

The Fluted Fragment with Transitivity

Ian Pratt-Hartmann 

University of Warsaw, Poland
University of Opole, Poland
University of Manchester, UK
ipratt@cs.man.ac.uk

Lidia Tendera 

University of Opole, Poland
tendera@math.uni.opole.pl

Abstract

We study the satisfiability problem for the fluted fragment extended with transitive relations. We show that the logic enjoys the finite model property when only one transitive relation is available. On the other hand we show that the satisfiability problem is undecidable already for the two-variable fragment of the logic in the presence of three transitive relations.

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

Keywords and phrases First-Order logic, Decidability, Satisfiability, Transitivity, Complexity

Digital Object Identifier 10.4230/LIPIcs.MFCS.2019.18

Related Version An extended version of the paper is available at <https://arxiv.org/abs/1906.09131>.

Funding This work is supported by the Polish National Science Centre grant 2018/31/B/ST6/03662.

1 Introduction

The *fluted fragment*, here denoted \mathcal{FL} , is a fragment of first-order logic in which, roughly speaking, the order of quantification of variables coincides with the order in which those variables appear as arguments of predicates. The allusion is presumably architectural: we are invited to think of arguments of predicates as being “lined up” in columns. The following formulas are sentences of \mathcal{FL}

$$\begin{aligned} &\text{No student admires every professor} \\ &\forall x_1(\text{student}(x_1) \rightarrow \neg \forall x_2(\text{prof}(x_2) \rightarrow \text{admires}(x_1, x_2))) \end{aligned} \tag{1}$$

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

with the “lining up” of variables illustrated in Fig. 1. By contrast, none of the formulas

$$\begin{aligned} &\forall x_1.r(x_1, x_1) \\ &\forall x_1 \forall x_2(r(x_1, x_2) \rightarrow r(x_2, x_1)) \\ &\forall x_1 \forall x_2 \forall x_3(r(x_1, x_2) \wedge r(x_2, x_3) \rightarrow r(x_1, x_3)), \end{aligned}$$

expressing, respectively, the reflexivity, symmetry and transitivity of the relation r , is fluted, as the atoms involved cannot be arranged so that their argument sequences “line up” in the fashion of Fig. 1.

The history of this fragment is somewhat tortuous. The basic idea of *fluted logic* can be traced to a paper given by W.V. Quine to the 1968 *International Congress of Philosophy* [19], in which the author defined the *homogeneous m-adic formulas*. Quine later relaxed this



© Ian Pratt-Hartmann and Lidia Tendera;

licensed under Creative Commons License CC-BY

44th International Symposium on Mathematical Foundations of Computer Science (MFCS 2019).

Editors: Peter Rossmanith, Pinar Heggernes, and Joost-Pieter Katoen; Article No. 18; pp. 18:1–18:15

Leibniz International Proceedings in Informatics



LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

$$\begin{array}{ccc}
& \forall x_1 & \\
& \text{(student}(x_1)) & \\
& \rightarrow \neg \forall x_2 & \\
& \text{(prof}(x_2)) & \\
\rightarrow \text{admires}(x_1, x_2))) & & \\
& & \forall x_1 \\
& & \text{(lecturer}(x_1)) \\
& & \rightarrow \neg \exists x_2 \\
& & \text{(prof}(x_2)) \\
& & \wedge \forall x_3 \\
& & \text{(student}(x_3))
\end{array}$$

■ **Figure 1** The “lining up” of variables in the fluted formulas (1) and (2); all quantification is executed on the right-most available column.

fragment, in the context of a discussion of predicate-functor logic, to what he called “fluted” quantificational schemata [20], claiming that the satisfiability problem for the relaxed fragment is decidable. The viability of the proof strategy sketched by Quine was explicitly called into question by Noah [12], and the subject then taken up by W.C. Purdy [17], who gave his own definition of “fluted formulas”, proving decidability. It is questionable whether Purdy’s reconstruction is faithful to Quine’s intentions: the matter is clouded by differences in the definitions of predicate functors between [12] and [20], both of which Purdy cites. In fact, Quine’s original definition of “fluted” quantificational schemata appears to coincide with a logic introduced – apparently independently – by A. Herzig [6]. Rightly or wrongly, however, the name “fluted fragment” has now attached itself to Purdy’s definition in [17]; and we shall continue to use it in that way in the present article. See Sec. 2 for a formal definition.

To complicate matters further, Purdy claimed in [18] that \mathcal{FL} (i.e. the fluted fragment, in our sense, and his) has the exponential-sized model property: if a fluted formula φ is satisfiable, then it is satisfiable over a domain of size bounded by an exponential function of the number of symbols in φ . Purdy concluded that the satisfiability problem for \mathcal{FL} is NEXPTIME-complete. These latter claims are false. It was shown in [14] that, although \mathcal{FL} has the finite model property, there is no elementary bound on the sizes of the models required, and the satisfiability problem for \mathcal{FL} is non-elementary. More precisely, define \mathcal{FL}^m to be the subfragment of \mathcal{FL} in which at most m variables (free or bound) appear. Then the satisfiability problem for \mathcal{FL}^m is $\lfloor m/2 \rfloor$ -NEXPTIME-hard for all $m \geq 2$ and in $(m-2)$ -NEXPTIME for all $m \geq 3$ [15]. It follows that the satisfiability problem for \mathcal{FL} is TOWER-complete, in the framework of [21]. These results fix the exact complexity of satisfiability of \mathcal{FL}^m for small values of m . Indeed, the satisfiability problem for FO^2 , the two-variable fragment of first-order logic, is known to be NEXPTIME-complete [5], whence the corresponding problem for \mathcal{FL}^2 is certainly in NEXPTIME. Moreover, for $0 \leq m \leq 1$, \mathcal{FL}^m coincides with the m -variable fragment of first-order logic, whence its satisfiability problem is NP-TIME-complete. Thus, taking 0-NEXPTIME to mean NP-TIME, we see that the satisfiability problem for \mathcal{FL}^m is $\lfloor m/2 \rfloor$ -NEXPTIME-complete, at least for $m \leq 4$.

The focus of the present paper is what happens when we add to the fluted fragment the ability to stipulate that certain designated binary relations are *transitive*, or are *equivalence relations*. The motivation comes from analogous results obtained for other decidable fragments of first-order logic. Consider basic propositional modal logic K. Under the standard translation into first-order logic (yielded by Kripke semantics), we can regard K as a fragment of first-order logic – indeed as a fragment of \mathcal{FL}^2 . From basic modal logic K, we obtain the logic K4 under the supposition that the accessibility relation on possible worlds is transitive, and the logic S5 under the supposition that it is an equivalence relation: it is well-known that the satisfiability problems for K and K4 are PSPACE-complete, whereas that for S5

is NPTIME-complete [11]. (For analogous results on *graded* modal logic, see [7].) Closely related are also description logics (cf. [2]) with *role hierarchies* and *transitive roles*. In particular, the description logic \mathcal{SH} , which has the finite model property, is an EXPTIME-complete fragment of \mathcal{FL} with transitivity. Similar investigations have been carried out in respect of FO^2 , which has the finite model property and whose satisfiability problem, as just mentioned, is NEXPTIME-complete. The finite model property is lost when one transitive relation or two equivalence relations are allowed. For equivalence, everything is known: the (finite) satisfiability problem for FO^2 in the presence of a single equivalence relation remains NEXPTIME-complete, but this increases to 2-NEXPTIME-complete in the presence of two equivalence relations [8, 9], and becomes undecidable with three. For transitivity, we have an incomplete picture: the *finite* satisfiability problem for FO^2 in the presence with a single transitive relation is decidable in 3-NEXPTIME [13], while the decidability of the satisfiability problem remains open (cf. [23]); the corresponding problems with two transitive relations are both undecidable [10].

Adding equivalence relations to the fluted fragment poses no new problems. Existing results on FO^2 with two equivalence relations can be used to show that the satisfiability and finite satisfiability problems for \mathcal{FL} (not just \mathcal{FL}^2) with *two* equivalence relations are decidable. Furthermore, the proof that the corresponding problems for FO^2 in the presence of *three* equivalence relations are undecidable can easily be seen to apply also to \mathcal{FL}^2 . On the other hand, the situation with transitivity is much less clear.

We show in the sequel that \mathcal{FL} in the presence of a single transitive relation has the finite model property. On the other hand, \mathcal{FL} with three transitive relations admits axioms of infinity and the corresponding satisfiability problem is undecidable even for the intersection of \mathcal{FL}^2 with the *guarded fragment* [1] (and the same holds even when one of these transitive relations is the identity). The status of \mathcal{FL} with just two transitive relations remains open. These can be contrasted with DLs where, to lose the finite model property, one needs to add to \mathcal{S} either both *inverses* and *number restrictions*, or the *self* operator (none expressible in \mathcal{FL}). We also want to point to another paper in this volume [4] and the references therein where the impact of adding transitivity to the *unary negation fragment* is discussed.

