the $\mathcal{FLC}^{k}$-formula $\exists_{\leq M}\beta$ can be equivalently written as $Q(k, \Sigma, M+1, \Theta)$, where Θ is the set of $(M+1)$-bounded, fluted k-profiles ζ satisfying the condition

$$\sum \{\zeta(\tau) : \tau \in \mathrm{Ftp}^{\Sigma}_{k+1},\ \models \tau \to \beta\} \leq M.$$

Thus, computable counting quantifiers generalize ordinary counting quantifiers applied to quantifier-free formulas. Conversely, computable quantifiers can be expressed in terms of counting quantifiers, using the fact that all numbers concerned are subject to the finite bound M. Thus, $Q(k, \Sigma, M, \Theta)$ can be written as the (huge) $\mathcal{FLC}^{k+1}$-formula

$$\bigvee_{\zeta \in \Theta} \left(\bigwedge \{\exists_{[=\zeta(\tau)]}\, \tau \mid \tau \in \mathrm{Ftp}^{\Sigma}_{k+1} \text{ s.t. } \zeta(\tau) < M\} \wedge \bigwedge \{\exists_{[\geq M]}\, \tau \mid \tau \in \mathrm{Ftp}^{\Sigma}_{k+1} \text{ s.t. } \zeta(\tau) = M\} \right).$$

Hence there is no difference in expressive power between ordinary counting quantifiers and computable counting quantifiers. However, there may be a complexity-theoretic difference: the condition Θ may be presented in a more compact way than ordinary counting quantifiers make possible. One such way will be of particular significance in this paper. Let $\bar{v}$ be a collection of variables v_τ (in some fixed order) as τ ranges over $\text{Ftp}^{\Sigma}_{k+1}$, and let $\Theta(\bar{v})$ be a formula of Presburger arithmetic in the variables $\bar{v}$. Now, any tuple of numbers chosen from $\{0, \ldots, M\}$ satisfying Θ corresponds to a function mapping each fluted $(k{+}1)$-type τ to the number assigned to the variable v_τ. Thus, Θ specifies a set of M-bounded, fluted k-profiles in a natural way. A computable counting quantifier in which Θ is presented as a formula of Presburger arithmetic in this way will be called a *Presburger quantifier.* If Θ is an existential Presburger formula, i.e. a formula of the form $\exists \bar{w}\, \Xi(\bar{v}, \bar{w})$, then we speak of an *existential Presburger quantifier.*

These considerations lead us to define a sequence of languages based on existential Presburger quantifiers. For any $k \geq 2$, let $\mathcal{FLU}^k$ be the collection of formulas of the form

$$\bigwedge_{h=1}^{m} \forall^{k-1}(\alpha_h \to Q(k-1, \Sigma, M, \Theta_h)) \tag{4}$$

where Σ is a purely relational signature, m a positive integer, the α_h quantifier-free $\mathcal{FLC}^{k-1}$-formulas over Σ, M a non-negative integer and the Θ_h existential Presburger formulas with free variables $\bar{v}$ indexed by the fluted k-types over Σ. The signature of such a formula is taken to be Σ. We need a careful parametrization of $\mathcal{FLU}^k$+1Tr-formulas in the sequel. If ψ is of the form (4), define the *effective size of* ψ, denoted $\#(\psi)$, to be the quantity $|\Sigma| + \log M + m$. We define the fragment $\mathcal{FLU}^k$+1Tr to consist of the set of formulas of the form (4), over a signature featuring $\mathfrak{t}$, again required to be interpreted as a transitive relation.

It is simple to re-write any normal-form $\mathcal{FLC}^k$-formula φ of the form (3) over a signature Σ as a logically equivalent formula ψ of the form (4), where $M = \max(M_1, \ldots M_m) + 1$. Moreover, ψ can be computed in exponential time. (Exponential time is required, because the variables in the embedded Presburger formulas are indexed by fluted 1-types, of which there are exponentially many.) Note however that $\#(\psi) \leq \|\varphi\|$. Thus, using Lemma 1, we see that any decision procedure for the problem $\text{Sat}(\mathcal{FLU}^k)$ yields a decision procedure for the problem $\text{Sat}(\mathcal{FLC}^k)$. To obtain such a procedure, we show how to reduce $\text{Sat}(\mathcal{FLU}^{k+1}\text{+1Tr})$ to $\text{Sat}(\mathcal{FLU}^k\text{+1Tr})$ (but with exponential blow-up), where $k \geq 2$. The principal difficulty is to then establish the base case, namely, $\text{Sat}(\mathcal{FLU}^2\text{+1Tr})$.

We approach our task in two stages. We consider first a sub-fragment of $\mathcal{FLU}^2$+1Tr obtained by restricting attention to signatures featuring no predicates of arity higher than 1 except for the distinguished predicate $\mathfrak{t}$. Sec. 5 is devoted to a small model property for this sub-fragment. Then, in Sec. 6, we first lift this result to the whole of $\mathcal{FLU}^2$+1Tr, and then generalize to $\text{Sat}(\mathcal{FLU}^k\text{+1Tr})$ for all $k \geq 2$. We remark that the same basic strategy was employed in [27] to decide $\text{Sat}(\mathcal{FLC})$. There, however, one has the luxury of an easy base case: the logic $\mathcal{FLC}^2$ is contained in $\mathcal{C}^2$, the two-variable fragment of first-order logic with counting, for which decidability of satisfiability is known. Indeed, the same argument shows the decidability of $\text{Sat}(\mathcal{FLC})$ even in the presence of a single distinguished binary predicate required to be interpreted as an *equivalence relation*, since the corresponding extension of $\mathcal{C}^2$ is known to have a decidable satisfiability problem [26]. Unfortunately, adding a single *transitive* relation to $\mathcal{C}^2$ yields a logic with undecidable satisfiability problem. Thus, the main work of this paper is to establish the decidability of $\text{Sat}(\mathcal{FLU}^2\text{+1Tr})$ from scratch.

Since satisfaction of a Presburger quantifier by a tuple $\bar{a}$ is determined by the profile of $\bar{a}$, the following observation is immediate:

▶ **Lemma 3.** *Let φ be a formula of $\mathcal{FLU}^k$ (or of $\mathcal{FLU}^k$+1Tr) of the form (4), with $k \geq 2$, over a signature Σ. Let $\mathfrak{A}$, $\mathfrak{B}$ be structures interpreting Σ such that every M-bounded, fluted $(k-1)$-star-type realized in $\mathfrak{B}$ is realized in $\mathfrak{A}$. Then $\mathfrak{A} \models \varphi$ implies $\mathfrak{B} \models \varphi$.*

5 Two variables, unary predicates

Define $\mathcal{FLU}^2_u$+1Tr to be the sub-fragment of $\mathcal{FLU}^2$+1Tr in which only unary predicates or the distinguished binary predicate $\mathfrak{t}$ occur. In this section we show that $\mathcal{FLU}^2_u$+1Tr has the finite model property, and that Sat($\mathcal{FLU}^2_u$+1Tr) is in 2-NExpTime.

We briefly outline the technique employed. Because $\mathcal{FLU}^2_u$+1Tr is not guarded, we cannot easily build tree-like models of satisfiable formulas as with modal or description logics. Because $\mathcal{FLU}^2_u$+1Tr features a transitive relation, we cannot easily employ equational techniques as with the two-variable fragment with counting, $\mathcal{C}^2$. Instead, we show how, from an arbitrary model $\mathfrak{A}$ of some $\mathcal{FLU}^2$+1Tr-formula φ, we can extract a model $\mathfrak{B}$ of cardinality bounded by some doubly exponential function in the size of φ. *A priori*, the directed graph $(A, \mathfrak{t}^{\mathfrak{A}})$ may contain infinite or even dense paths, so that it is hard to know where to start removing vertices. Our solution, paradoxically, is to *add* vertices. By a judicious choice of such additions, we are subsequently able to remove edges in such a way that no long paths remain. It is then straightforward to obtain a model subject to the desired size bound.