2 Preliminaries

Unless explicitly stated to the contrary, the fragments of first-order logic considered here do not contain equality. We employ purely relational signatures, i.e. no individual constants or function symbols. We do, however, allow 0-ary relations (proposition letters).

Let $\bar{x}_\omega = x_1, x_2, \dots$ be a fixed sequence of variables. We define the sets of formulas $\mathcal{FL}^{[m]}$ (for $m \geq 0$) by structural induction as follows: (i) any atom $\alpha(x_\ell, \dots, x_m)$, where x_ℓ, \dots, x_m is a contiguous subsequence of \bar{x}_ω , is in $\mathcal{FL}^{[m]}$; (ii) $\mathcal{FL}^{[m]}$ is closed under boolean combinations; (iii) if φ is in $\mathcal{FL}^{[m+1]}$, then $\exists x_{m+1}\varphi$ and $\forall x_{m+1}\varphi$ are in $\mathcal{FL}^{[m]}$. The set of *fluted formulas* is defined as $\mathcal{FL} = \bigcup_{m \geq 0} \mathcal{FL}^{[m]}$. A *fluted sentence* is a fluted formula with no free variables. Thus, when forming Boolean combinations in the fluted fragment, all the combined formulas must have as their free variables some suffix of some prefix x_1, \dots, x_m of \bar{x}_ω ; and, when quantifying, only the last variable in this sequence may be bound. Note also that proposition letters (0-ary predicates) may be combined freely with formulas: if φ is in $\mathcal{FL}^{[m]}$, then so, for example, is $\varphi \wedge P$, where P is a proposition letter.

Denote by \mathcal{FL}^m the sub-fragment of \mathcal{FL} consisting of those formulas featuring at most m variables, free or bound. Do not confuse \mathcal{FL}^m (the set of fluted formulas with at most m variables, free or bound) with $\mathcal{FL}^{[m]}$ (the set of fluted formulas with *free* variables x_ℓ, \dots, x_m).

These are, of course, quite different. For example, (1) is in \mathcal{FL}^2 , and (2) is in \mathcal{FL}^3 , but they are both in $\mathcal{FL}^{[0]}$. Note that \mathcal{FL}^m cannot include predicates of arity greater than m .

For $m \geq 2$, denote by $\mathcal{FL}^m kT$ the m -variable fluted fragment \mathcal{FL}^m together with k distinguished transitive relations. In addition, denote by $\mathcal{FL}^2 kT^u$ the sub-fragment of $\mathcal{FL}^2 kT$ in which no binary predicates occur except the k distinguished transitive ones.

3 The decidability of fluted logic with one transitive relation

In this section, we show that the logic $\mathcal{FL}1T$, the fluted fragment together with a single distinguished transitive relation t , has the finite model property. We proceed in stages. First, we show that $\mathcal{FL}^2 1T^u$ has a doubly exponential-sized model property. Next, we show that $\mathcal{FL}^2 1T$ has a triply exponential-sized model property, via an exponential-sized reduction to $\mathcal{FL}^2 1T^u$. Finally, for $m \geq 2$, we provide an exponential-sized reduction of the satisfiability problem for $\mathcal{FL}^{m+1} 1T$ to the corresponding problem for $\mathcal{FL}^m 1T$, showing that, if the target of the reduction has a model of size N , then the source has a model of size $O(2^N)$. The satisfiability problems considered here will all have at least exponential complexity. Therefore, we may assume without loss of generality in this section that all signatures feature no 0-ary predicates, since their truth values can simply be guessed.

3.1 The logic $\mathcal{FL}^2 1T^u$

Fix some signature Σ of unary predicates. We consider $\mathcal{FL}^2 1T^u$ -formulas over the signature $\Sigma \cup \{t\}$, where t is the distinguished transitive predicate. (Thus, $t \notin \Sigma$.) By a *1-type over Σ* , we mean a maximal consistent conjunction of literals $\pm p(x)$, where $p \in \Sigma$. If \mathfrak{A} is a structure interpreting $\Sigma \cup \{t\}$, any element $a \in A$ satisfies a unique 1-type over Σ ; we denote it $\text{tp}^{\mathfrak{A}}[a]$. Since Σ will not vary, we typically omit reference to it when speaking of 1-types. We use the letters π and π' always to range over 1-types and μ always to range over arbitrary quantifier-free Σ -formulas involving just the variable x . We write $\pi(y)$ to indicate the result of substituting y everywhere for x in π , and similarly for π' and μ .

Call a $\mathcal{FL}^2 1T^u$ -formula over $\Sigma \cup \{t\}$ *basic* if it is of one of the forms

$$\exists x.\mu \quad \forall x.\mu \quad \forall x(\pi \rightarrow \exists y(\mu(y) \wedge \pm t(x, y))) \quad \forall x(\pi \rightarrow \forall y(\pi'(y) \rightarrow \pm t(x, y))).$$

The following Lemma is a version of the familiar ‘‘Scott normal form’’ for \mathcal{FL}^2 from [22].

► **Lemma 1.** *Let φ be a $\mathcal{FL}^2 1T^u$ -sentence. There exists a set Ψ of basic $\mathcal{FL}^2 1T^u$ -formulas with the following properties: (i) $\models (\bigwedge \Psi) \rightarrow \varphi$; (ii) if φ has a model, then so has Ψ ; (iii) $\|\Psi\|$ is bounded by a polynomial function of $\|\varphi\|$.*

We say that a *super-type over Σ* is a pair $\langle \pi, \Pi \rangle$, where π is a 1-type over Σ and Π a set of 1-types over Σ . If \mathfrak{A} is a structure interpreting the signature $\Sigma \cup \{t\}$ and $a \in A$, the super-type of a in \mathfrak{A} , denoted $\text{stp}^{\mathfrak{A}}[a]$, is the pair $\langle \text{tp}^{\mathfrak{A}}[a], \Pi \rangle$, where $\Pi = \{\text{tp}^{\mathfrak{A}}[b] \mid \mathfrak{A} \models t[a, b]\}$. Intuitively, a super-type is a description of an element in a structure specifying that element’s 1-type together with the 1-types of those elements to which it is related by t . If S is a set of super-types, we write $\text{tp}(S) = \{\pi \mid \langle \pi, \Pi \rangle \in S \text{ for some } \Pi\}$. We usually omit Σ when speaking of super-types. By a *certificate*, we mean a pair $C = (S, \ll)$, where S is a set of super-types and \ll is a transitive relation on $\text{tp}(S)$ satisfying the following conditions:

(C1) if $\langle \pi, \Pi \rangle \in S$ and $\pi' \in \Pi$, then there exists $\langle \pi', \Pi' \rangle \in S$ with $\Pi' \subseteq \Pi$;

(C2) if $\pi \ll \pi'$, $\langle \pi, \Pi \rangle \in S$ and $\langle \pi', \Pi' \rangle \in S$, then $\{\pi'\} \cup \Pi' \subseteq \Pi$.

For a structure \mathfrak{A} , the *certificate of \mathfrak{A}* , denoted $C(\mathfrak{A})$, is the pair (S, \ll) , where $S = \{\text{stp}^{\mathfrak{A}}[a] \mid a \in A\}$ is the set of super-types realized in \mathfrak{A} , and $\pi \ll \pi'$ if and only if π and π' are realized in \mathfrak{A} and $\mathfrak{A} \models \forall x(\pi \rightarrow \forall y(\pi'(y) \rightarrow \mathfrak{t}(x, y)))$. Intuitively, a certificate is a description of a structure listing the realized super-types and containing a partial order which specifies when all elements realizing one 1-type are related by \mathfrak{t} to all elements realizing another 1-type.

► **Lemma 2.** *If \mathfrak{A} is any structure interpreting $\Sigma \cup \{\mathfrak{t}\}$, $C(\mathfrak{A})$ is a certificate.*

Proof. Write $C(\mathfrak{A}) = (S, \ll)$. Obviously \ll is transitive. We must check **(C1)** and **(C2)**.

(C1): Suppose $\langle \pi, \Pi \rangle \in S$ and $\pi' \in \Pi$. Let a be such that $\text{stp}^{\mathfrak{A}}[a] = \langle \pi, \Pi \rangle$. Then there exists a $b \in A$ such that $\text{tp}^{\mathfrak{A}}[b] = \pi'$ and $\mathfrak{A} \models \mathfrak{t}[a, b]$. Let $\text{stp}^{\mathfrak{A}}[b] = \langle \pi', \Pi' \rangle$.

(C2): Suppose $\langle \pi, \Pi \rangle \in S$ and $\langle \pi', \Pi' \rangle \in S$ with $\pi \ll \pi'$, and let $a, b \in A$ be such that $\text{stp}^{\mathfrak{A}}[a] = \langle \pi, \Pi \rangle$ and $\text{stp}^{\mathfrak{A}}[b] = \langle \pi', \Pi' \rangle$. Since $\pi \ll \pi'$, by construction of $C(\mathfrak{A})$, we have $\mathfrak{A} \models \forall x(\pi \rightarrow \forall y(\pi'(y) \rightarrow \mathfrak{t}(x, y)))$, whence it is immediate that $\pi' \in \Pi$ and $\Pi' \subseteq \Pi$. ◀

If $C = (S, \ll)$ is a certificate, and ψ a basic \mathcal{FL}^21T^u -formula, we define the relation $C \models \psi$ to hold provided the following six conditions are satisfied. The motivation for this definition is provided by Lemmas 3 and 4.

- (i) if ψ is of the form $\forall x(\pi \rightarrow \exists y(\mu(y) \wedge \mathfrak{t}(x, y)))$, then, for all Π such that $\langle \pi, \Pi \rangle \in S$, there exists $\pi' \in \Pi$ such that $\models \pi' \rightarrow \mu$;
- (ii) if ψ is of the form $\forall x(\pi \rightarrow \forall y(\pi'(y) \rightarrow \mathfrak{t}(x, y)))$ and $\pi, \pi' \in \text{tp}(S)$, then $\pi \ll \pi'$;
- (iii) if ψ is of the form $\forall x(\pi \rightarrow \exists y(\mu(y) \wedge \neg \mathfrak{t}(x, y)))$, then, for all $\langle \pi, \Pi \rangle \in S$, there exists $\langle \pi', \Pi' \rangle \in S$ such that $\models \pi' \rightarrow \mu$ and there exists no $\alpha \in \{\pi\} \cup \Pi$ such that $\alpha \ll \pi'$;
- (iv) if ψ is of the form $\forall x(\pi \rightarrow \forall y(\pi'(y) \rightarrow \neg \mathfrak{t}(x, y)))$, then, for all $\langle \pi, \Pi \rangle \in S$, $\pi' \notin \Pi$;
- (v) if ψ is of the form $\exists x.\mu$, then there exists $\langle \pi, \Pi \rangle \in S$ such that $\models \pi \rightarrow \mu$;
- (vi) if ψ is of the form $\forall x.\mu$, then, for all $\langle \pi, \Pi \rangle \in S$, $\models \pi \rightarrow \mu$.

► **Lemma 3.** *Let \mathfrak{A} be a structure interpreting $\Sigma \cup \{\mathfrak{t}\}$ and let ψ be a basic \mathcal{FL}^21T^u -formula over $\Sigma \cup \{\mathfrak{t}\}$. If $\mathfrak{A} \models \psi$, then $C(\mathfrak{A}) \models \psi$.*

Proof. We write $C(\mathfrak{A}) = (S, \ll)$ and consider the possible forms of ψ in turn.

- $\psi = \forall x(\pi \rightarrow \exists y(\mu(y) \wedge \mathfrak{t}(x, y)))$: Suppose $\langle \pi, \Pi \rangle \in S$. Then there exists $a \in A$ with $\text{stp}^{\mathfrak{A}}[a] = \langle \pi, \Pi \rangle$. Since $\mathfrak{A} \models \psi$, choose $b \in A$ such that $\mathfrak{A} \models \mu[b]$ and $\mathfrak{A} \models \mathfrak{t}[a, b]$, and let $\text{tp}^{\mathfrak{A}}[b] = \pi'$. Then $\models \pi' \rightarrow \mu$ and $\pi' \in \Pi$, as required.
- $\psi = \forall x(\pi \rightarrow \forall y(\pi'(y) \rightarrow \mathfrak{t}(x, y)))$: It is immediate by the construction of $C(\mathfrak{A})$ that, if $\pi, \pi' \in \text{tp}(S)$, then $\pi \ll \pi'$;
- $\psi = \forall x(\pi \rightarrow \exists y(\mu(y) \wedge \neg \mathfrak{t}(x, y)))$: Suppose $\langle \pi, \Pi \rangle \in S$. Then there exists $a \in A$ with $\text{stp}^{\mathfrak{A}}[a] = \langle \pi, \Pi \rangle$. Since $\mathfrak{A} \models \psi$, choose $b \in A$ such that $\mathfrak{A} \models \mu[b]$ and $\mathfrak{A} \not\models \mathfrak{t}[a, b]$, and let $\text{tp}^{\mathfrak{A}}[b] = \pi'$, so that $\models \pi' \rightarrow \mu$. We require only to show that there exists no $\alpha \in \{\pi\} \cup \Pi$ such that $\alpha \ll \pi'$. Suppose, for contradiction, that such an α exists. By **(C1)**, $\alpha \in \text{tp}(S)$. If $\alpha = \pi$, then, by the definition of \ll , we have $\mathfrak{A} \models \forall x(\pi \rightarrow \forall y(\pi'(y) \rightarrow \mathfrak{t}(x, y)))$, which contradicts the supposition that $\mathfrak{A} \not\models \mathfrak{t}[a, b]$. If $\alpha \in \Pi$, then, by the definition of Π and \ll , we have an element $a' \in A$ such that $\text{tp}^{\mathfrak{A}}[a'] = \alpha$, $\mathfrak{A} \models \mathfrak{t}[a, a']$ and $\mathfrak{A} \models \forall x(\alpha \rightarrow \forall y(\pi'(y) \rightarrow \mathfrak{t}(x, y)))$, which again contradicts the supposition that $\mathfrak{A} \not\models \mathfrak{t}[a, b]$.
- $\psi = \forall x(\pi \rightarrow \forall y(\pi'(y) \rightarrow \neg \mathfrak{t}(x, y)))$: Suppose $\langle \pi, \Pi \rangle \in S$ and let $a \in A$ be such that $\text{stp}^{\mathfrak{A}}[a] = \langle \pi, \Pi \rangle$. Since $\mathfrak{A} \models \psi$, we have $\pi' \notin \Pi$.
- $\psi = \exists x.\mu$ or $\psi = \forall x.\mu$. Immediate by construction of S . ◀

► **Lemma 4.** *If $C = (S, \ll)$ is a certificate, then there exists a structure \mathfrak{A} over a domain of cardinality $2|S|$ such that, for any basic \mathcal{FL}^21T^u -formula ψ over Σ , $C \models \psi$ implies $\mathfrak{A} \models \psi$.*

Proof. Define $A^+ = \{a_{\pi, \Pi}^+ \mid \langle \pi, \Pi \rangle \in S\}$ and $A^- = \{a_{\pi, \Pi}^- \mid \langle \pi, \Pi \rangle \in S\}$, where the various $a_{\pi, \Pi}^+$ and $a_{\pi, \Pi}^-$ are some objects (assumed distinct), and set $A = A^+ \cup A^-$. Define the binary relations $T_1 = \{\langle a_{\pi, \Pi}^\pm, a_{\pi', \Pi'}^\pm \rangle \mid \{\pi'\} \cup \Pi' \subseteq \Pi\}$ and $T_2 = \{\langle a_{\pi, \Pi}^\pm, a_{\pi', \Pi'}^\pm \rangle \mid \pi \ll \pi'\}$, and let T be the transitive closure of $T_1 \cup T_2$. Intuitively, we may think of the elements $a_{\pi', \Pi'}^\pm$ as witnessing existential formulas of the form $\exists y(\mu(y) \wedge \mathfrak{t}(x, y))$, where $\models \pi' \rightarrow \mu$, and of the elements $a_{\pi', \Pi'}^-$ as witnessing existential formulas of the form $\exists y(\mu(y) \wedge \neg \mathfrak{t}(x, y))$. Now define \mathfrak{A} on the domain A by setting $\text{tp}^{\mathfrak{A}}[a_{\pi, \Pi}^\pm] = \pi$ for all $\langle \pi, \Pi \rangle \in S$, and by setting $\mathfrak{t}^{\mathfrak{A}} = T$.

We observe that if $a = a_{\pi, \Pi}^\pm$ and $b = a_{\pi', \Pi'}^\pm$ are elements of A such that a is related to b by either T_1 or T_2 , then $\{\pi'\} \cup \Pi' \subseteq \Pi$. Indeed, for T_1 , this is immediate by definition; and for T_2 , it follows from property **(C2)** of certificates. It follows by induction that, if a is related to b by T , then $\{\pi'\} \cup \Pi' \subseteq \Pi$. To prove the lemma, we consider the possible forms of ψ in turn.