Fix some signature Σ featuring only unary predicates and the distinguished binary predicate $\mathfrak{t}$. Observe that the fluted 2-types over Σ are simply the fluted 1-types over Σ together with either of the fluted literals $\mathfrak{t}$ or $\neg\mathfrak{t}$. For the remainder of this section, we generally suppress reference to Σ. We find it helpful to think of an M-bounded fluted 1-profile ζ as a pair of functions (ζ^+, ζ^-), each mapping Ftp^Σ_1 to $[0, M]$, given by:

$$\zeta^+(\pi) = \zeta(\pi \wedge \mathfrak{t}); \qquad \zeta^-(\pi) = \zeta(\pi \wedge \neg\mathfrak{t}).$$

Thus, the components ζ^+ and ζ^- simply separate out the "positive" and "negative" 2-types, from the point of view of the fluted atom $\mathfrak{t}$. If $\mathfrak{A}$ is a structure interpreting $\mathfrak{t}$, we call any maximal subset Q of A having the property that $\mathfrak{A} \models \mathfrak{t}[a, b]$ for all distinct $a, b \in Q$, a *clique*. Thus, A is partitioned into cliques. If Q is a clique with $|Q| > 1$, then, by transitivity, $\mathfrak{A} \models \mathfrak{t}[a, a]$ for all $a \in Q$. If $Q = \{a\}$, then it may or may not be the case that $\mathfrak{A} \models \mathfrak{t}[a, a]$. If Q and Q' are cliques and c an element, we write $\mathfrak{A} \models \mathfrak{t}[Q, c]$ to mean $\mathfrak{A} \models \mathfrak{t}[a, c]$ for some (equivalently, any) $a \in Q$; similarly for $\mathfrak{A} \models \mathfrak{t}[c, Q]$ and $\mathfrak{A} \models \mathfrak{t}[Q, Q']$.

Let $\zeta = (\zeta^+, \zeta^-)$ and $\eta = (\eta^+, \eta^-)$ be M-bounded fluted 1-profiles. We write $\zeta \preceq \eta$ if, for all $\pi \in \mathrm{Ftp}^\Sigma_1$, $\zeta^+(\pi) \geq \eta^+(\pi)$ and $\zeta^-(\pi) \leq \eta^-(\pi)$. Evidently, if $\mathfrak{A}$ is a structure and $\mathfrak{A} \models \mathfrak{t}[a, b]$, then $\mathrm{fpr}^{\mathfrak{A}}[a] \preceq \mathrm{fpr}^{\mathfrak{A}}[b]$. If ζ is an M-bounded, fluted 1-profile realized in $\mathfrak{A}$, we call the set of elements $B = \{b \in A : \mathrm{fpr}^{\mathfrak{A}}[b] = \zeta\}$ a *profile set*. A *cluster* is a weakly connected component of the directed graph $(B, \mathfrak{t}^{\mathfrak{A}} \cap B^2)$, where B is a profile set – that is, a maximal subset $C \subseteq B$ such that, for all distinct $a, b \in C$, there exists a sequence of elements $a = a_1, \ldots, a_s = b$ such that, for all i $(1 \leq i < s)$, either $\mathfrak{A} \models \mathfrak{t}[a_i, a_{i+1}]$ or $\mathfrak{A} \models \mathfrak{t}[a_{i+1}, a_i]$. Obviously, a cluster is a union of cliques. Observe that, for any cluster C, there is no triple of elements c, d, e with $c, e \in C$, $d \notin C$ such that $\mathfrak{A} \models \mathfrak{t}[c, d]$ and $\mathfrak{A} \models \mathfrak{t}[d, e]$. Indeed, otherwise, we have $\mathrm{fpr}^{\mathfrak{A}}[c] \preceq \mathrm{fpr}^{\mathfrak{A}}[d] \preceq \mathrm{fpr}^{\mathfrak{A}}[e] = \mathrm{fpr}^{\mathfrak{A}}[c]$ contradicting the supposition that $d \notin C$.

▶ **Lemma 4.** *Let C be a cluster in a structure $\mathfrak{A}$ whose elements have M-bounded, fluted 1-profile ζ, let $\pi \in \mathrm{Ftp}^\Sigma_1$ and suppose that either $\zeta^+(\pi) < M$ or $\zeta^-(\pi) < M$. Then any two elements of C are related by $\mathfrak{t}$ in $\mathfrak{A}$ to the same elements of A having fluted 1-type π.*

Proof. Since C is a weakly connected component of the directed graph $(B, \mathfrak{t}^{\mathfrak{A}} \cap B^2)$, where B is a profile set, it suffices to show that any two elements $c, d \in C$ such that $\mathfrak{A} \models \mathfrak{t}[c, d]$ are related by $\mathfrak{t}$ in $\mathfrak{A}$ to the same elements of A having fluted 1-type π. By transitivity, we have $\{b \in A : \mathfrak{A} \models \mathfrak{t}[c, b] \text{ and } \mathrm{ftp}^{\mathfrak{A}}[b] = \pi\} \supseteq \{b \in A : \mathfrak{A} \models \mathfrak{t}[d, b] \text{ and } \mathrm{ftp}^{\mathfrak{A}}[b] = \pi\}$, and similarly $\{b \in A : \mathfrak{A} \not\models \mathfrak{t}[c, b] \text{ and } \mathrm{ftp}^{\mathfrak{A}}[b] = \pi\} \subseteq \{b \in A : \mathfrak{A} \not\models \mathfrak{t}[d, b] \text{ and } \mathrm{ftp}^{\mathfrak{A}}[b] = \pi\}$. If $\zeta^+(\pi) < M$, then the former two sets have equal finite cardinality, and hence are equal; if $\zeta^-(\pi) < M$, then the latter two sets are. Either way, c and d are related by $\mathfrak{t}$ in $\mathfrak{A}$ to the same elements of A having fluted 1-type π. ◀

A clique Q in a cluster C in a structure $\mathfrak{A}$ will be called *the superior clique* of C if $\mathfrak{A} \models \mathfrak{t}[a, Q]$ for every $a \in C$. It is possible for there to be no superior clique in C, but if it exists, it is unique.

▶ **Lemma 5.** *Let C be a cluster in a structure $\mathfrak{A}$ whose elements have M-bounded, fluted 1-profile ζ. Let $\pi \in \mathrm{Ftp}_1^\Sigma$. If either $\zeta^+(\pi) < M$ or $\zeta^-(\pi) < M$ and there exists $b \in C$ such that $\mathrm{ftp}^\mathfrak{A}[b] = \pi$ and $a \in C$ such that $\mathfrak{A} \models \mathfrak{t}[a, b]$, then the superior clique of C exists and contains every element $c \in C$ such that $\mathrm{ftp}^\mathfrak{A}[c] = \pi$ and $\mathfrak{A} \models \mathfrak{t}[a, c]$ for some $a \in C$.*

Proof. By Lemma 4, if such a and b exist, then every element of C is related to b by $\mathfrak{t}^\mathfrak{A}$, whence b is in the superior clique. But there can only be one superior clique. ◀

An important notion in the ensuing argument is that of a *stable* clique. If $\mathfrak{A}$ is a structure interpreting Σ, π a fluted 1-type and Q a clique of $\mathfrak{A}$ included in a cluster C with M-bounded, fluted 1-profile ζ, we say that π is *sensitive* for Q if

$$|\{b \in (A \setminus C) \cup Q : \mathfrak{A} \models \mathfrak{t}[Q, b] \text{ and } \mathrm{ftp}^\mathfrak{A}[b] = \pi\}| < \zeta^+(\pi).$$

Intuitively, π is sensitive for Q if the elements of Q need to be related by $\mathfrak{t}$ to elements having fluted 1-type π belonging to *other* cliques of C in order to make up the quota of witnesses demanded by ζ^+ – in other words, if there are not enough witnesses either outside C or inside Q. We call a clique Q *stable* if, for every fluted 1-type π sensitive for Q and every element $a \in A \setminus C$ such that $\mathfrak{A} \not\models \mathfrak{t}[a, Q]$, $|\{c \in C : \mathfrak{A} \not\models \mathfrak{t}[a, c] \text{ and } \mathrm{ftp}^\mathfrak{A}[c] = \pi\}| \geq M$. That is, Q is stable if, for every fluted 1-type π that is sensitive for Q, every element not related by $\mathfrak{t}$ to Q is not related to at least M elements of C satisfying π. A special case of stability is where Q has no sensitive fluted 1-types; in that case, we say that Q is *trivially stable.*

▶ **Lemma 6.** *If $\mathfrak{A}$ is a structure, and Q a clique of $\mathfrak{A}$ included in a cluster C, then either Q is stable, or there exists a stable clique $Q' \subseteq C$ such that $\mathfrak{A} \models \mathfrak{t}[Q, Q']$.*

Proof. Given the clique Q and cluster $C \supseteq Q$, define the procedure `stabilize` as shown in Fig. 1, in which the auxiliary variable f stores a function mapping the set of fluted 1-types Ftp_1^Σ to $[0, M]$. The goal of `stabilize` is to find the clique Q' guaranteed by the lemma. The idea is to examine the list Π of sensitive fluted 1-types for the clique currently under consideration (initially Q). By selecting a $\pi \in \Pi$ which has been encountered least often, we move along a $\mathfrak{t}$-edge to another clique of C in which π is realized. Before any execution of line 6, Π will have been assigned the set of fluted 1-types sensitive for P; hence $\pi \in \Pi$ implies that there exists $b \in (C \setminus P)$ such that $\mathfrak{A} \models \mathfrak{t}[P, b]$ and $\mathrm{ftp}^\mathfrak{A}[b] = \pi$, whence the instruction in line 7 can be executed. Furthermore, any run of `stabilize` terminates, since one of the values $f(\pi)$ less than $M + 1$ is incremented by every execution of line 8. Let Q' be the clique returned as the value of P in line 11. We claim that Q' is stable. Indeed, suppose π is sensitive for Q'. Then there is a clique path $Q_1, \ldots Q_M, Q'$ with Q_i containing an element of type π, say c_i, for all i $(1 \leq i \leq M)$. Hence, if $a \in A \setminus C$ with $\mathfrak{A} \not\models \mathfrak{t}[a, Q']$, then $\mathfrak{A} \not\models \mathfrak{t}[a, c_i]$ for all i $(1 \leq i \leq M)$, as claimed. ◀

We can now prove the key lemma concerning $\mathcal{FLU}_u^2$+1Tr. We say that a *clique path* in a structure $\mathfrak{A}$ is a sequence of distinct cliques $Q_1, \ldots, Q_s$ $(s \geq 1)$ such that $\mathfrak{A} \models \mathfrak{t}[Q_i, Q_{i+1}]$ for all i $(1 \leq i < s)$; the *length* of the clique path is $s - 1$. The *depth* of $\mathfrak{A}$ is the maximal length of any clique path plus 1 (∞ if there is no maximum).

```
begin stabilize
    for all fluted 1-types π do f(π) ← 0
    P ← Q
    Π ← the set of fluted 1-types sensitive for P
    until f(π) = M + 1 for all π ∈ Π
        choose π ∈ Π for which f(π) is smallest
        choose a clique P′ ⊆ C containing some element of 1-type π
            with P′ ≠ P and 𝔄 ⊨ 𝔱[P, P′]
        f(π) ← f(π) + 1
        P ← P′
        let Π be the set of fluted 1-types sensitive for P
    return P
end stabilize
```

Figure 1 Procedure `stabilize` of Lemma 6.

▶ **Lemma 7.** *Let φ be a formula of $\mathcal{FLU}^2_u$+1Tr of the form (4) with $k = 2$, featuring the constant M. Let Σ be the signature of φ. If φ is satisfiable, then φ has a model of depth $(M+1) \cdot 2^{|\Sigma|+1}$.*

Proof. Suppose $\mathfrak{A} \models \varphi$. By the downward Löwenheim-Skolem theorem, we may assume that $\mathfrak{A}$ is finite or countably infinite.

Stage 1. The goal of this stage is to flatten any clusters having superior cliques. We modify $\mathfrak{A}$, proceeding cluster by cluster. Number the clusters of $\mathfrak{A}$ as $C_0, C_1, \ldots$, and let $\mathfrak{A}_0 = \mathfrak{A}$. Supposing $\mathfrak{A}_i$ to have been defined, we define $\mathfrak{A}_{i+1}$, over the same domain, A. If C_i lacks a superior clique, we do nothing, and set $\mathfrak{A}_{i+1} = \mathfrak{A}_i$. Otherwise, let Q be the superior clique of C_i; thus, $\mathfrak{A} \models \mathfrak{t}[b, Q]$ for all $b \in C_i$. We define $\mathfrak{A}_{i+1}$ to be the same as $\mathfrak{A}_i$, except that we take $\mathfrak{t}^{\mathfrak{A}_{i+1}}$ on Q to be $C_i \times Q$: i.e. $\mathfrak{A}_{i+1} \models \mathfrak{t}[c, e]$ if and only if either (i) $\mathfrak{A}_i \models \mathfrak{t}[c, e]$ and $\{c, e\} \not\subseteq C_i$, or (ii) $c \in C_i$ and $e \in Q$. Thus, $\mathfrak{t}^{\mathfrak{A}_{i+1}} \subseteq \mathfrak{t}^{\mathfrak{A}_i}$, and the two relations agree on pairs of elements which do not both lie in C_i. That $\mathfrak{t}^{\mathfrak{A}_{i+1}}$ is transitive follows from the fact that there is no triple of elements c, d, e with $c, e \in C_i$, $d \notin C_i$ such that $\mathfrak{A}_i \models \mathfrak{t}[c, d]$ and $\mathfrak{A}_i \models \mathfrak{t}[d, e]$.

We claim that the M-bounded, fluted 1-profiles are preserved in the transition from $\mathfrak{A}_i$ to $\mathfrak{A}_{i+1}$. Consider $a \in A$ and define $\zeta = \mathrm{fpr}^{\mathfrak{A}_i}_M[a]$ and $\eta = \mathrm{fpr}^{\mathfrak{A}_{i+1}}_M[a]$; we show that $\zeta = \eta$. We may assume that $a \in C_i$, for otherwise there is nothing to show. Fix a fluted 1-type π. Since $\mathfrak{t}^{\mathfrak{A}_{i+1}} \subseteq \mathfrak{t}^{\mathfrak{A}_i}$, $\zeta^+(\pi) \geq \eta^+(\pi)$ and $\zeta^-(\pi) \leq \eta^-(\pi)$. To show the reverse comparisons, observe that we can find $\zeta^+(\pi)$ elements $b \in (A \setminus C_i) \cup Q$ such that $\mathrm{ftp}^{\mathfrak{A}_i}[b] = \pi$ and $\mathfrak{A}_i \models \mathfrak{t}[Q, b]$. But for each such b, $\mathfrak{A}_{i+1} \models \mathfrak{t}[a, b]$, by construction of $\mathfrak{A}_{i+1}$. Hence $\eta^+(\pi) \geq \zeta^+(\pi)$. Turning now to ζ^-, suppose $\zeta^-(\pi) < M$. Then, by Lemma 5 the only elements $b \in C_i$ such that $\mathfrak{A}_i \models \mathfrak{t}[a, b]$ and $\mathrm{ftp}^{\mathfrak{A}_i}[b] = \pi$ must lie in Q, and these witnesses are preserved in $\mathfrak{A}_{i+1}$. Hence the number of elements b such that $\mathfrak{A}_{i+1} \not\models \mathfrak{t}[a, b]$ and $\mathrm{ftp}^{\mathfrak{A}_{i+1}}[b] = \pi$ is at most $\zeta^-(\pi)$. This establishes our claim that $\zeta = \eta$, and hence, that the M-bounded, fluted 1-profiles are preserved. In particular, $\mathfrak{A}_{i+1} \models \varphi$.

From the above construction, for any pair of integers i, j, the induced substructure $\mathfrak{A}_k \restriction (A_i \cup A_j)$ stays the same for all $k > \max(i, j)$. Thus, it makes sense to speak about the limit structure $\mathfrak{A}'$ obtained by running the above process to infinity. Moreover, since each $\mathfrak{t}^{\mathfrak{A}_i}$ is transitive, $\mathfrak{t}^{\mathfrak{A}'}$ is also transitive, and $\mathfrak{A}' \models \varphi$. Finally, in the model $\mathfrak{A}'$, no cluster having a superior clique contains any clique path of length greater than 1.

Stage 2. The goal of this stage is to add elements to stable cliques in clusters lacking a superior clique. We modify $\mathfrak{A}'$, proceeding clique by clique. Number the cliques of $\mathfrak{A}'$ as $Q_0, Q_1, \ldots$, and let $\mathfrak{A}'_0 = \mathfrak{A}'$. Supposing $\mathfrak{A}'_i$ to have been defined, we build the structure $\mathfrak{A}'_{i+1}$. If Q_i either occurs in a cluster having a superior clique, or is not stable, we do nothing, and set $\mathfrak{A}'_{i+1} = \mathfrak{A}'_i$. Otherwise, we form $\mathfrak{A}'_{i+1}$ as follows. Letting Π be the set of fluted 1-types sensitive for Q_i, for each $\pi \in \Pi$, we add M elements to Q_i with fluted 1-type π. More precisely, let X be a set of $M \cdot |\Pi|$ fresh elements, and let $A'_{i+1} = A_i \cup X$. We define $\mathfrak{A}'_{i+1}$ on this set by declaring M elements of X to have fluted 1-type π for each $\pi \in \Pi$, setting $\mathfrak{t}^{\mathfrak{A}'_{i+1}}$ to be total on $Q_i \cup X$, and finally setting $\mathfrak{A}'_{i+1} \models \mathfrak{t}[a, Q_i \cup X]$ if and only if $\mathfrak{A}'_i \models \mathfrak{t}[a, Q_i]$ and $\mathfrak{A}'_{i+1} \models \mathfrak{t}[Q_i \cup X, a]$ if and only if $\mathfrak{A}'_i \models \mathfrak{t}[Q_i, a]$, for all $a \in A'_i \setminus Q_i$.