- $\psi = \forall x(\pi \rightarrow \exists y(\mu(y) \wedge \mathfrak{t}(x, y)))$: Suppose $a = a_{\pi, \Pi}^\pm$. Since $C \models \psi$, there exists $\pi' \in \Pi$ such that $\pi' \rightarrow \mu$. By **(C1)**, there exists $\langle \pi', \Pi' \rangle \in S$ such that $\Pi' \subseteq \Pi$. Letting $b = a_{\pi', \Pi'}^\pm$, we have that a is related to b by T_1 . But then $\mathfrak{A} \models \mathfrak{t}[a, b]$ and $\mathfrak{A} \models \mu[b]$ by construction of \mathfrak{A} .
- $\psi = \forall x(\pi \rightarrow \forall y(\pi'(y) \rightarrow \mathfrak{t}(x, y)))$: Since $C \models \psi$, we have $\pi \ll \pi'$. Suppose now $a = a_{\pi, \Pi}^\pm$ and $b = a_{\pi', \Pi'}^\pm$. Thus, a is related to b by T_2 , and so by construction of \mathfrak{A} , $\mathfrak{A} \models \mathfrak{t}[a, b]$.
- $\psi = \forall x(\pi \rightarrow \exists y(\mu(y) \wedge \neg \mathfrak{t}(x, y)))$: Suppose $a = a_{\pi, \Pi}^\pm$. Since $C \models \psi$, there exists $\langle \pi', \Pi' \rangle \in S$ such that $\pi' \rightarrow \mu$, and such that there is no $\alpha \in \{\pi\} \cup \Pi$ with $\alpha \ll \pi'$. Now let $b = a_{\pi', \Pi'}^-$. By construction of \mathfrak{A} , $\mathfrak{A} \models \mu[b]$. It suffices to show that $\mathfrak{A} \not\models \mathfrak{t}[a, b]$. For otherwise, by the definition of T , there exists a chain of elements $a = a_1, \dots, a_m = b$ with each related to the next by either T_1 or T_2 and with a_{m-1} related to a_m by T_2 . (Notice that nothing can be related by T_1 to $b = a_{\pi', \Pi'}^-$.) Writing $a_{m-1} = a_{\alpha, \Pi'}^\pm \in S$, we see that $\alpha \ll \pi'$, and, moreover, that a is either identical to a_{m-1} , or related to it by T . As we observed above, if $a_{\pi, \Pi}^\pm$ is related to $a_{\alpha, \Pi'}^\pm$ by T , then $\alpha \in \Pi$. Thus, either way, $\alpha \in \{\pi\} \cup \Pi$. But we are supposing that no such α exists.
- $\psi = \forall x(\pi \rightarrow \forall y(\pi'(y) \rightarrow \neg \mathfrak{t}(x, y)))$: Suppose $a = a_{\pi, \Pi}^\pm$ and $b = a_{\pi', \Pi'}^\pm$ are elements of A . We observed above that, if a is related to b by T , then $\pi' \in \Pi$, contradicting the assumption that $C \models \psi$. Thus, by construction of \mathfrak{A} , $\mathfrak{A} \not\models \mathfrak{t}[a, b]$.
- $\psi = \exists x.\mu$ or $\psi = \forall x.\mu$. Immediate by construction of \mathfrak{A} . ◀

Since the number of super-types over Σ is bounded by $2^{(2^{|\Sigma|} + |\Sigma|)}$, and a structure \mathfrak{A} can be guessed and verified to be a model of any m -variable first-order formula φ in time $O(|\varphi| \cdot |A|^m)$ [25], Lemmas 1–4 yield:

► **Lemma 5.** *If φ is a satisfiable formula of $\mathcal{FL}^2\text{1T}^u$, then φ has a model of size at most doubly exponential in $\|\varphi\|$. Hence the satisfiability problem for $\mathcal{FL}^2\text{1T}^u$ is in 2-NEXPTIME.*

3.2 The logics $\mathcal{FL}^m\text{1T}$ for $m \geq 2$

Let Σ be a signature of predicates of positive arity, excluding \mathfrak{t} . An atomic formula of $\mathcal{FL}^m\text{1T}$ involving a predicate from $\Sigma \cup \{\mathfrak{t}\}$ will be called a *fluted m -atom* over $\Sigma \cup \{\mathfrak{t}\}$. A *fluted m -literal* is a fluted m -atom or the negation thereof. A *fluted m -clause* is a disjunction of fluted m -literals. We allow the absurd formula \perp (i.e. the empty disjunction) to count as a fluted m -clause. Thus, any literal of a fluted m -clause has arguments x_h, \dots, x_m , in that order, for some h ($1 \leq h \leq m$). When writing fluted m -clauses, we silently remove bracketing, re-order literals and delete duplicated literals as necessary. A *fluted m -type* is a maximal consistent set of fluted m -literals; where convenient, we identify fluted m -types with their conjunctions. If \mathfrak{A} is a structure interpreting $\Sigma \cup \{\mathfrak{t}\}$, any tuple a_1, \dots, a_m from A satisfies a unique fluted m -type; we denote it $\text{ftp}^{\mathfrak{A}}[a_1, \dots, a_m]$. Note that a fluted 1-type over

$\Sigma \cup \{\mathfrak{t}\}$ coincides with what we earlier called a 1-type over Σ . Reference to the signature $\Sigma \cup \{\mathfrak{t}\}$ will as usual be suppressed when clear from context. Predicates in Σ will be referred to as *non-distinguished*. Our strategy will be to reduce the (finite) satisfiability problem for $\mathcal{FL}^m 1T$ to that for $\mathcal{FL}^2 1T$ (Lemma 11), and thence to that for $\mathcal{FL}^2 1T^u$ (Lemma 9), which we have already dealt with (Lemma 5).

A $\mathcal{FL}^m 1T$ -formula φ ($m \geq 2$) is in *clause normal form* if it is of the form

$$\forall x_1 \cdots x_m. \Omega \wedge \bigwedge_{i=1}^s \forall x_1 \cdots x_{m-1} (\alpha_i \rightarrow \exists x_m. \Gamma_i) \wedge \bigwedge_{j=1}^t \forall x_1 \cdots x_{m-1} (\beta_j \rightarrow \forall x_m. \Delta_j), \quad (3)$$

where $\Omega, \Gamma_1, \dots, \Gamma_s, \Delta_1, \dots, \Delta_t$ are sets of fluted m -clauses, and $\alpha_1, \dots, \alpha_s, \beta_1, \dots, \beta_t$ fluted $(m-1)$ -atoms. We refer to $\forall x_1 \cdots x_m. \Omega$ as the *static conjunct* of φ , to conjuncts of the form $\forall x_1 \cdots x_{m-1} (\alpha_i \rightarrow \exists x_m. \Gamma_i)$ as the *existential conjuncts* of φ , and to conjuncts of the form $\forall x_1 \cdots x_{m-1} (\beta_j \rightarrow \forall x_m. \Delta_j)$ as the *universal conjuncts* of φ .

Using the same techniques as for Lemma 1, we can transform any $\mathcal{FL}^m 1T$ -formula into clause normal form.

► **Lemma 6.** *Let φ be an $\mathcal{FL}^m 1T$ -formula, $m \geq 2$. There exists an $\mathcal{FL}^m 1T$ -formula ψ in clause normal form such that: (i) $\models \psi \rightarrow \varphi$; and (ii) if φ has a model then so has ψ ; (iii) $\|\psi\|$ is bounded by a polynomial function of $\|\varphi\|$.*

For fragments of first-order logic not involving equality, we are free to duplicate any element a in a structure \mathfrak{A} . More formally, we have the following lemma, which will be used as a step in the ensuing argument.

► **Lemma 7.** *Let \mathfrak{A} be any structure, and let $z > 0$. There exists a structure \mathfrak{B} such that (i) if φ is any first-order formula without equality, then $\mathfrak{A} \models \varphi$ if and only if $\mathfrak{B} \models \varphi$; (ii) $|B| = z \cdot |A|$; and (iii) if $\psi(x_1, \dots, x_{m-1}) = \exists x_m. \chi(x_1, \dots, x_m)$ is a first-order formula without equality, and $\mathfrak{B} \models \psi[b_1, \dots, b_{m-1}]$, then there exist at least z distinct elements b of B such that $\mathfrak{B} \models \chi[b_1, \dots, b_{m-1}, b]$.*

Keeping the signature Σ fixed, we employ the standard apparatus of resolution theorem-proving to eliminate non-distinguished predicates of arity 2 or more. Suppose $p \in \Sigma$ is a predicate of arity m , and let γ' and δ' be fluted m -clauses over Σ . Then, $\gamma = p(x_1, \dots, x_m) \vee \gamma'$ and $\delta = \neg p(x_1, \dots, x_m) \vee \delta'$ are also fluted m -clauses, as indeed is $\gamma' \vee \delta'$. In that case, we call $\gamma' \vee \delta'$ a *fluted resolvent* of γ and δ , and we say that $\gamma' \vee \delta'$ is *obtained by fluted resolution* from γ and δ on $p(x_1, \dots, x_m)$. Thus, fluted resolution is simply a restriction of the familiar resolution rule from first-order logic to the case where the resolved-on literals have maximal arity, m , and (in the case $m = 2$) do not feature the distinguished predicate \mathfrak{t} . It may be helpful to note the following at this point: (i) if γ and δ resolve to form ϵ , then $\models \forall x_1 \cdots \forall x_m (\gamma \wedge \delta \rightarrow \epsilon)$; (ii) the fluted resolvent of two fluted m -clauses may or may not involve predicates of arity m ; (iii) in fluted resolution, the arguments of the literals in the fluted m -clauses undergo no change when forming the resolvent; (iv) if the fluted m -clause γ involves no predicates of arity m , then it cannot undergo fluted resolution at all.