We again claim that the M-bounded fluted profiles of existing elements do not change. Given $a \in A'_i$, define $\zeta = \mathrm{fpr}_M^{\mathfrak{A}'_i}[a]$ and $\eta = \mathrm{fpr}_M^{\mathfrak{A}'_{i+1}}[a]$, and fix a fluted 1-type $\pi \in \Pi$; we show that $\zeta(\pi) = \eta(\pi)$. Since elements having fluted 1-type π are added in the construction of $\mathfrak{A}'_{i+1}$, $\eta^+(\pi) \geq \zeta^+(\pi)$ and $\eta^-(\pi) \geq \zeta^-(\pi)$; we establish the reverse inclusions.

If $\mathfrak{A}'_i \models \mathfrak{t}[a, Q_i]$, then a is unrelated by $\mathfrak{t}$ to exactly the same elements in both structures, whence $\eta^-(\pi) = \zeta^-(\pi)$. Now, let C be the cluster containing Q_i. Since π is sensitive for Q_i, there exists an element $b \in C \setminus Q_i$ such that $\mathrm{ftp}^{\mathfrak{A}'_i}[b] = \pi$ and $\mathfrak{A}'_i \models \mathfrak{t}[Q_i, b]$. And since C has no superior clique, by Lemma 5 we have $\zeta^+(\pi) = M \geq \eta^+(\pi)$. If, on the other hand, $\mathfrak{A}'_i \not\models \mathfrak{t}[a, Q_i]$, then a is related by $\mathfrak{t}$ to exactly the same elements in both structures, whence $\eta^+(\pi) = \zeta^+(\pi)$. Furthermore, since $\pi \in \Pi$ and Q_i is stable, a is unrelated by $\mathfrak{t}^{\mathfrak{A}'_i}$ to at least M elements of fluted 1-type π, whence $\zeta^-(\pi) = M \geq \eta^-(\pi)$. This establishes that $\zeta = \eta$, and hence, our claim that the M-bounded fluted 1-profiles are preserved.

Since the elements of X, i.e. the new elements of $\mathfrak{A}'_{i+1}$, are part of the clique $Q_i \cup X$ of $\mathfrak{A}'_{i+1}$, we see that $\mathfrak{A}'_{i+1}$ and $\mathfrak{A}'_i$ realize the same M-bounded fluted 1-profiles. Indeed, they realize the same M-bounded fluted 1-star-types. For even if Q_i contains no elements of fluted 1-type $\pi \in \Pi$, some other element in the same cluster as Q_i will. Hence, $\mathfrak{A}_{i+1} \models \varphi$. We remark also that no fluted 1-types are sensitive for the clique $Q_i \cup X$ of $\mathfrak{A}'_{i+1}$: any required witnesses are provided by X. That is, the clique $Q_i \cup X$ is trivially stable. Using the same reasoning as in Stage 1, we may speak about the limit structure $\mathfrak{A}''$ obtained by running the above process to infinity; evidently, $\mathfrak{t}^{\mathfrak{A}''}$ is transitive, and $\mathfrak{A}'' \models \varphi$. Finally, in the model $\mathfrak{A}''$, no cluster having a superior clique contains any clique path of length greater than 1, and, in all remaining clusters, all stable cliques are trivially stable.

Stage 3. The goal of this stage is to flatten any remaining clusters. We finally modify $\mathfrak{A}''$, proceeding cluster by cluster. Number the clusters of $\mathfrak{A}''$ as $D_0, D_1, \ldots$ and let $\mathfrak{A}''_0 = \mathfrak{A}''$. Supposing $\mathfrak{A}''_i$ to have been defined, we build the structure $\mathfrak{A}''_{i+1}$. If D_i has a superior clique, we do nothing, and set $\mathfrak{A}''_{i+1} = \mathfrak{A}''_i$. Otherwise, we call those elements of D_i lying in a stable clique *stable*, and the remainder *unstable*, and we define $\mathfrak{A}''_{i+1}$ to be the same as $\mathfrak{A}''_i$ except that the relation $\mathfrak{t}^{\mathfrak{A}''_{i+1}}$ on D_i is taken to be

$$\{\langle a, b\rangle : \mathfrak{A}''_i \models \mathfrak{t}[a, b] \text{ and } \mathfrak{A}''_i \models \mathfrak{t}[b, a]\} \cup \{\langle a, b\rangle : \mathfrak{A}''_i \models \mathfrak{t}[a, b],\ a \text{ is unstable, and } b \text{ is stable}\}.$$

Thus, $\mathfrak{t}^{\mathfrak{A}''_{i+1}} \subseteq \mathfrak{t}^{\mathfrak{A}''_i}$, and the two relations agree on pairs of elements which do not both lie in D_i. To check that $\mathfrak{t}^{\mathfrak{A}''_{i+1}}$ is transitive, we recall that there is no triple of elements c, d, e with $c, e \in D_i$, $d \notin D_i$ such that $\mathfrak{A}''_i \models \mathfrak{t}[c, d]$ and $\mathfrak{A}''_i \models \mathfrak{t}[d, e]$. Furthermore, it is immediate from the above construction of $\mathfrak{A}''_{i+1}$ that no clique path in D_i has length greater than 1. Consider $a \in A''$ and define $\zeta = \mathrm{fpr}_M^{\mathfrak{A}''_i}[a]$ and $\eta = \mathrm{fpr}_M^{\mathfrak{A}''_{i+1}}[a]$; we show that $\zeta = \eta$. We may assume that $a \in D_i$ and D_i has no superior clique, for otherwise there is nothing to show. Since $\mathfrak{t}^{\mathfrak{A}''_{i+1}} \subseteq \mathfrak{t}^{\mathfrak{A}''_i}$, $\zeta^+(\pi) \geq \eta^+(\pi)$ and $\zeta^-(\pi) \leq \eta^-(\pi)$.

To show the reverse comparisons, consider first ζ^+ and η^+, and fix a fluted 1-type π. If the element a lies in a stable clique Q, then, by the construction of $\mathfrak{A}''$, Q is trivially stable, whence π is not sensitive for Q. In other words we can find $\zeta^+(\pi)$ elements $b \in (A'' \setminus D_i) \cup Q$ such that $\text{ftp}^{\mathfrak{A}''_i}[b] = \pi$ and $\mathfrak{A}''_i \models \mathfrak{t}[Q, b]$. But for each such b, $\mathfrak{A}''_{i+1} \models \mathfrak{t}[a, b]$, by construction of $\mathfrak{A}''_{i+1}$. On the other hand, if a lies in an unstable clique Q, by Lemma 6, there is a stable clique $Q' \subseteq D_i$ such that $\mathfrak{A}''_i \models \mathfrak{t}[Q, Q']$, whence the same reasoning applies. Hence $\eta^+(\pi) \geq \zeta^+(\pi)$.

Turning now to ζ^- and η^-, suppose $\zeta^-(\pi) < M$ (for otherwise $\zeta^-(\pi) \geq \eta^-(\pi)$ trivially). Then, since D_i by assumption contains no superior clique, it follows by Lemma 5 that there is no $b \in D_i$ such that $\mathfrak{A}''_i \models \mathfrak{t}[a, b]$ and $\text{ftp}^{\mathfrak{A}''_i}[b] = \pi$. That is: although a is unrelated to more elements in $\mathfrak{A}''_{i+1}$ than in $\mathfrak{A}''_i$, none of them has fluted 1-type π. This implies $\eta^-(\pi) \leq \zeta^-(\pi) < M$, since $\mathfrak{t}$ is unchanged on pairs of elements at least one of which is not in D_i. This establishes that $\zeta = \eta$. Hence, the same M-bounded fluted 1-star-types are realized in $\mathfrak{A}''_i$ and $\mathfrak{A}''_{i+1}$, and therefore $\mathfrak{A}''_{i+1} \models \varphi$.