If Γ is a set of fluted m -clauses, denote by Γ^* the smallest set of fluted m -clauses including Γ and closed under fluted resolution. If $\Gamma = \Gamma^*$, we say that it is *closed under fluted resolution*. We further denote by Γ° the result of deleting from Γ^* any clause involving a *non-distinguished* predicate of arity m . Observe that, since all fluted m -atoms involving predicates of *non-maximal* arity are of the form $p(x_h, \dots, x_m)$ for some $h \geq 2$, it follows that Γ° features the variable x_1 *only* in the case $m = 2$, and even then only in literals of the form $\pm \mathfrak{t}(x_1, x_2)$.

The following lemma is, in effect, nothing more than the familiar completeness theorem for (ordered) propositional resolution. The proof is omitted due to space limits.

► **Lemma 8.** *Let Γ be a set of fluted m -clauses over a signature $\Sigma \cup \{\mathfrak{t}\}$, let Σ' be the result of removing all predicates of maximal arity m from Σ , and let τ^- be a fluted m -type over $\Sigma' \cup \{\mathfrak{t}\}$. If τ^- is consistent with Γ° , then there exists a fluted m -type τ over the signature $\Sigma \cup \{\mathfrak{t}\}$ such that $\tau \supseteq \tau^-$ and τ is consistent with Γ .*

The following lemma employs a technique from [13] to eliminate binary predicates.

► **Lemma 9.** *Let φ be an $\mathcal{FL}^2\text{1T}$ -formula in clause normal form over a signature $\Sigma \cup \{\mathfrak{t}\}$, and suppose that φ has s existential and t universal conjuncts. Then there exists a clause normal form $\mathcal{FL}^2\text{1T}^u$ -formula φ' over a signature $\Sigma' \cup \{\mathfrak{t}\}$ such that: (i) φ' has at most $2^t s$ existential and 2^t universal conjuncts; (ii) $|\Sigma'| \leq |\Sigma| + 2^t(s+1)$; (iii) if φ has a model, so does φ' ; and (iv) if φ' has a model of size M , then φ has a model of size at most sM .*

Proof. Let $\varphi = \forall x_1 x_2. \Omega \wedge \bigwedge_{i=1}^s \forall x_1 (p_i(x_1) \rightarrow \exists x_2. \Gamma_i) \wedge \bigwedge_{j=1}^t \forall x_1 (q_j(x_1) \rightarrow \forall x_2. \Delta_j)$, where $\Omega, \Gamma_1, \dots, \Gamma_s, \Delta_1, \dots, \Delta_t$ are sets of fluted 2-clauses, and $p_1, \dots, p_s, q_1, \dots, q_t$ unary predicates. Write $T = \{1, \dots, t\}$. For all i ($1 \leq i \leq s$) and all $J \subseteq T$, let $p_{i,J}$ and q_J be new unary predicates. The intended interpretation of $p_{i,J}(x_1)$ is “ x_1 satisfies p_i , and also satisfies q_j for every $j \in J$,” and the intended interpretation of $q_J(x_1)$ is “ x_1 satisfies q_j for every $j \in J$.” Let φ' be the conjunction of the sentences:

- (a) $\bigwedge_{i=1}^s \bigwedge_{J \subseteq T} \forall x_2 ((p_i(x_2) \wedge \bigwedge_{j \in J} q_j(x_2)) \rightarrow p_{i,J}(x_2))$,
- (b) $\bigwedge_{J \subseteq T} \forall x_2 ((\bigwedge_{j \in J} q_j(x_2)) \rightarrow q_J(x_2))$,
- (c) $\bigwedge_{i=1}^s \bigwedge_{J \subseteq T} \forall x_1 (p_{i,J}(x_1) \rightarrow \exists x_2 (\Gamma_i \cup \Omega \cup \bigcup \{\Delta_j \mid j \in J\})^\circ)$, and
- (d) $\bigwedge_{J \subseteq T} \forall x_1 (q_J(x_1) \rightarrow \forall x_2 (\Omega \cup \bigcup \{\Delta_j \mid j \in J\})^\circ)$.

Observe that φ' contains no non-distinguished binary predicates, and hence is in $\mathcal{FL}^2\text{1T}^u$. Clearly, φ' satisfies properties (i) and (ii). To show (iii), suppose $\mathfrak{A} \models \varphi$, and let \mathfrak{A}' be the structure obtained by interpreting the predicates $p_{i,J}$ and q_J as suggested above. To show (iv), suppose φ' has a model of size M . By Lemma 7, φ' has a model \mathfrak{B} of size sM in which witnesses for all the conjuncts in (c) are duplicated s times. We need to show that \mathfrak{B} can be expanded to a model of φ . Fix $a \in B$ and suppose a satisfies p_1 . Let J be the set of indices j such that a satisfies q_j . By (a), putting $i = 1$, a satisfies $p_{1,J}$, whence, by (c), there exists b such that the pair $\langle a, b \rangle$ satisfies $(\Gamma_1 \cup \Omega \cup \bigcup \{\Delta_j \mid j \in J\})^\circ$. But Lemma 8 guarantees that we can expand \mathfrak{B} by interpreting the non-distinguished binary predicates so that $\langle a, b \rangle$ satisfies $\Gamma_1 \cup \Omega \cup \bigcup \{\Delta_j \mid j \in J\}$. Because of the duplication of witnesses, we can repeat with p_2, \dots, p_s , choosing a fresh witness each time, so as to avoid clashes. Do this for all elements a . At the end of the process, the partially defined expansion of \mathfrak{B} satisfies all the existential conjuncts of φ , and violates none of the universal or static conjuncts. A precisely similar argument shows that we may complete the expansion so that no universal or static conjuncts of φ are violated. ◀

Thus, at the expense of an exponentially larger signature, we have reduced the (finite) satisfiability problem for $\mathcal{FL}^2\text{1T}$ to that for $\mathcal{FL}^2\text{1T}^u$. By Lemmas 5 and 9, we obtain

► **Lemma 10.** *Let φ be a $\mathcal{FL}^2\text{1T}$ -formula. If φ is satisfiable, then φ has a model of size at most triply exponential in $\|\varphi\|$. Hence the satisfiability problem for $\mathcal{FL}^2\text{1T}$ is in 3-NEXPTIME.*

We now establish the finite model property for the whole of $\mathcal{FL}\text{1T}$ by eliminating variables from $\mathcal{FL}^{m+1}\text{1T}$, where $m \geq 2$, one at a time. The proof of the following Lemma is similar to the proof of Lemma 9 and is omitted due to space limits.

► **Lemma 11.** *Let φ be a clause normal form $\mathcal{FL}^{m+1}\text{T}$ -formula ($m \geq 2$) over a signature $\Sigma \cup \{\mathfrak{t}\}$, and suppose that φ has s existential conjuncts and t universal conjuncts. Then there exists a clause normal form $\mathcal{FL}^m\text{T}$ -formula φ' over a signature $\Sigma' \cup \{\mathfrak{t}\}$ such that the following hold: (i) φ' has at most $2^t s$ existential and 2^t universal conjuncts; (ii) $|\Sigma'| \leq |\Sigma| + 2^t(s+1)$; (iii) if φ has a model, so does φ' ; and (iv) if φ' has a model of size M , then φ has a model of size at most sM .*

► **Theorem 12.** *Let φ be a $\mathcal{FL}^m\text{T}$ -formula for $m \geq 2$. If φ is satisfiable, then φ has a model of size at most $(m+1)$ -tuply exponential in $\|\varphi\|$. Hence the satisfiability problem for $\mathcal{FL}^m\text{T}$ is in non-deterministic $(m+1)$ -tuply exponential time.*

Proof. Induction on m . The case $m = 2$ is Lemma 10. The inductive step is Lemma 11. ◀

We mentioned in Sec. 1 that [14] establishes a lower bound of $\lfloor m/2 \rfloor$ -NEXPTIME-hard for the satisfiability problem for \mathcal{FL}^m . For $m \geq 3$, this appears to be the best available lower bound on the corresponding problem for $\mathcal{FL}^m\text{T}$. Thus, a gap remains between the best available upper and lower complexity bounds. Certainly, it follows that the satisfiability problem for $\mathcal{FL}\text{T}$ is TOWER-complete, as for \mathcal{FL} .

4 Fluted Logic with more Transitive Relations

In this section we show two undecidability results for the fluted fragment with two variables, \mathcal{FL}^2 , extended with more transitive relations, that have been informally announced in [24]. We employ the apparatus of tiling systems.

A *tiling system* is a tuple $\mathcal{C} = (\mathcal{C}, \mathcal{C}_H, \mathcal{C}_V)$, where \mathcal{C} is a finite set of *tiles*, and $\mathcal{C}_H, \mathcal{C}_V \subseteq \mathcal{C} \times \mathcal{C}$ are the *horizontal* and *vertical* constraints.

Let S be any of the spaces $\mathbb{N} \times \mathbb{N}$, $\mathbb{Z} \times \mathbb{Z}$ or $\mathbb{Z}_t \times \mathbb{Z}_t$. A tiling system \mathcal{C} *tiles* S , if there exists a function $\rho : S \rightarrow \mathcal{C}$ such that for all $(p, q) \in S$: $(\rho(p, q), \rho(p+1, q)) \in \mathcal{C}_H$ and $(\rho(p, q), \rho(p, q+1)) \in \mathcal{C}_V$. The following problems are known to be undecidable (cf. e.g. [3]):

- Given a tiling system \mathcal{C} determine if \mathcal{C} tiles $\mathbb{Z} \times \mathbb{Z}$, or $\mathbb{N} \times \mathbb{N}$.
- Given a tiling system \mathcal{C} determine if \mathcal{C} tiles $\mathbb{Z}_t \times \mathbb{Z}_t$, for some $t \geq 1$.

In this section we first prove the following theorem.

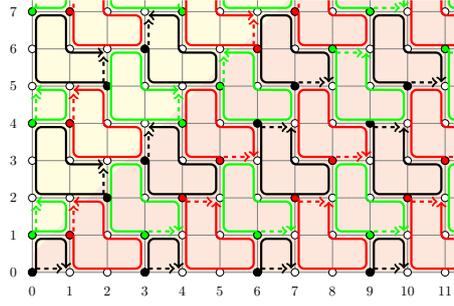
► **Theorem 13.** *The satisfiability problem for $\mathcal{FL}^2\text{3T}$, the two-variable fluted fragment with three transitive relations, is undecidable.*

Proof. Suppose the signature contains transitive relations b (black), g (green) and r (red), and additional unary predicates $e, e', f, l, c_{i,j}$ ($0 \leq i \leq 5, 0 \leq j \leq 2$) and $d_{i,j}$ ($0 \leq i \leq 2, 0 \leq j \leq 5$); we refer to the $c_{i,j}$'s and to the $d_{i,j}$'s as *colours*.

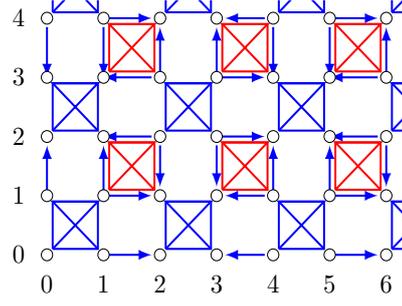
We reduce from the $\mathbb{N} \times \mathbb{N}$ tiling problem. We first write a formula φ_{grid} that captures several properties of the intended expansion of the $\mathbb{N} \times \mathbb{N}$ grid as shown in Fig. 2a. There the predicates $c_{i,j}$ and $d_{i,j}$ together define a partition of the universe as follows: an element (k, k') with $k' > k$ (i.e. in the yellow region, above the diagonal) satisfies $c_{i,j}$ with $i = k \bmod 6, j = k' \bmod 3$, and an element (k, k') with $k \geq k'$ (i.e. in the pink region, on or below the diagonal) satisfies $d_{i,j}$ with $i = k \bmod 3, j = k' \bmod 6$. Paths of the same transitive relation have length at most 7 and follow one of four designated patterns. Remaining unary predicates mark the following elements: l – left column, f – bottom row, e – main diagonal, and e' – elements with coordinates $(k, k+1)$.

The formula φ_{grid} comprises a large number of conjuncts. We have organized these conjuncts into groups, each of which secures a particular property (or collection of properties) exhibited by its models. The first two properties are very simple:

18:10 The Fluted Fragment with Transitivity

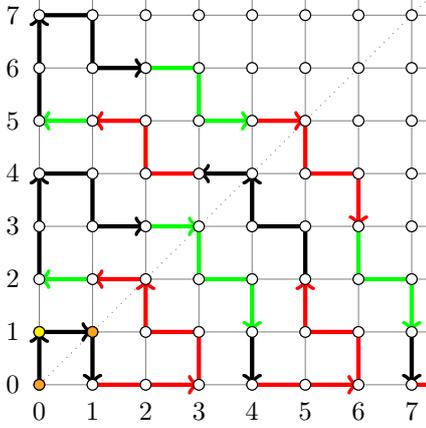


(a) Three transitive relations: b , g and r . Filled nodes depict the beginning of a transitive path of the same colour; dotted lines connect the first element with the last element on such path.

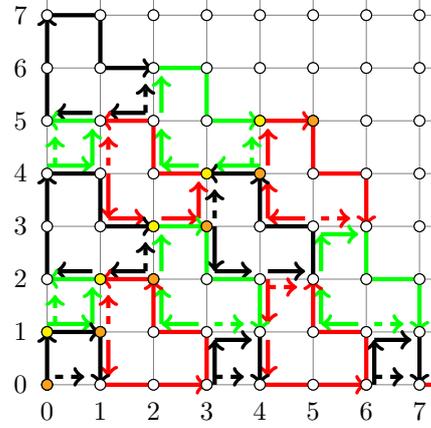


(b) Two transitive relations: b and r . Edges without arrows depict connections in both direction.

■ **Figure 2** Expansions of the $\mathbb{N} \times \mathbb{N}$ grid in the proofs of Theorem 13 (a) and Theorem 16 (b).



(a) Path starting with an initial element and generated by witnesses for conjuncts of group (3).



(b) Additional edges arising from conjuncts of group (4) (solid lines drawn inside grid cells) and transitivity (dashed lines). Nodes on the diagonals are marked orange (e) and yellow (e').

■ **Figure 3** Construction of the intended model of φ_{grid} in the proof of Theorem 13.

(1) There is an “initial” element satisfying $d_{00}(x) \wedge e(x) \wedge l(x) \wedge f(x)$.

(2) The predicates $c_{i,j}$ and $d_{i,j}$ together partition the universe.

The third property generates the path shown in Fig. 3a:

(3) Each element has a b - r - or g - successor as shown in the path shown in Fig. 3a, and satisfying the appropriate predicates $c_{i,j}$ or $d_{i,j}$. Specifically, if a node in this path has coordinates (x, y) with $y > x$, then it satisfies $c_{i,j}$ where $i = x \bmod 3$ and $j = y \bmod 6$; and when $y \leq x$, then the node satisfies $d_{i,j}$ where $i = x \bmod 3$ and $j = y \bmod 6$.

The conjuncts enforcing this property have the form

$$\forall x(\text{colour}(x) \wedge \text{diag}(x) \wedge \text{border}(x) \rightarrow \exists y(t(x, y) \wedge \text{colour}'(y))), \quad (3a)$$

where colour and colour' stand for one of the predicate letters $c_{i,j}$ or $d_{i,j}$, $\text{diag}(x)$ stands for one of the literals $e'(x)$, $\neg e'(x)$, $e(x)$, $\neg e(x)$ or \top (i.e. the logical constant true), $\text{border}(x)$ stands for one of the literals $l(x)$, $\neg l(x)$, $f(x)$, $\neg f(x)$ or \top , and t stands for one of the

transitive predicate letters b , r or g . The precise combinations of the literals and predicate letters in these conjuncts can be read from Fig. 3a (cf. [16] for details).

To connect all pairs of elements that are neighbours in the standard grid we require a fourth property, which we give in schematic form as follows:

- (4) *Certain* pairs of elements connected by *one* transitive relation are also connected by *another*, as indicated in Fig. 3b.

Here are some examples of the conjuncts enforcing this property:

$$\forall x(c_{01}(x) \rightarrow \forall y(b(x, y) \wedge (c_{11}(y) \vee d_{11}(y)) \rightarrow g(x, y))), \quad (4a)$$

$$\forall x(d_{11}(x) \rightarrow \forall y(b(x, y) \wedge d_{10}(y) \rightarrow r(x, y))), \quad (4b)$$

$$\forall x(d_{11}(x) \rightarrow \forall y(r(x, y) \wedge (c_{12}(y) \vee d_{12}(y)) \rightarrow g(x, y))), \quad (4c)$$

The role of these conjuncts can be explained referring to Fig. 3b. For example, employing (4b) for the element $(1, 1)$ in the intended model \mathfrak{G} , we get $\mathfrak{G} \models r((1, 1), (1, 0))$; hence by transitivity of r , also $\mathfrak{G} \models r((1, 1), (1, 2))$. This, applying (4c), implies $\mathfrak{G} \models g((1, 1), (1, 2))$. By (4a), we get $\mathfrak{G} \models g((0, 1), (1, 1))$ and, by transitivity of g , $\mathfrak{G} \models g((0, 1), (0, 2))$. The process is illustrated in Fig. 3b; when carried on along the zig-zag path, it constructs a grid-like structure.