Using the same reasoning as in Stage 1, we may speak about the limit structure $\mathfrak{B}$ obtained by running the above process to infinity; evidently, $\mathfrak{t}^{\mathfrak{B}}$ is transitive, and $\mathfrak{B} \models \varphi$. Finally, in the model $\mathfrak{B}$, no cluster contains any clique path of length greater than 1. Since, in any clique path $Q_1, \ldots, Q_s$, $\text{fpr}^{\mathfrak{B}}[Q_i] \preceq \text{fpr}^{\mathfrak{B}}[Q_{i+1}]$ for all i $(1 \leq i < s)$, and each $\text{fpr}^{\mathfrak{B}}[Q_i]$ consists of $2^{|\Sigma|}$ integers in the range $[0, M]$, the depth of $\mathfrak{B}$ is at most $2(M+1)2^{|\Sigma|}$. ◀

▶ **Lemma 8.** *Let φ be a formula of $\mathcal{FLU}^2_u$+1Tr of the form (4) with $k = 2$, featuring the constant M. If φ has a model of depth D, then it has a model of depth D in which all cliques contain at most M elements having any given fluted 1-type π.*

Proof. Let $\mathfrak{A}$ be a model of depth D. For every clique Q in $\mathfrak{A}$ and fluted 1-type π, select M elements (all if there are fewer) in Q having fluted 1-type π, and let $\mathfrak{B}$ be the restriction of $\mathfrak{A}$ to the selected elements. It is obvious that $\text{fst}^{\mathfrak{A}}_M[b] = \text{fst}^{\mathfrak{B}}_M[b]$ for all $b \in B$ and indeed that $\mathfrak{A}$ and $\mathfrak{B}$ realize the same M-bounded, fluted 1-star-types. ◀

▶ **Lemma 9.** *Let φ be a formula of $\mathcal{FLU}^2_u$+1Tr of the form (4) with $k = 2$. If φ is satisfiable, then it has a model of size bounded by $2^{2^{O(\#(\varphi))}}$.*

Proof. Let Σ be the signature of φ. We may assume that the constant M featured in φ as given by (4) is at least 1. Write $L = 2^{|\Sigma|-1}$ for the number of fluted 1-types over Σ. By Lemmas 7 and 8, let $\mathfrak{A} \models \varphi$ with $\mathfrak{A}$ of depth at most $D = (M+1) \cdot 2^{|\Sigma|+1}$, and in which all cliques contain at most M elements having any given fluted 1-type π. Define the *level* of any clique Q of $\mathfrak{A}$ to be the length of the longest clique path in $\mathfrak{A}$ ending in Q plus 1. Thus, the minimum possible level is 1, and the maximum, D. For each i in this range, let A_i be the union of those cliques at level i.

We select cliques from A_i as follows. At each level i $(1 \leq i \leq D)$, select, for each fluted 1-type π, up to $(M+1)$ distinct cliques containing elements whose fluted 1-type is π (if there are fewer than $(M+1)$, then select all of them). The total number of cliques thus selected for each level i is therefore at most $(M+1)L \leq 2ML$. Next, considering successive levels i in the range $[2, D]$ (starting with $i = 2$), for each already selected clique $Q \subseteq A_1 \cup \cdots \cup A_{i-1}$, and for each fluted 1-type π, select up to M distinct cliques $Q' \subseteq A_i \cup \cdots \cup A_D$ containing elements whose fluted 1-type is π such that $\mathfrak{A} \models \mathfrak{t}[Q, Q']$, and select at least M distinct cliques $Q' \subseteq A_i \cup \cdots \cup A_D$ containing elements whose fluted 1-type is π such that $\mathfrak{A} \not\models \mathfrak{t}[Q, Q']$ (in both cases, if there are fewer than M, then select all of them). The number of selected cliques on level i is thus at most

$$(2ML) + (2ML)^2 + \cdots + (2ML)^i \leq (2ML)^{i+1}.$$

Let the sub-structure consisting of the selected cliques be $\mathfrak{B}$. It is easy to see that every element of $\mathfrak{B}$ realizes the same M-bounded fluted 1-star-type that it realizes in $\mathfrak{A}$. By Lemma 3, $\mathfrak{B} \models \varphi$. It remains to establish the size of $\mathfrak{B}$. Summing over all levels i from 1 to D, the total number of selected cliques is at most

$$(2ML)^2 + (2ML)^3 + \cdots + (2ML)^{D+1} \leq (2ML)^{D+2}.$$

Bearing in mind that each clique contains at most $M < 2M$ elements of each of the L fluted 1-types, we obtain $|B| < (2ML)^{D+3} = \left(M \cdot 2^{|\Sigma|}\right)^{(M+1)\cdot 2^{|\Sigma|+1}+3}$. ◀

6 Unrestricted signatures and any number of variables

We begin this section by showing that $\mathcal{FLU}^2$+1Tr has the finite model property, and that Sat($\mathcal{FLU}^2$+1Tr) is in 3-NExpTime. We proceed to show that the entire logic $\mathcal{FLU}$+1Tr has the finite model property, and that Sat($\mathcal{FLU}^k$+1Tr) is in (k+1)-NExpTime for all $k \geq 2$.

▶ **Lemma 10.** *Given an $\mathcal{FLU}^2$+1Tr-formula φ, there exists an $\mathcal{FLU}^2_u$+1Tr-formula ψ such that ψ and φ are satisfiable over the same domains, and $\#(\psi) \leq 2^{\#(\varphi)}$.*

Proof. Let φ, over a signature Σ, be given by

$$\bigwedge_{h=1}^{m} \forall(\alpha_h \rightarrow Q(1, \Sigma, M, \Theta_h)),$$

where m is a positive integer, the α_h quantifier-free $\mathcal{FLC}^1$-formulas over Σ, M a non-negative integer and the Θ_h existential Presburger formulas with free variables $\bar{v}$ indexed by the fluted 2-types over Σ. Let Σ^- be the result of removing from Σ all binary predicates other than the distinguished predicate $\mathfrak{t}$, and let $\bar{u}$ be the set of variables u_ρ indexed by the fluted 2-types over Σ^-. For any index set $H \subseteq [1, m]$, define the Presburger formula $\Theta_H(\bar{u})$ to be

$$\exists(\bar{v} \leq M)\left(\bigwedge_{h\in H}\Theta_h(\bar{v}) \wedge \bigwedge_{\rho\in \mathrm{Ftp}_2^{\Sigma^-}}\left(u_\rho \geq \sum\{v_\tau : \tau \in \mathrm{Ftp}_2^\Sigma \text{ s.t. } \models \tau \rightarrow \rho\}\right) \wedge \right.$$
$$\bigwedge_{\rho\in \mathrm{Ftp}_2^{\Sigma^-}}\left(\bigwedge\{v_\tau < M : \tau \in \mathrm{Ftp}_2^\Sigma \text{ s.t. } \models \tau \rightarrow \rho\} \rightarrow \right.$$
$$\left.\left. u_\rho \leq \sum\{v_\tau : \tau \in \mathrm{Ftp}_2^\Sigma \text{ s.t. } \models \tau \rightarrow \rho\}\right)\right).$$

(We adopt the usual convention that $\bigwedge \emptyset = \top$.) Intuitively, $\Theta_H(\bar{u})$ is intended to characterize those fluted 1-profiles over Σ^- which can be consistently extended to a fluted 1-profile over Σ satisfying all of the existential Presburger formulas Θ_h, for $h \in H$. By renaming bound variables, any existential quantifiers in Θ_H can be moved to the front. Thus we may regard Θ_H as an existential Presburger formula. Now let $L = 2^{(|\Sigma|-|\Sigma^-|)}$, and let ψ be

$$\bigwedge_{H\subseteq[1,m]} \forall\left(\left(\bigwedge_{h\in H}\alpha_h\right) \rightarrow Q(1, \Sigma^-, LM, \Theta_H)\right).$$

It is immediate that $\#(\psi) = |\Sigma^-| + \log(LM) + 2^m \leq |\Sigma^-| + \left(|\Sigma| - |\Sigma^-| + \log M\right) + 2^m \leq 2^{\#(\varphi)}$. Note that the signature of ψ is Σ^-; in other words, ψ is an $\mathcal{FLU}^2_u$+1Tr-formula. We show that φ and ψ are satisfiable over the same domains.

Suppose $\mathfrak{A} \models \varphi$ and let $\mathfrak{A}^-$ be the reduct of $\mathfrak{A}$ to Σ^-. We show that $\mathfrak{A}^- \models \psi$. Fix $a \in A$, and suppose $H \subseteq \{h \in [1,m] \colon \mathfrak{A} \models \alpha_h[a]\}$. It suffices to show that the LM-bounded, fluted profile of a in $\mathfrak{A}^-$ satisfies $\Theta_H(\bar{u})$. Let us therefore assign values to the various u_ρ (for $\rho \in \text{Ftp}_2^{\Sigma^-}$) accordingly: $u_\rho \leftarrow \min\big(|\{b \in A \colon \mathfrak{A}^- \models \rho[a,b]\}|, LM\big)$. To show satisfaction of Θ_H, we must find values for the existentially quantified variables. For each $\tau \in \text{Ftp}_2^{\Sigma}$, we let the variable v_τ be assigned the value given by the M-bounded, fluted profile of a in $\mathfrak{A}$, namely: $v_\tau \leftarrow \min\big(|\{b \in A \colon \mathfrak{A} \models \tau[a,b]\}|, M\big)$. We consider the three groups of conjuncts in the body of Θ_H in turn. Since $\mathfrak{A} \models \varphi$, we certainly have (under this valuation) $\Theta_h(\bar{v})$ for all $h \in H$. For the second group of conjuncts, fix $\rho \in \text{Ftp}_2^{\Sigma^-}$, and suppose first that $u_\rho < LM$. Then

$$u_\rho = |\{b \in A \colon \mathfrak{A}^- \models \rho[a,b]\}| = \sum\{|\{b \in A \colon \mathfrak{A} \models \tau[a,b]\}| \ \colon \tau \in \text{Ftp}_2^{\Sigma} \text{ s.t.} \models \tau \to \rho\}$$
$$\geq \sum\{v_\tau \colon \tau \in \text{Ftp}_2^{\Sigma} \text{ s.t.} \models \rho \to \tau\},$$

as required. And of course, if $u_\rho \geq LM$, then the conjunct follows from the fact that $\bar{v} \leq M$ and $|\{\tau \in \text{Ftp}_2^{\Sigma} \colon \models \tau \to \rho\}| = L$. For the final group of conjuncts, again fix $\rho \in \text{Ftp}_2^{\Sigma^-}$, and suppose that $v_\tau < M$ for all v_τ such that $\tau \in \text{Ftp}_2^{\Sigma}$ and $\models \rho \to \tau$. Then, applying essentially similar reasoning,

$$u_\rho \leq |\{b \in A \colon \mathfrak{A}^- \models \rho[a,b]\}| = \sum\{|\{b \in A \colon \mathfrak{A} \models \tau[a,b]\}| \ \colon \tau \in \text{Ftp}_2^{\Sigma} \text{ s.t.} \models \tau \to \rho\}$$
$$= \sum\{v_\tau \colon \tau \in \text{Ftp}_2^{\Sigma} \text{ s.t.} \models \rho \to \tau\}.$$

Thus, a satisfies the formula $Q(1, \Sigma^-, LM, \Theta_H)$ in $\mathfrak{A}^-$, and hence $\mathfrak{A}^- \models \psi$.

Finally, suppose $\mathfrak{B} \models \psi$, with $\mathfrak{B}$ interpreting Σ^-. We expand to a model $\mathfrak{B}^+ \models \varphi$ by interpreting the predicates in $\Sigma \setminus \Sigma^-$. All such predicates are of course binary. Consider any $a \in B$, and let $H = \{h \in [1,m] \colon \mathfrak{B} \models \alpha_h[a]\}$. Thus, the LM-bounded, fluted profile of a in $\mathfrak{B}$ satisfies $\Theta_H(\bar{u})$. To avoid notational clutter, we write the names of the variables in $\bar{u}$ to denote their values under this assignment, so that u_ρ denotes the value $\big(\text{fpr}^{\mathfrak{B}}_{LM}[a]\big)(\rho)$. Similarly, we write the names of the variables in $\bar{v}$ to denote some collection of values witnessing the existentially quantified statement in Θ_H. Now fix some $\rho \in \text{Ftp}_2^{\Sigma^-}$ and let $B_\rho = \{b \in B \colon \text{ftp}^{\mathfrak{B}}[a,b] = \rho\}$. Thus, $u_\rho = \min(LM, |B_\rho|)$. We have two cases to consider. If $v_\tau < M$ for all $\tau \in \text{Ftp}_2^{\Sigma}$ such that $\models \tau \to \rho$, then, by the second and third conjunct groups of Θ_H, $u_\rho = \sum\{v_\tau \colon \tau \in \text{Ftp}_2^{\Sigma} \text{ s.t.} \models \tau \to \rho\}$, whence we may partition B_ρ into sets B_τ of cardinality v_τ, with τ varying over the fluted 2-types over Σ extending ρ. If, on the other hand, $v_\tau = M$ for some $\tau \in \text{Ftp}_2^{\Sigma}$ such that $\models \tau \to \rho$, pick any such τ, say τ_0. By the second conjunct group of Θ_H, $u_\rho \geq \sum\{v_\tau \colon \tau \in \text{Ftp}_2^{\Sigma} \text{ s.t.} \models \tau \to \rho\}$, so we may partition B_ρ into sets B_τ, such that B_τ has cardinality v_τ for all τ *except* τ_0, with B_{τ_0} mopping up any remaining elements. Observe that $|B_{\tau_0}| \geq M$. Having constructed the various sets B_τ, we partially define $\mathfrak{B}^+$ by interpreting the predicates of $\Sigma \setminus \Sigma^-$ (all binary) in such a way that $\mathfrak{B} \models \tau[a,b]$ for every $b \in B_\tau$, and every $\tau \in \text{Ftp}_2^{\Sigma}$ such that $\models \tau \to \rho$. However $\mathfrak{B}^+$ is completed, we are assured that $\text{ftp}^{\mathfrak{B}^+}[a] = \bar{v}$, and thus, by Θ_H, satisfies Θ_h for every $h \in H$. Now repeat this process for every $a \in B$. Because only *fluted* 2-types are being defined here, no pair of elements is re-assigned, and once all elements have been considered, $\mathfrak{B}^+$ will be fully defined, and will satisfy $\mathfrak{B}^+ \models \varphi$. ◀

It remains to show that $\mathcal{FLU}$+1Tr has the finite model property, and that Sat($\mathcal{FLU}^k$+1Tr) is in (k+1)-NExpTime for all $k \geq 2$. We proceed by reducing the problem Sat($\mathcal{FLU}^{k+1}$+1Tr) (with exponential blow-up) to the problem Sat($\mathcal{FLU}^k$+1Tr). Modulo some technical simplifications, this work repeats [27, Section 4], where a similar reduction was carried out in the absence of a distinguished predicate interpreted as a transitive relation. Since this distinguished predicate is binary, it is not affected by the reduction in question. Due to space limits, the proof of the following Lemma is therefore omitted.

▶ **Lemma 11.** *Given an $\mathcal{FLU}^{k+1}$+1Tr-formula φ ($k \geq 2$), there exists an $\mathcal{FLU}^k$+1Tr-formula ψ such that ψ and φ are satisfiable over the same domains, and $\#(\psi)$ is $2^{O(\#(\varphi))}$.*

▶ **Theorem 12.** *Let φ be a $\mathcal{FLU}^k$+1Tr-formula ($k \geq 2$). If φ is satisfiable, then it has a model of cardinality bounded by some fixed (k+1)-tuply exponential function of $\#(\varphi)$.*

Proof. Induction on k: the base case ($k = 2$) follows from Lemmas 9 and 10; the inductive case is Lemma 11. ◀

▶ **Corollary 13.** *$\mathcal{FLC}$+1Tr has the finite model property, and Sat($\mathcal{FLC}^k$+1Tr) is in (k+1)-NExpTime for all $k \geq 2$.*

Proof. Let a formula $\varphi \in \mathcal{FLC}^k$+1Tr be given. By Lemma 1, we may assume that φ is in normal form. Now re-write φ as a logically equivalent $\mathcal{FLU}^k$+1Tr-formula ψ with $\#(\psi) \leq \|\varphi\|$. By Theorem 12, ψ – and hence φ – has a model of size bounded by a (k+1)-tuply exponential function of $\#(\psi)$. The result then follows by standard model-checking techniques. ◀

7 $\mathcal{FLC}$ and two transitive relations

In this section we show that the satisfiability and finite satisfiability problems for $\mathcal{FLC}^2$+2Tr are both undecidable. The result holds when the signature features – besides the two distinguished transitive relations – only unary predicates. Thus, we strengthen the undecidability result for $\mathcal{SHQ}$ [18] where three roles were used. The proof proceeds by reduction from undecidable tiling problems that are typical for two-variable logics. For instance, this technique was used in [30] to show undecidability of the (finite) satisfiability problems for $\mathcal{FL}^2$ in the presence of three transitive relations.

A *tiling system* is a tuple $\boldsymbol{\mathcal{C}} = (\mathcal{C}, H, V)$, where $\mathcal{C}$ is a finite set of *tiles*, and $H, V \subseteq \mathcal{C} \times \mathcal{C}$ are the *horizontal* and *vertical* constraints. A *tiling* of $\mathbb{N}^2$ for $\boldsymbol{\mathcal{C}}$ is a function $f : \mathbb{N}^2 \to C$, such that for all $i, j \in \mathbb{N}$, $(f(i,j), f(i+1,j)) \in H$ and $(f(i,j), f(i,j+1)) \in V$. A tiling is *periodic* if there exist m, n such that, for all i and j, $f(i+m,j) = f(i,j+n) = f(i,j)$. The *infinite* (*periodic*) *tiling problem* is the following: given a tiling system $\boldsymbol{\mathcal{C}}$, does there exist a (periodic) tiling of $\mathbb{N}^2$ for $\boldsymbol{\mathcal{C}}$? Our proof relies on the following result (see e.g. [4, p. 90]).