These conjuncts depend on having available the predicates marking the borders and the diagonals. Specifically, we require the following property:

- (5) the predicates l , f , e and e' are distributed to mark the left-most column, the first row, the diagonal and the “super-diagonal” of the grid, as indicated above. To secure this property, we add to φ_{grid} several conjuncts, for instance:

$$\bigwedge_{0 \leq i \leq 2, 0 \leq j \leq 5} \forall x(d_{i,j}(x) \wedge \pm e(x) \rightarrow \forall y((b(x, y) \vee g(x, y) \vee r(x, y)) \wedge d_{i+1,j+1}(y) \rightarrow \pm e(y))), \quad (5a)$$

where $\pm e(x)$ denotes uniformly $e(x)$ or $\neg e(x)$. Similar conjuncts are added for the super-diagonal, left column and bottom row; and also for the connection with and between e and e' . The conjuncts ensuring properties (4) and (5) work in tandem. For instance, applying (5a) to $(1, 1)$ we get e is true at $(2, 2)$; then, following the zig-zag path and applying more conjuncts from the group (4), we get that $g((2, 2), (3, 3))$ holds, so the node $(3, 3)$ will be marked by e ; this will propagate along the main diagonal.

The structure \mathfrak{G} depicted in Fig. 2a is a model of φ_{grid} . In fact, φ_{grid} is an infinity axiom. To see this, let $\mathfrak{A} \models \varphi_{grid}$ and define an injective embedding ρ of the standard grid on $\mathbb{N} \times \mathbb{N}$ into \mathfrak{A} as follows. Let $next : \mathbb{N} \times \mathbb{N} \mapsto \mathbb{N} \times \mathbb{N}$ be the successor function defined on $\mathbb{N} \times \mathbb{N}$ as depicted by the zig-zag path in the left-hand picture of Fig. 3 starting at $(0, 0)$ (ignoring any colours). Denote $s_0 = (0, 0)$, $s_n = next(s_{n-1})$ and $S_n = \{s_0, \dots, s_n\}$. Let $a_0 \in A$ be an element such that $\mathfrak{A} \models d_{00}(a) \wedge e(a) \wedge l(a) \wedge f(a)$ that exists by condition (1). Define $\rho(s_0) = a_0$. Now, we proceed inductively: suppose $\rho(s_{n-1})$ has already been defined in step $n-1$ of the induction and $\rho(s_{n-1}) = a_{n-1}$. Let a_n be the witness of a_{n-1} for the appropriate conjunct from the group (3), i.e. where the unary literals for x agree with the unary literals satisfied by a_{n-1} in \mathfrak{A} . Define $\rho(s_n) = a_n$. Using induction one can prove that ρ is indeed injective: in the inductive step we assume that $\mathfrak{A} \upharpoonright \{a_0, \dots, a_{n-1}\}$ is isomorphic to $\mathfrak{G} \upharpoonright S_{n-1}$, and we show that $a_n \notin \{a_0, \dots, a_{n-1}\}$ and $\mathfrak{A} \upharpoonright \{a_0, \dots, a_n\}$ and $\mathfrak{G} \upharpoonright S_n$ are again isomorphic. In the proof one considers several cases depending on the 1-type realized by a_n . The formula φ_{grid} ensures that a_0, \dots, a_{18} are all distinct, and any eight consecutive elements of the sequence a_0, \dots, a_n are always distinct. Consider $a_{18} = \rho(4, 2)$ that requires a witness $b \in A$ for a conjunct from the group (3) such that $\mathfrak{A} \models r(a_{18}, b) \wedge d_{11}(b)$. Suppose, $b = a_2 = \rho(1, 1)$, since $\mathfrak{A} \models d_{11}(a_2)$. Then, by transitivity of g , $\mathfrak{A} \models g(a_{18}, a_{10})$, which is a contradiction with $\mathfrak{G} \models \neg g((4, 2), (1, 1))$. Other cases are similar and due to page limits have been omitted.

18:12 The Fluted Fragment with Transitivity

We are now ready to define the horizontal and vertical successors in models of φ_{grid} . In fact, instead of defining the horizontal grid successor h as one binary relation, we define two disjoint binary relations $rt(x, y)$ and $lt(x, y)$ such that rt and the inverse of lt together give the expected horizontal grid successor; they are defined respecting the “direction” of the transitive edges in the models. In the intended model $rt((x_1, y_1), (x_2, y_2))$ holds if $x_2 = x_1 + 1$, $y_2 = y_1$ and (x_1, y_1) and (x_2, y_2) are connected by b , g or r ; and for $lt((x_1, y_1), (x_2, y_2))$ to hold we require $x_2 = x_1 - 1$ instead of $x_2 = x_1 + 1$. We present the definition of $rt(x, y)$ in detail below¹

$$\begin{aligned} rt(x, y) := & (b(x, y) \vee g(x, y) \vee r(x, y)) \wedge \\ & (c_{01}(x) \wedge d_{11}(y)) \vee (c_{20}(x) \wedge d_{03}(y)) \vee (c_{42}(x) \wedge d_{25}(y)) \vee \\ & \bigvee_{(i,j) \notin \{(0,2), (1,2), (2,1), (3,1), (4,0), (5,0)\}} (c_{ij}(x) \wedge c_{i+1,j}(y)) \vee \bigvee_{(i,j) \notin \{(2,1), (1,3), (0,5)\}} (d_{ij}(x) \wedge d_{i+1,j}(y)) \end{aligned}$$

The relation rt connects elements that are connected by b , g or r and satisfy one of the possible combinations of colours: in the second line the combinations for crossing the diagonal are listed, in the third line the left disjunction describes combinations when both elements are located above the diagonal, and in the right disjunction – when both elements are located on and below the diagonal. The definition of $lt(x, y)$ complements that of rt . Analogously, we define relations up and dn that together define the vertical grid successor.

Now we are ready to write formulas that properly assign tiles to elements of the model. We do this with a formula φ_{tile} , which again features several conjuncts enforcing various properties of its models. Fortunately, the properties in question are much simpler this time:

(6) Each node encodes precisely one tile: $\forall x (\bigvee_{C \in \mathcal{C}} C(x) \wedge \bigwedge_{C \neq D} (\neg C(x) \vee \neg D(x)))$.

(7) Adjacent tiles respect \mathcal{C}_H . This is secured by the conjunct

$$\bigwedge_{C \in \mathcal{C}} \forall x (C(x) \rightarrow \forall y ((rt(x, y) \rightarrow \bigvee_{C': (C, C') \in \mathcal{C}_H} C'(y)) \wedge (lt(x, y) \rightarrow \bigvee_{C': (C', C) \in \mathcal{C}_H} C'(y))))$$

(8) Adjacent tiles respect \mathcal{C}_V (written as above using up and dn).

We remark that these latter formulas are not strictly fluted but can be rewritten as fluted using classical tautologies.

Finally, let $\eta_{\mathcal{C}}$ be the conjunction of φ_{grid} and φ_{tile} . We complete the proof showing

▷ **Claim 14.** $\eta_{\mathcal{C}}$ is satisfiable iff \mathcal{C} tiles $\mathbb{N} \times \mathbb{N}$.

Proof. (\Leftarrow) If \mathcal{C} tiles $\mathbb{N} \times \mathbb{N}$ then to show that $\eta_{\mathcal{C}}$ is satisfiable we can expand our intended model \mathfrak{G} for φ_{grid} assigning to every element of the grid a unique $C \in \mathcal{C}$ given by the tiling.

(\Rightarrow) Let $\mathfrak{A} \models \eta_{\mathcal{C}}$. Let ρ be the embedding of the standard $\mathbb{N} \times \mathbb{N}$ grid into \mathfrak{A} defined above. One can inductively show that ρ maps neighbours in the grid to elements connected by one of the relations lt, rt, up, dn as follows ($i, j \geq 0$):

$$\mathfrak{A} \models rt(\rho(i, j), \rho(i+1, j)) \dot{\vee} lt(\rho(i+1, j), \rho(i, j)) \wedge up(\rho(i, j), \rho(i, j+1)) \dot{\vee} dn(\rho(i, j+1), \rho(i, j)).$$

(Here, $\dot{\vee}$ is exclusive disjunction.) So, we can define a tiling of the standard grid assigning to every node (i, j) the unique tile C such that $\mathfrak{A} \models C(\rho(i, j))$. Conditions (7) and (8) together with the above observation ensure that this assignment satisfies the tiling conditions. ◁

¹ Addition in subscripts of the $c_{i,j}$'s is always understood modulo 6 in the first position, and modulo 3 in the second position, i.e. $c_{i+a,j+b}$ denotes $c_{(i+a) \bmod 6, (j+b) \bmod 3}$. Similarly, addition in subscripts of the $d_{i,j}$'s is understood modulo 3 in the first position, and modulo 6 in the second position.

We remark that the formula φ_{grid} in the proof of Theorem 13 is an axiom of infinity, hence the satisfiability and the finite satisfiability problems do not coincide. Moreover, all formulas used in the proof are either guarded or can easily be rewritten as guarded. Furthermore, in the proof it would suffice to assume that b , g and r are interpreted as equivalence relations. Hence, we can strengthen the above theorem as follows.

► **Corollary 15.** *The satisfiability problem for the intersection of the fluted fragment with the two-variable guarded fragment is undecidable in the presence of three transitive relations (or three equivalence relations).*

Now we improve the undecidability result to the case of $\mathcal{FL}^2\mathbf{2T}$ with equality.

► **Theorem 16.** *The (finite) satisfiability problem for the two-variable fluted fragment with equality is undecidable in the presence of two transitive relations.*

Proof. We write a formula φ_{grid} over a signature consisting of transitive relations b and r , and unary predicates $c_{i,j}$ ($0 \leq i, j \leq 3$). The formula φ_{grid} captures several properties of the intended expansion of the $\mathbb{Z} \times \mathbb{Z}$ grid as shown Fig. 2b:

- (1) there is an initial element: $\exists x.c_{00}(x)$.
- (2) the predicates $c_{i,j}$ partition the universe.
- (3) transitive paths do not connect distinct elements of the same colour: $\bigwedge_{0 \leq i, j \leq 3} \forall x(c_{ij}(x) \rightarrow \forall y((b(x, y) \vee r(x, y)) \wedge c_{ij}(y) \rightarrow x = y))$
- (4) each element belongs to a 4-element blue clique and to a 4-element red clique.
- (5) *certain* pairs of elements connected by r are also connected by b , and *certain* pairs of elements connected by b are also connected by r .

We have given property (5) only schematically, of course; its role is analogous to that of property (4) in the proof of Theorem 13. The remainder of the proof is similar to the one presented for Theorem 13 and is omitted due to space limits. We note that φ_{grid} has also finite models expanding a toroidal grid structure $\mathbb{Z}_{4m} \times \mathbb{Z}_{4m}$ ($m > 0$) obtained by identifying elements from columns 0 and $4m$ and from rows 0 and $4m$. Hence, the proof gives undecidability for both the satisfiability and the finite satisfiability problems. ◀

Again, the formulas used in the above proof are guarded or can be rewritten as guarded. Also it suffices to assume that r is an equivalence relation. Hence we get the following

► **Corollary 17.** *The (finite) satisfiability problem for the intersection of the fluted fragment with equality with the two-variable guarded fragment is undecidable in the presence of two transitive relations (or one transitive and one equivalence relation).*

5 Conclusions

In this paper, we considered the (m -variable) fluted fragment in the presence of different numbers of transitive relations. We showed that $\mathcal{FL}1\mathbf{T}$ has the finite model property, but $\mathcal{FL}^2\mathbf{3T}$ admits axioms of infinity and the satisfiability problem for $\mathcal{FL}^2\mathbf{3T}$ is undecidable. This contrasts with known results for other decidable fragments, in particular, \mathbf{FO}^2 , where the satisfiability and finite satisfiability problems are undecidable in the presence of two transitive relations, and where the *finite* satisfiability problem is decidable in the presence of one transitive relation. It is open whether the (finite) satisfiability problem for \mathcal{FL} in the presence of *two* transitive relations, t_1 and t_2 , is decidable. We point out that Lemma 11 in Section 3 could be generalized to normal form formulas from $\mathcal{FL}^{m+1}\mathbf{2T}$. Hence, the (finite) satisfiability problem for \mathcal{FL} in the presence of two transitive relations is decidable

if and only if the corresponding problem for \mathcal{FL}^2 with two transitive relations is decidable. Unfortunately neither the method of Sec. 3 (to show decidability) nor that of Sec. 4 (to show undecidability) appears to apply here. The barrier in the former case is that pairs of elements can be related by both t_1 and t_2 via divergent t_1 - and t_2 -chains, so that simple certificates of the kind employed for \mathcal{FL}^21T^u do not guarantee the existence of models. The barrier in the latter case is that the grid construction has to build models featuring transitive paths of *bounded* length, and this seems not to be achievable with just two transitive relations. Finally, we expect that the undecidability result for \mathcal{FL}^23T can be extended to get undecidability of the corresponding finite satisfiability problem.

References

- 1 H. Andréka, J. van Benthem, and I. Németi. Modal languages and bounded fragments of predicate logic. *Journal of Philosophical Logic*, 27(3):217–274, 1998.
- 2 F. Baader, D. Calvanese, D. L. McGuinness, D. Nardi, and P. F. Patel-Schneider, editors. *The Description Logic Handbook: Theory, Implementation, and Applications*. Cambridge University Press, 2003.
- 3 E. Börger, E. Grädel, and Y. Gurevich. *The Classical Decision Problem*. Springer, 1997.
- 4 D. Danielski and E. Kieroński. Finite Satisfiability of Unary Negation Fragment with Transitivity. In *Proceedings of MFCS 2019*, volume 138, pages 17:1–17:15, 2019.
- 5 E. Grädel, P. Kolaitis, and M. Vardi. On the decision problem for two-variable first-order logic. *Bulletin of Symbolic Logic*, 3(1):53–69, 1997.
- 6 A. Herzig. A new decidable fragment of first order logic. In *Abstracts of the 3rd Logical Biennial Summer School and Conference in Honour of S. C. Kleene*, June 1990.
- 7 Y. Kazakov and I. Pratt-Hartmann. A note on the complexity of the satisfiability problem for graded modal logic. In *Logic in Computer Science*, pages 407–416. IEEE, 2009.
- 8 E. Kieroński, J. Michaliszyn, I. Pratt-Hartmann, and L. Tendera. Two-variable first-order logic with equivalence closure. *SIAM Journal on Computing*, 43(3):1012–1063, 2014.
- 9 E. Kieroński and M. Otto. Small Substructures and Decidability Issues for First-Order Logic with Two Variables. *Journal of Symbolic Logic*, 77:729–765, 2012.
- 10 E. Kieroński and L. Tendera. On finite satisfiability of two-variable first-order logic with equivalence relations. In *Logic in Computer Science*, pages 123–132. IEEE, 2009.
- 11 R. Ladner. The computational complexity of provability in systems of modal propositional logic. *SIAM Journal on Computing*, 6:467–480, 1980.
- 12 A. Noah. Predicate-functors and the limits of decidability in logic. *Notre Dame Journal of Formal Logic*, 21(4):701–707, 1980.
- 13 I. Pratt-Hartmann. Finite satisfiability for two-variable, first-order logic with one transitive relation is decidable. *Mathematical Logic Quarterly*, 64(3):218–248, 2018.
- 14 I. Pratt-Hartmann, W. Szwaast, and L. Tendera. Quine’s fluted fragment is non-elementary. In *25th EACSL Annual Conference on Computer Science Logic, CSL*, volume 62 of *LIPICs*, pages 39:1–39:21. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 2016.
- 15 I. Pratt-Hartmann, W. Szwaast, and L. Tendera. The fluted fragment revisited. *Journal of Symbolic Logic*, 2019. (Forthcoming).
- 16 I. Pratt-Hartmann and L. Tendera. The fluted fragment with transitivity. *ArXiv*, 2019. extended version of MFCS’19 paper. URL: <https://arxiv.org/abs/1906.09131>.
- 17 W. C. Purdy. Fluted formulas and the limits of decidability. *Journal of Symbolic Logic*, 61(2):608–620, 1996.
- 18 W. C. Purdy. Complexity and nicety of fluted logic. *Studia Logica*, 71:177–198, 2002.
- 19 W. V. Quine. On the limits of decision. In *Proceedings of the 14th International Congress of Philosophy*, volume III, pages 57–62. University of Vienna, 1969.
- 20 W. V. Quine. The variable. In *The Ways of Paradox*, pages 272–282. Harvard University Press, revised and enlarged edition, 1976.

- 21 S. Schmitz. Complexity hierarchies beyond Elementary. *ACM Transactions on Computation Theory*, 8(1:3):1–36, 2016.
- 22 D. Scott. A decision method for validity of sentences in two variables. *Journal of Symbolic Logic*, 27:477, 1962.
- 23 W. Szwast and L. Tendera. On the satisfiability problem for fragments of the two-variable logic with one transitive relation. *Journal of Logic and Computation*, 2019. Forthcoming.
- 24 L. Tendera. Decidability frontier for fragments of first-order logic with transitivity. In *Proceedings of the 31st International Workshop on Description Logics co-located with 16th International Conference on Principles of Knowledge Representation and Reasoning (KR 2018)*, 2018. URL: <http://ceur-ws.org/Vol-2211/paper-02.pdf>.
- 25 M. Vardi. On the complexity of bounded-variable queries. In *Proceedings of the Fourteenth ACM SIGACT-SIGMOD-SIGART Symposium on Principles of Database Systems*, pages 266–276, 1995.