▶ **Proposition 14.** *The periodic tiling problem and the complement of the infinite tiling problem are recursively inseparable.*

To achieve the goal of this section, given a tiling system $\boldsymbol{\mathcal{C}}$, we construct an $\mathcal{FLC}^2$+2Tr-formula $\eta_{\boldsymbol{\mathcal{C}}}$ such that: (i) if $\mathbb{N}^2$ has a periodic tiling for $\boldsymbol{\mathcal{C}}$, then $\eta_{\boldsymbol{\mathcal{C}}}$ is finitely satisfiable; and (ii) if $\eta_{\boldsymbol{\mathcal{C}}}$ is satisfiable, then $\mathbb{N}^2$ has a tiling for $\boldsymbol{\mathcal{C}}$. The result then follows from Prop. 14.

The formula $\eta_{\boldsymbol{\mathcal{C}}}$ features a conjunct φ_{grid} whose canonical model, shown in Fig. 2, has the domain $\mathbb{N} \cup \{-1\} \times \mathbb{N}$. The signature of φ_{grid} consists of the two distinguished binary predicates b (blue) and r (red), together with the unary predicates: $a_{i,j}$ ($0 \leq i, j \leq 3$) and c_i ($0 \leq i \leq 3$). The formula φ_{grid} is a conjunction enforcing the following properties: (a) there exists an initial element satisfying c_0, and the unary predicates enforce a partition of the universe; (b) witness requirements for elements forming the leftmost column; (c) witness requirements for elements not on the leftmost column; (d) confluence. Properties (a)-(c) are typical formulas of $\mathcal{FL}^2$. The only conjuncts where counting is used are the confluence conjuncts presented in Fig. 2. These conjuncts ensure that certain witnesses connected by the b and r relations must be identical. In this way we get a grid-like structure with short transitive paths that connect elements corresponding to both horizontal and vertical neighbours in a standard grid. One can also obtain finite models of φ_{grid} over a nearly toroidal grid structure $(\{-1\} \cup \mathbb{Z}_{4m}) \times \mathbb{Z}_{4m}$ $(m > 0)$ by identifying elements from columns 0 and $4m$ and from rows 0 and $4m$ of the canonical model.

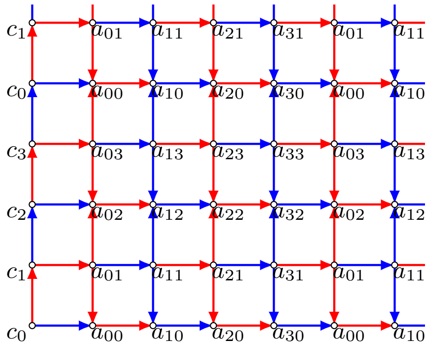

- Confluence conjuncts of φ_{grid}:

$$\bigwedge_{i\in\{1,3\}} \forall\,(c_i \to \exists_{[=1]}((b \vee r) \wedge a_{0,i+1}))$$

$$\bigwedge_{i\in\{0,2\}} \forall\,(c_i \to \exists_{[=1]}((b \vee r) \wedge a_{0,i}))$$

$$\bigwedge_{\substack{i\in\{0,1,2,3\}\\ j\in\{1,3\}}} \forall\,\big(a_{i,j} \to \exists_{[=1]}((b \vee r) \wedge a_{i+1,j+1}) \wedge \exists_{[=1]}((b \vee r) \wedge a_{i+1,j-1})\big).$$

Figure 2 Intended model of φ_{grid} (left) and the confluence conjuncts (right). The transitive relations b and r are depicted by blue and red arrows. Nodes with the coordinates $(-1, Y)$ satisfy the predicates $c_{Y \bmod 4}$; nodes with coordinates (X, Y) $(X \geq 0)$ satisfy the predicates $a_{X \bmod 4, Y \bmod 4}$. Addition and subtraction in indices of the confluence formulas are understood modulo 4.

The rest of the reduction is done in a standard fashion: using unary predicates representing tiles from $\mathcal{C}$ one adds to $\eta_{\mathcal{C}}$ conjuncts assigning tiles to elements of a model in such a way that the horizontal and vertical constraints are preserved. As a result one shows that: (i) from a periodic tiling of $\mathbb{N}^2$ a finite model of $\eta_{\mathcal{C}}$ can be built, and, (ii) from any model of $\eta_{\mathcal{C}}$ a tiling of $\mathbb{N}^2$ for $\mathcal{C}$ can be constructed. Hence we have

▶ **Theorem 15.** *The satisfiability problem and the finite satisfiability problem for $\mathcal{FLC}^2$+2Tr are both undecidable.*

We note that our proof of Theorem 15 is also valid when the two distinguished transitive relations are required to be *partial orders*. The same proof strategy does not work, however, if they are required to be *equivalence relations*. Nevertheless, it was shown in [26] that the satisfiability and finite satisfiability problems for the logic $\mathcal{C}^2$ with two equivalence relations are undecidable; and the formulas securing undecidability can easily be written as fluted formulas. The undecidability of the (finite) satisfiability problem for $\mathcal{FLC}^2$ with two equivalence relations then follows.

References

1 Bartosz Bednarczyk. Exploiting forwardness: Satisfiability and query-entailment in forward guarded fragment. In Wolfgang Faber, Gerhard Friedrich, Martin Gebser, and Michael Morak, editors, *Logics in Artificial Intelligence - 17th European Conference, JELIA 2021, Virtual Event, May 17-20, 2021, Proceedings*, volume 12678 of *Lecture Notes in Computer Science*, pages 179–193. Springer, 2021. doi:10.1007/978-3-030-75775-5_13.

2 Bartosz Bednarczyk, Emanuel Kieroński, and Piotr Witkowski. Completing the picture: Complexity of graded modal logics with converse. *Theory Pract. Log. Program.*, 21(4):493–520, 2021. doi:10.1017/S1471068421000065.

3 Michael Benedikt, Egor V. Kostylev, and Tony Tan. Two variable logic with ultimately periodic counting. In Artur Czumaj, Anuj Dawar, and Emanuela Merelli, editors, *47th International Colloquium on Automata, Languages, and Programming, ICALP 2020, July 8-11, 2020, Saarbrücken, Germany (Virtual Conference)*, volume 168 of *LIPIcs*, pages 112:1–112:16. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2020. doi:10.4230/LIPIcs.ICALP.2020.112.

4 Egon Börger, Erich Grädel, and Yuri Gurevich. *The Classical Decision Problem.* Perspectives in Mathematical Logic. Springer, 1997.

5 Diego Calvanese, Thomas Eiter, and Magdalena Ortiz. Answering regular path queries in expressive description logics via alternating tree-automata. *Information and Computation*, 237:12–55, 2014. doi:10.1016/j.ic.2014.04.002.

6 Ashok K. Chandra, Dexter Kozen, and Larry J. Stockmeyer. Alternation. *J. ACM*, 28(1):114–133, 1981. doi:10.1145/322234.322243.

7 Daniel Danielski and Emanuel Kieroński. Unary negation fragment with equivalence relations has the finite model property. In Anuj Dawar and Erich Grädel, editors, *Proceedings of the 33rd Annual ACM/IEEE Symposium on Logic in Computer Science, LICS 2018, Oxford, UK, July 09-12, 2018*, pages 285–294. ACM, 2018. doi:10.1145/3209108.3209205.

8 Daniel Danielski and Emanuel Kieroński. Finite satisfiability of unary negation fragment with transitivity. In Peter Rossmanith, Pinar Heggernes, and Joost-Pieter Katoen, editors, *44th International Symposium on Mathematical Foundations of Computer Science, MFCS 2019, August 26-30, 2019, Aachen, Germany*, volume 138 of *LIPIcs*, pages 17:1–17:15. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2019. doi:10.4230/LIPIcs.MFCS.2019.17.

9 Thomas Eiter, Magdalena Ortiz, and Mantas Simkus. Conjunctive query answering in the description logic SH using knots. *Journal of Computer and System Sciences*, 78(1):47–85, 2012. doi:10.1016/j.jcss.2011.02.012.

10 Birte Glimm, Ian Horrocks, and Ulrike Sattler. Unions of conjunctive queries in SHOQ. In Gerhard Brewka and Jérôme Lang, editors, *Principles of Knowledge Representation and Reasoning: Proceedings of the Eleventh International Conference, KR 2008, Sydney, Australia, September 16-19, 2008*, pages 252–262. AAAI Press, 2008. URL: http://www.aaai.org/Library/KR/2008/kr08-025.php.

11 Birte Glimm, Carsten Lutz, Ian Horrocks, and Ulrike Sattler. Conjunctive query answering for the description logic SHIQ. *J. Artif. Intell. Res.*, 31:157–204, 2008. doi:10.1613/jair.2372.

12 Tomasz Gogacz, Víctor Gutiérrez-Basulto, Yazmín Ibáñez-García, Jean Christoph Jung, and Filip Murlak. On finite and unrestricted query entailment beyond SQ with number restrictions on transitive roles. In Sarit Kraus, editor, *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence, IJCAI 2019, Macao, China, August 10-16, 2019*, pages 1719–1725. ijcai.org, 2019. doi:10.24963/ijcai.2019/238.

13 Víctor Gutiérrez-Basulto, Yazmín Angélica Ibáñez-García, and Jean Christoph Jung. Number restrictions on transitive roles in description logics with nominals. In Satinder Singh and Shaul Markovitch, editors, *Proceedings of the Thirty-First AAAI Conference on Artificial Intelligence, February 4-9, 2017, San Francisco, California, USA*, pages 1121–1127. AAAI Press, 2017. URL: http://aaai.org/ocs/index.php/AAAI/AAAI17/paper/view/14357.

14 Víctor Gutiérrez-Basulto, Yazmín Angélica Ibáñez-García, and Jean Christoph Jung. Answering regular path queries over SQ ontologies. In Sheila A. McIlraith and Kilian Q. Weinberger, editors, *Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence, (AAAI-18), the 30th innovative Applications of Artificial Intelligence (IAAI-18), and the 8th AAAI Symposium on Educational Advances in Artificial Intelligence (EAAI-18), New Orleans, Louisiana, USA, February 2-7, 2018*, pages 1845–1852. AAAI Press, 2018. URL: https://www.aaai.org/ocs/index.php/AAAI/AAAI18/paper/view/16242.

15 Ian Horrocks, Ulrike Sattler, and Stephan Tobies. Practical reasoning for very expressive description logics. *Logic Journal of the IGPL*, 8(3):239–263, 2000. doi:10.1093/jigpal/8.3.239.

16 Mark Kaminski and Gert Smolka. Terminating tableaux for $\mathcal{SOQ}$ with number restrictions on transitive roles. In Cristian S. Calude and Vladimiro Sassone, editors, *Theoretical Computer Science – 6th IFIP TC 1/WG 2.2 International Conference, TCS 2010, Held as Part of WCC 2010, Brisbane, Australia, September 20-23, 2010. Proceedings*, volume 323 of *IFIP Advances in Information and Communication Technology*, pages 213–228. Springer, 2010. doi:10.1007/978-3-642-15240-5_16.

17 Yevgeny Kazakov and Ian Pratt-Hartmann. A note on the complexity of the satisfiability problem for graded modal logics. In *Proceedings of the 24th Annual IEEE Symposium on Logic in Computer Science, LICS 2009, 11-14 August 2009, Los Angeles, CA, USA*, pages 407–416. IEEE Computer Society, 2009. `doi:10.1109/LICS.2009.17`.

18 Yevgeny Kazakov, Ulrike Sattler, and Evgeny Zolin. How many legs do I have? Non-simple roles in number restrictions revisited. In Nachum Dershowitz and Andrei Voronkov, editors, *Proc. of the 14th Int. Conf. on Logic for Programming, Artificial Intelligence, and Reasoning (LPAR 2007)*, volume 4790 of *Lecture Notes in Computer Science*, pages 303–317. Springer, 2007. `doi:10.1007/978-3-540-75560-9_23`.

19 Emanuel Kieroński. Results on the guarded fragment with equivalence or transitive relations. In C.-H. Luke Ong, editor, *Computer Science Logic, 19th International Workshop, CSL 2005, 14th Annual Conference of the EACSL, Oxford, UK, August 22-25, 2005, Proceedings*, volume 3634 of *Lecture Notes in Computer Science*, pages 309–324. Springer, 2005. `doi:10.1007/11538363_22`.

20 Emanuel Kieroński. On the complexity of the two-variable guarded fragment with transitive guards. *Inf. Comput.*, 204(11):1663–1703, 2006. `doi:10.1016/j.ic.2006.08.001`.

21 Emanuel Kieroński. Decidability issues for two-variable logics with several linear orders. In Marc Bezem, editor, *Computer Science Logic, 25th International Workshop / 20th Annual Conference of the EACSL, CSL 2011, September 12-15, 2011, Bergen, Norway, Proceedings*, volume 12 of *LIPIcs*, pages 337–351. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2011. `doi:10.4230/LIPIcs.CSL.2011.337`.

22 Emanuel Kieroński and Adam Malinowski. The triguarded fragment with transitivity. In Elvira Albert and Laura Kovács, editors, *LPAR 2020: 23rd International Conference on Logic for Programming, Artificial Intelligence and Reasoning, Alicante, Spain, May 22-27, 2020*, volume 73 of *EPiC Series in Computing*, pages 334–353. EasyChair, 2020. `doi:10.29007/z359`.

23 Emanuel Kieroński and Sebastian Rudolph. Finite model theory of the triguarded fragment and related logics. In *36th Annual ACM/IEEE Symposium on Logic in Computer Science, LICS 2021, Rome, Italy, June 29 - July 2, 2021*, pages 1–13. IEEE, 2021. `doi:10.1109/LICS52264.2021.9470734`.

24 Emanuel Kieroński and Lidia Tendera. Finite satisfiability of the two-variable guarded fragment with transitive guards and related variants. *ACM Trans. Comput. Log.*, 19(2):8:1–8:34, 2018. `doi:10.1145/3174805`.

25 Richard E. Ladner. The computational complexity of provability in systems of modal propositional logic. *SIAM J. Comput.*, 6(3):467–480, 1977. `doi:10.1137/0206033`.

26 Ian Pratt-Hartmann. The two-variable fragment with counting and equivalence. *Mathematical Logic Quarterly*, 61(6):474–515, 2015. `doi:10.1002/malq.201400102`.

27 Ian Pratt-Hartmann. Fluted logic with counting. In Nikhil Bansal, Emanuela Merelli, and James Worrell, editors, *48th International Colloquium on Automata, Languages, and Programming, ICALP 2021, July 12-16, 2021, Glasgow, Scotland (Virtual Conference)*, volume 198 of *LIPIcs*, pages 141:1–141:17. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2021. `doi:10.4230/LIPIcs.ICALP.2021.141`.

28 Ian Pratt-Hartmann, Wiesław Szwast, and Lidia Tendera. Quine's fluted fragment is non-elementary. In Jean-Marc Talbot and Laurent Regnier, editors, *25th EACSL Annual Conference on Computer Science Logic, CSL 2016, August 29 - September 1, 2016, Marseille, France*, volume 62 of *LIPIcs*, pages 39:1–39:21. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2016. `doi:10.4230/LIPIcs.CSL.2016.39`.

29 Ian Pratt-Hartmann, Wiesław Szwast, and Lidia Tendera. The fluted fragment revisited. *Journal of Symbolic Logic*, 84(3):1020–1048, 2019. `doi:10.1017/jsl.2019.33`.

30 Ian Pratt-Hartmann and Lidia Tendera. The fluted fragment with transitive relations. *Annals of Pure and Applied Logic*, 173(1):103042, 2022. `doi:10.1016/j.apal.2021.103042`.

31 William C. Purdy. Fluted formulas and the limits of decidability. *Journal of Symbolic Logic*, 61(2):608–620, 1996. `doi:10.2307/2275678`.

32 William C. Purdy. Complexity and nicety of fluted logic. *Studia Logica*, 71:177–198, 2002. doi:10.1023/A:1016596721799.

33 Sebastian Rudolph. Undecidability results for database-inspired reasoning problems in very expressive description logics. In Chitta Baral, James P. Delgrande, and Frank Wolter, editors, *Principles of Knowledge Representation and Reasoning: Proceedings of the Fifteenth International Conference, KR 2016, Cape Town, South Africa, April 25-29, 2016*, pages 247–257. AAAI Press, 2016. URL: http://www.aaai.org/ocs/index.php/KR/KR16/paper/view/12787.

34 Wiesław Szwast and Lidia Tendera. The guarded fragment with transitive guards. *Ann. Pure Appl. Log.*, 128(1-3):227–276, 2004. doi:10.1016/j.apal.2004.01.003.

35 Stephan Tobies. Complexity results and practical algorithms for logics in knowledge representation, 2001. PhD Thesis, LuFG Theoretical Computer Science, RWTH-Aachen, Germany.

36 Stephan Tobies. PSPACE reasoning for graded modal logics. *J. of Logic and Computation*, 11(1):85–106, 2001. doi:10.1093/logcom/11.1.85.

37 Evgeny Zolin. Undecidability of the transitive graded modal logic with converse. *J. Log. Comput.*, 27(5):1399–1420, 2017. doi:10.1093/logcom/exw026.